Characterization of a 100-kW Class Nested-Channel Hall Thruster

by

Scott James Hall

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Aerospace Engineering)
in the University of Michigan
2018

Doctoral Committee:

Professor Alec D. Gallimore, Co-Chair
Assistant Professor Benjamin A. Jorns, Co-Chair
Professor John E. Foster
Dr. Richard R. Hofer, Jet Propulsion Laboratory
Dr. Hani Kamhawi, NASA Glenn Research Center

It will not be we who reach Alpha Centauri and the other nearby stars. It will be a species very like us, but with more of our strengths and fewer of our weaknesses, a species returned to circumstances more like those for which it was originally evolved, more confident, farseeing, capable, and prudent—the sorts of beings we would want to represent us in a Universe that, for all we know, is filled with species much older, much more powerful, and very different.

The vast distances that separate the stars are providential. Beings and worlds are quarantined from one another. The quarantine is lifted only for those with sufficient self-knowledge and judgment to have safely traveled from star to star.

–Carl Sagan, *Pale Blue Dot*

50 years of high-power electric propulsion at NASA Glenn Research Center: a 1.5-meter diameter mercury ion thruster tested to 200 kW in 1967 (which has since been modified into a display unit) on the left and the X3 NHT, which was tested to 102 kW in 2017, on the right. Image courtesy NASA.

Scott James Hall

sjhall@umich.edu

ORCID iD: 0000-0002-8899-2993

For Katherine, my test buddy for life.

ACKNOWLEDGMENTS

Theodore von Karman very poignantly said, "It is interesting to note how again and again a human relationship can create history," a sentiment that is very true of the work described here (arguments about the level of "history" notwithstanding). An incredible number of people helped me along on this journey, and I owe them each a debt of gratitude.

I must first thank my advisor, Dean Alec Gallimore, who took in a student whose electric propulsion knowledge extended to a single magazine article on VASIMR and let him work on arguably the most visible project in the lab. Alec provided me with guidance and with freedom, and with a level-headed approach to all things. His vision of the big picture is something I admired and attempted to learn from my entire time in graduate school. His unwavering support of me and this project are the reasons that we accomplished so much.

Dr. Hani Kamhawi at NASA Glenn Research Center deserves special recognition as well. He has provided an incredible amount of support for me as my fellowship mentor, but was working hard to get the X3 into VF-5 well before I even joined the project. I have learned so much from him, not only about Hall thrusters but about how research gets done. I owe not only my success in graduate school but any future success in my career to the lessons he taught me.

I made the choice to bring my committee together early in hopes that I could maximize the time I had to benefit from their contributions, and that proved to be true even more so than I expected. Professor John Foster and Dr. Rich Hofer provided excellent hardware and experimental advice. Though Professor Ben Jorns only recently joined the scrum, his contributions are likely the largest ones to this document on a sentence-by-sentence basis. In the short time we worked together he improved my writing and taught me to be a better researcher. Thank you to each of you for your support and help. Any errors that remain in this document are mine alone.

I received financial support from a number of sources during my graduate career. I was humbled and honored to serve as a NASA Space Technology Research Fellow (supported under grant number NNX14AL67H and, as mentioned above, the fearless guidance of Dr. Hani Kamhawi). The NSTRF program is one of a kind and my work benefited greatly from their support. My National Science Foundation Graduate Research Fellowship supported me early in my studies, and provided the funding I needed to start at PEPL. And finally, the Michigan Institute for Plasma Science and Engineering provided me with a graduate fellowship, which I was able to use to support our testing at GRC.

A number of other people contributed greatly to my success. Dan Goebel at JPL consistently provided me with world-class cathodes, even as I consistently found new ways to break them. Peter Peterson not only put his sweat into the initial development of the X3, but was by my side the entire time we were testing at GRC (well, figuratively, since he was mostly up

on the deck running Big Blue and manning the camera). Pete can read thrusters in a way that I can only hope to someday match. There are those that put work into the X3 development early, before I joined the project: Jay Polk, Dan Brown, James Haas, Christopher Griffiths, and of course R.E.F., on whose shoulders I stood throughout this work. Thank you all.

At the University of Michigan, I was assisted by: Aaron Borgman, Chris Chartier, Cindy Enoch, Tom Griffin, Terry Larrow, Dave McLean, Kristin Parrish, Denise Phelps, Julie Power, J.P. Sheehan, Tim Smith, and Eric Viges. Special thanks to Alec's assistants over the years: Colleen Root, Chanda Doxie, and most importantly Kathleen Grimes, who treated all of us like her own children. Thank you for making it so easy to work with Alec even as he became the Dean. Professor Mark Kushner taught the two best courses I had the good fortune to take at the University of Michigan. I will carry my MIPSE mug with pride. Bill Cleveland and the team at Precision Manufacturing Services in Ann Arbor not only machined most of the X3, but a lengthy list of my own parts including two thrust stands and uncountable other pieces, and they did so with good cheer and great quality.

At NASA Glenn Research Center, I was fortunate to work with an incredibly talented and dedicated team, including: Jim Gilland, Tom Haag, Jon Mackey, Wensheng Huang, Luis Piñero, Matt Baird, Taylor Seablom, Chad Joppeck, Richard Senyitko, Jim Zakany, Nick Lalli, Jim Szologowski, Josh Gibson, Dave Yendriga, Larry Hambly, and Grandpa George Jacynycz. Kevin Blake in particular saved me from myself multiple times. Eric Pencil was the first to welcome me into the GRC family, and worked tirelessly to keep us all in line.

A number of past PEPL students influenced me as well. Mike Sekerak and Ray Liang were a constant source of insight, information, and inspiration. Ray, I never ended up needing to use that hammer, so I passed it on to the next crop of students. Adam Shabshelowitz and I did not overlap much, but his encouragement to start writing early and to use Grad Tools to track my progress are reflected in the quality of this document. Chris Bellant constantly reminded me that the most important thing is to love what you do, and I wish I had the opportunity to work longer with both him and Ingrid Goglio. I never stopped learning from Tom Liu while working with him during my early days in the lab.

My lab mates at PEPL were a huge component of my success. Thank you in particular to Frans Ebersohn, Sarah Cusson, Ethan Dale, and Marcel Georgin, for being hugely supportive and a joy to work alongside. You guys are the reason my years at PEPL have been the best of my life. Good luck to each of you as you finish up or, in Frans's case, continue to make real money. Josh Woods contributed to my diagnostics at GRC and sat through many hours of thruster testing. I hope I was able to teach you something. To the others in the lab, including Tim Collard, Zach Brown, Shad Hepner, Ben Wachs, Matt Byrne, and Angela Kimber, good luck as you finish (and in some cases start) your research. Especially best wishes to those starting new projects. Keep the spirit of the lab vital. I also benefited greatly from working and studying alongside Horatiu Dragnea and Astrid Raisanen from NGPDL. Horatiu and I never got to simultaneously characterize the X3 like we had dreamed of as pre-candidates, but he and Sarah served as my NHT cohort throughout my time in grad school.

I have had the strong support of a number of people in my life outside of the lab. The Turnbulls (John Sr., John Jr., Tyler, and now Tucker) have been a second family for me long before I started graduate school. They showed me the benefits of taking care of equipment, and of hard work, persistence, and loving what you do. My actual family was also

hugely instrumental in getting me here: my parents, for creating a household where learning and reading were expected and encouraged; Grandpa Freed, for teaching me patience and craftsmanship; my Grandma Freed and Grandparents Hall, all of whom made sacrifices in their lives to get me here; and my sister Megan, who was racing me to a post-graduate degree. I win, but it was close. Thank you to Megan as well for settling once and for all the "Proper possessive form of Mars" debate in the Hall household and answering my numerous other grammar and syntax questions.

And finally, of course, my wife Katherine, whose constant, unwavering, determined support has been what pushed me to do this, and to do it right. My favorite test buddy, my closest confidante, my editor in chief, my Captain Grafoil. You have supported me through late nights and early mornings, through chamber pumpdowns and thruster spark events. You were the first to volunteer to stand by the breakout box, the loudest vote of "Yes, go for 100 kW!" You even tolerate my more, er, bookish tendencies ("Look, honey, it's an *entire book* about *electrons*!"). This is as much yours as it is mine. From the bottom of my heart, thank you.

SJH
Ann Arbor, MI
2017

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF APPENDICES

NOMENCLATURE

Roman

a	fit parameter
A, B, C, I	regression coefficients (mass model)
A_{ingest}	ingestion area
A_{probe}	probe collection area
$\mathcal{A}$	area (mass model)
b	fit parameter
B	magnetic field strength (typically on channel centerline)
B_{nom}	nominal magnetic field strength
$B_{sat,design}$	design saturation magnetic field strength
c	fit parameter
C	channel combination
d	channel diameter (mass model); fit parameter
f	unitless weighting factor (mass model); fit result
f_b	breathing mode frequency
I_b	beam current
I_c	collected probe current
I_d	discharge current
I_e	electron current
I_{rms}	root mean square current
I_{sp}	specific impulse
$I_{sp,a}$	anode specific impulse
$\mathcal{I}$	current (mass model)
j	current density; CEX-corrected WFS current density
j_0	collected WFS current density
j_{ffp}	fixed Faraday probe current density
$\frac{j}{j_0}$	WFS CEX attenuation factor
j_{ref}	reference thruster current density
L	channel length

L_i	length of ionization region
m	mode number
m_f	final mass
m_0	initial mass
m_{prop}	propellant atomic mass
M	mass (mass model)
$\dot{m}$	mass flow rate
$\dot{m}_a$	anode mass flow rate
$\dot{m}_{a,t}$	total anode mass flow rate (summed across all channels)
$\dot{m}_b$	ion beam mass flow rate
$\dot{m}_c$	cathode mass flow rate
$\dot{m}_{ch}$	actual channel mass flow rate
$\dot{m}_{c,t}$	total cathode mass flow rate
$\dot{m}_{facility}$	facility ingestion mass flow rate
$\dot{m}_{ingest}$	ingestion mass flow rate
$\dot{m}_{max,design}$	designed maximum mass flow rate
$\dot{m}_{max,test}$	tested maximum mass flow rate
$\dot{m}_{min,design}$	designed minimum mass flow rate
$\dot{m}_{min,test}$	tested minimum mass flow rate
$\dot{m}_{pc}$	power-corrected mass flow rate
$\|\|\dot{m}\|\|$	normalized mass flow rate
n	number density
n_0	background neutral number density
N	number of thrusters (mass model)
p	pressure
p_b	facility background pressure
$\overline{p}$	average pressure
$\|\|p\|\|$	normalized pressure
$\|\|p\|\|_{max}$	maximum normalized pressure
$\|\|p\|\|_{min}$	minimum normalized pressure
pz	WFS charge exchange effect parameter
$\overline{pz}$	average WFS charge exchange effect parameter
P	power
P_d	discharge power
P_{elec}	electrical power
P_{jet}	jet power

r_{sun}	distance from the Sun
t	time
T	thrust
T_e	electron temperature
$T_{expected}$	expected thrust
T_g	gas temperature
T/P	thrust to power ratio
u_e	exhaust velocity
u_i	ion velocity
u_n	neutral velocity
V	voltage
V_a	acceleration voltage
V_{avg}	average voltage
V_b	probe bias
V_{cg}	cathode to ground voltage
$V_{coupling}$	thruster/cathode coupling voltage
V_d	discharge voltage
V_f	floating potential
V_{mp}	most probable voltage
V_p	plasma potential
w	channel width
x	independent fit variable
z	distance from thruster exit plane
Z	charge state

Greek

α_m	portion of mass utilization efficiency related to charge state information
β	ratio of discharge current to anode mass flow rate
γ_{ij}	cross-channel ingestion geometry factor
ΔV	velocity increment
ζ	ion species fraction
η	efficiency
η_a	anode efficiency
$\eta_{a,corr}$	corrected anode efficiency

$\eta_{a,lab}$	anode efficiency as measured in the laboratory
$\eta_{a,space}$	anode efficiency as experienced in space
η_b	current utilization efficiency
$\eta_{b,300M}$	current utilization efficiency of the NASA-300M
η_{bd}	divergence-weighted current utilization efficiency
$\eta_{bd,300M}$	divergence-weighted current utilization efficiency of the NASA-300M
η_d	divergence utilization efficiency
$\eta_{d,300M}$	divergence utilization efficiency of the NASA-300M
$\eta_{facility}$	efficiency due to facility background neutral ingestion
η_m	mass utilization efficiency
$\eta_{m,300M}$	mass utilization efficiency of the NASA-300M
η_{md}	divergence-weighted mass utilization efficiency
$\eta_{md,300M}$	divergence-weighted mass utilization efficiency of the NASA-300M
η_{NHT}	efficiency due to cross-channel ingestion in nested Hall thrusters
η_q	charge utilization efficiency
η_v	voltage utilization efficiency
$\overline{\eta}$	average efficiency
θ	beam divergence angle
σ	collisional cross section
σ_{max}	maximum deviation
σ_{min}	minimum deviation
$\overline{\sigma}$	average deviation
ϕ	angular location on thruster face
Φ	particle flux
Ω	ion species current fraction

Subscripts (Mass Model)

ac	active thruster
cab	cabling
gim	gimbal
i	channel number
NHT	Nested Hall Thruster
outer	outermost channel
rd	redundant thruster

s	structure
str	string
sys	system
$tank$	tankage
tg	thruster and gimbal
th	thruster
tot	total
Xe	xenon
XFS	xenon flow system

Constants

e	Elementary charge, 1.602×10^{-19} C
g	Earth's gravitational acceleration, 9.81 m/s^2
k_b	Boltzmann constant, 1.38×10^{-23} J/K
m_{kr}	Mass of krypton atom or ion, 1.39×10^{-25} kg
m_{xe}	Mass of xenon atom or ion, 2.18×10^{-25} kg
$\epsilon_{i,kr}$	First ionization energy of krypton, 14.0 eV
$\epsilon_{i,xe}$	First ionization energy of xenon, 12.1 eV

Abbreviations

AFRL	Air Force Research Laboratory
AR	Aerojet Rocketdyne
BN	Boron Nitride
BPT	Busek Primex Thruster
CEX	Charge EXchange
CFF	Cathode Flow Fraction
DSH	Deep Space Habitat
ELF	Electrodeless Lorentz force Thruster
EP	Electric Propulsion
FFP	Fixed Faraday Probe
GRC	Glenn Research Center
HERMeS	Hall Effect Rocket with Magnetic Shielding

HMT	High-Mass Thruster
HSC	High Speed Camera
HSDC	High Speed Discharge Current
I	Inner channel condition
IEDF	Ion Energy Distribution Function
IM	Inner + Middle channel condition
IMO	Inner + Middle + Outer channel condition
IO	Inner + Outer channel condition
JPL	Jet Propulsion Laboratory
LEO	Low Earth Orbit
LIF	Laser-Induced Fluorescence
LMT	Low-Mass Thruster
LP	Langmuir Probe
LVDT	Linear Variable Differential Transducer
LVTF	Large Vacuum Test Facility
M	Middle channel condition
MFC	Mass Flow Controller
MO	Middle + Outer channel condition
MPD	Magnetoplasmadynamic thruster
M1	Magnet 1
M2	Magnet 2
M3	Magnet 3
M4	Magnet 4
M5	Magnet 5
M6	Magnet 6
NASA	National Aeronautics and Space Administration
NEA	Near-Earth Asteroid
NEXT	NASA's Evolutionary Xenon Thruster
NextSTEP	Next Space Technologies for Exploration Partnerships
NHT	Nested Hall Thruster
NSTAR	NASA Solar electric propulsion Technology Applications Readiness ion engine
O	Outer channel condition
P2P	Peak-to-Peak
PDF	Probability Distribution Function
PEPL	Plasmadynamics and Electric Propulsion Laboratory
PID	Proportional Integral Derivative

PIT	Pulsed Inductive Thruster
PMA	Propellant Management Assembly
PPT	Pulsed Plasma Thruster
PPU	Power Processing Unit
PSD	Power Spectral Density
RMS	Root Mean Square
RPA	Retarding Potential Analyzer
SEP	Solar Electric Propulsion
SERT	Space Electric Rocket Test
SEV	Space Exploration Vehicle
SHT	Single-channel Hall Thruster
SNAP	Systems for Nuclear Auxiliary Power
SPT	Stationary Plasma Thruster
TAL	Thruster with Anode Layer
TCFF	Total Cathode Flow Fraction
TDU	Technology Demonstration Unit
UM	University of Michigan
VASIMR	VAriable Specific Impulse Magnetoplasma Rocket
VF5	Vacuum Facility 5
WFS	Wien Filter Spectrometer
XFS	Xenon Flow System
XR	Xenon Rocket

ABSTRACT

The next generation of electric propulsion consists of systems in excess of 300 kW of power. These systems enable a wide variety of missions, including crewed missions to near-Earth asteroids and Mars. Hall thrusters are a particularly attractive technology for these missions, but development and demonstration of 100-kW class devices has been limited to date. Here, a 100-kW class three-channel nested Hall thruster called the X3 was operated up to 102 kW total discharge power. The three channels of the X3 can be operated in any combination, providing seven unique configurations and a total throttling envelope of 2–200 kW. Previous testing of the X3 was limited to 30 kW and showed that it was not providing state-of-the-art performance.

Two low-power test campaigns were completed at the University of Michigan which identified potential mechanisms for this under-performance. Improvements to the thruster were made before a high-power performance characterization at NASA Glenn Research Center. There, the X3 was operated on xenon propellant from 5–102 kW total power. The thruster demonstrated stable operation in all seven channel combinations at discharge voltages from 300 V to 500 V and three different current densities. All seven channel combinations demonstrated similar performance at a given discharge voltage and current density. The largest thrust recorded was 5.4 N, and total efficiency and specific impulse ranged from 0.54 to 0.67 and 1800 seconds to 2650 seconds, respectively. For all channel combinations, total efficiency values greater than 0.63 were demonstrated.

In addition to the performance measurements, a suite of plasma diagnostics and a high-speed camera were used to study the operation of the thruster in greater detail. The probe results are compared against those in the literature and show that the X3, even in multi-channel operation, is producing similar charge, mass, current, and voltage utilization efficiencies as the NASA-300M 20-kW Hall thruster, a state of the art high-power thruster designed with similar design principles as the X3. High-speed camera analysis identified that the X3 operated in a similar mode of discharge current oscillations at nearly all conditions tested. This oscillatory behavior was characterized by the entire discharge channel oscillating as a whole (a so-called breathing mode oscillation) in a random, non-sinusoidal manner. Analysis indicated that when channels were operating together their oscillations did not correlate with one another either in sync or with a phase delay. Oscillatory behavior was also confirmed with high-speed discharge current analysis. Additionally, a preliminary calculation of cross-channel ingestion and its effect on thruster efficiency was made.

This work represents the highest total power (102 kW), thrust (5.4 N), and discharge current (247 A) demonstrated by a Hall thruster to date, improvements of 6%, 64%, and 119% respectively over previous values. These results are discussed in the context of continued high-power Hall thruster development and future mission applications.

CHAPTER 1

Introduction

"You'd be amazed at how fast you can get going with a tiny acceleration over a long time."

– *The Martian* by Andy Weir [7]

1.1 Problem Statement

Electric propulsion systems at power levels in excess of 300 kW enable a wide variety of space missions. These include missions to bodies with scientific and economic interest such as near-Earth asteroids and Mars. Hall thrusters have demonstrated performance and lifetime capability that is ideal for these types of missions [2, 8]. However, to date the highest-power Hall thruster to have flown is 5 kW [9]. Hall thrusters have been operated in excess of 50 kW previously [10], but the scaling techniques used in their development cause thrusters in the 100-kW class and above to become challengingly large.

An alternative design known as nesting places multiple discharge channels concentrically around a center-mounted cathode [1]. A three-channel, 100-kW class nested Hall thruster known as the X3 was developed using this technique [4]. This thruster was designed for operation up to 250 kW of discharge power at specific impulses ranging from 1800–3500 seconds on xenon propellant and up to 5000 seconds on krypton, all with state of the art total efficiencies in excess of 60%. However, early testing of the thruster was limited to 30

kW at steady-state and demonstrated performance lower than expected.

1.2 Objectives and Contributions

The primary objective of this work was to characterize the performance of the X3 across its throttling range. This consisted of two low-power tests at the University of Michigan, a full thruster disassembly and inspection to implement improvements based on the results of these tests, and then a full characterization up to 102 kW total power at NASA Glenn Research Center. This characterization included performance and detailed plasma plume measurements as well as a study of discharge current oscillations using a high-speed camera.

Contributions of this work include:

1. Expansion of a mass and cost model to capture the benefits of nested Hall thrusters. These results demonstrate the need for high-power electric propulsion and that nested Hall thrusters provide performance and efficiency well-matched to many missions. Additionally, the system-level mass and cost savings of nested Hall thrusters are identified.

2. Development of a new inverted-pendulum thrust stand designed specifically for the X3. This stand is capable of supporting thrusters in excess of 250 kg and measuring thrusts in excess of 8 N with a measurement uncertainty of approximately 1%. The stand features an optical displacement sensor for improved resolution. Versions of the stand were built at both the University of Michigan and NASA Glenn Research Center.

3. Demonstration of new capabilities of Hall thrusters in three metrics: total power, thrust, and discharge current. Previously, the NASA-457Mv1 demonstrated a maximum total power of 96 kW, 112 A discharge current, and 3.3 N thrust [11]. In this work, we operated the X3 to 102 kW and 247 A discharge current and measured a maximum thrust of 5.4 N. These new accomplishments demonstrate high-power Hall thruster operation at specific impulses more relevant for crewed missions.

4. The first study of nested Hall thruster operation up to 100 kW discharge power. Lingering questions about the viability of the technology, including whether the thruster would have thermal issues or incur cathode coupling problems, as well as whether it would produce comparable performance to the state of the art, have been answered by this work.

5. The first plasma plume measurements of a three-channel nested Hall thruster and of a Hall thruster operating in excess of 50 kW. Previous nested Hall thruster plume measurements did not exceed 6 kW [1] and were for only two channels operating together. These results include plasma measurements for up to three channels operating together at power levels up to 80 kW. Furthermore, the Outer channel of the X3 is the largest Hall thruster discharge channel to ever be built. The plasma plume measurements from it represent validation of the scaling techniques used to develop the thruster. These measurements provide insight not only into the operation of each channel of the X3 but also how the channels operate together.

6. The first detailed study of the oscillations in a 100-kW class Hall thruster via high-speed camera analysis. The highest-power Hall thruster to date to be studied by high-speed camera techniques was 20 kW [12]. This work provides measurements on the Middle and Outer channels of the X3, which on their own represent much larger discharges than those previously studied. Additionally, significant insight is provided into how the oscillation character of the X3 varies for a given channel between single- and multi-channel operation. These measurements provide valuable information for future testing of the X3 and future development of high-power Hall thrusters.

7. The first calculations of the propellant ingestion occurring between channels of a three-channel nested Hall thruster and the impact this ingestion has on thruster performance. Though further work is needed, these results suggest that the ingestion of neutral propellant between channels of nested Hall thrusters contributes 1–5% to the overall

thruster efficiency.

1.3 Organization

Chapter 2 provides background information on electric propulsion, as well as a description and history of the Hall thruster. It also provides a description of the X3 thruster and details of the early work performed on it. Chapter 3 motivates in great detail the need for high-power electric propulsion, including a survey of potential missions and results from a system mass and cost model that we expanded to include nested Hall thrusters. Chapter 4 describes the various experimental apparatus used in this work, which spanned three tests at two state of the art testing facilities. Data analysis techniques for the plasma diagnostics are also included.

Chapter 5 details the low-power test campaigns of this work, which gathered valuable information in preparation for high-power testing. We performed a thorough thruster inspection and made a series of improvements to its design and construction before the high-power test, details of which are provided in Chapter 6. The high-power thruster characterization is split into two chapters. Chapter 7 reports the performance of the thruster, including thrust, specific impulse, and efficiency values; it also compares operation in single-channel and multi-channel configurations. Chapter 8 then reports results from the suite of plasma diagnostics as well as the high-speed camera analysis. From the plasma diagnostics, a series of phenomenological efficiencies are calculated for the X3 in different operating conditions. Results from both techniques are compared to relevant results in the literature, including a comparison of the efficiency results to the 20-kW class NASA-300M thruster. Chapter 9 summarizes the findings from these experiments and provides recommendations for future work.

Appendix A reports a study that we performed of X3 operation at low total cathode flow fraction, including thruster performance and plasma plume measurements. Appendices B

and C present the detailed telemetry and performance results for the low- and high-power tests, respectively.

CHAPTER 2

Background

"For I have discovered that [the professional engineer in the rocket business] is frequently abysmally ignorant of the history of his own profession, and, unless forcibly restrained, is almost certain to do something which, as we learned fifteen years ago, is not only stupid but likely to result in catastrophe."

– John D. Clark [13]

2.1 Introduction

This chapter provides background information on electric propulsion, as well as details of Hall thrusters and their performance metrics. Section 2.2 describes electric propulsion, including details of its origin, governing principles, and various categories of devices. We then focus on Hall thrusters in Section 2.3, where we provide basic details of those devices as well as a brief history and description of techniques used to scale them to high power. Section 2.4 offers a history and description of the particular high-power Hall thruster used in this investigation, the X3. We then detail the performance metrics typically used to evaluate Hall thrusters (which we use in this work to evaluate the X3) in Section 2.5. There, we also describe the modifications to a typical Hall thruster efficiency analysis necessary to accommodate the diagnostic setup used on the X3, as well as develop a method for accounting for propellant ingestion between the multiple channels of a nested Hall thruster. Finally,

Section 2.6 describes the expected performance of the X3 thruster based on theory and historical values, as well as details of the initial testing of the X3 undertaken prior to this work.

2.2 Electric Propulsion

2.2.1 Province

The basic mathematics of space travel concisely motivates the need for propulsion systems with exhaust velocities greater than what is attainable with traditional chemical means. The fraction of a spacecraft's mass that is not used for propulsion (the dry mass, m_f) can be related to the exhaust velocity of the propulsion system by the following expression:

$$\frac{m_f}{m_0} = \exp\left(-\frac{\Delta V}{u_e}\right) \tag{2.1}$$

This expression, called the Rocket Equation, was first developed by Russian rocket pioneer Konstantin Tsiolkovsky in 1895 [14]. Typically the fraction m_f/m_0 is referred to as the payload fraction, which defines the fraction of the craft's mass at launch that is delivered to the final destination. This includes useful payload (such as orbiters, landers, rovers, supplies, and human crew) as well as non-expendable infrastructure such as the spacecraft chassis, propellant tanks, and the propulsion system itself. A driving goal of space mission design is maximizing payload fraction [15] (and thus minimizing propellant mass fraction, which is one minus the payload fraction), and the Rocket Equation explains how to go about this. It is clear that to maximize payload fraction, the fraction $\Delta V/u_e$ within the exponential of Equation 2.1 must go to zero. The term ΔV is the change in velocity required to enact a given maneuver (e.g., to change orbits or transit from Earth to Mars) and thus defines the mission [15, 16] (and is fixed). Values for ΔV are typically 1–10 km/s for interplanetary travel. The exhaust velocity of the propulsion system, u_e, is thus the free variable available to

the propulsion engineer, and it is clear from Equation 2.1 that maximizing exhaust velocity is required to maximize payload fraction.

Propulsion systems are often compared using a parameter called the specific impulse:

$$I_{sp} = \frac{u_e}{g} \tag{2.2}$$

which describes how effectively propellant is converted into thrust. By dividing by the gravitational acceleration constant g, specific impulse is expressed in units of seconds and conveniently provides no conflict between the imperial and metric unit systems. Additionally, if the metric system is used, u_e is conveniently approximately ten times larger than I_{sp}.

The value of u_e that a particular propulsion system can obtain is related to the mechanism by which the propellant acceleration occurs. Traditional chemical rockets, which release energy stored in the bonds of their propellant through chemical reaction and then accelerate these products through a converging-diverging nozzle, are limited to specific impulses of a few hundred seconds, regardless of the propellants used or the particular nozzle design [17]. For many missions, such as those with large ΔV values, chemical propulsion methods are ill-suited. Their low specific impulses require large fractions of a spacecraft's mass to be devoted to propellant. Yet the physics of chemical propulsion simply preclude substantial increases in specific impulse: an alternative thrust generation mechanism is needed.

Electric propulsion (EP), which has been speculated about since 1906 [14, 18], given treatment in texts since 1929 [19], and flown in space in various forms since the Space Electric Rocket Test 1 spacecraft fired one of its ion thrusters for 31 minutes in 1964 [20], provides an alternative means of thrust generation. The classic text on EP, *The Physics of Electric Propulsion* by Professor Robert Jahn [20], defines EP as:

> **The acceleration of gases for propulsion by electrical heating and/or by electric and magnetic body forces.**

In other words, in electric thrusters the thrust-generation mechanism is decoupled from the

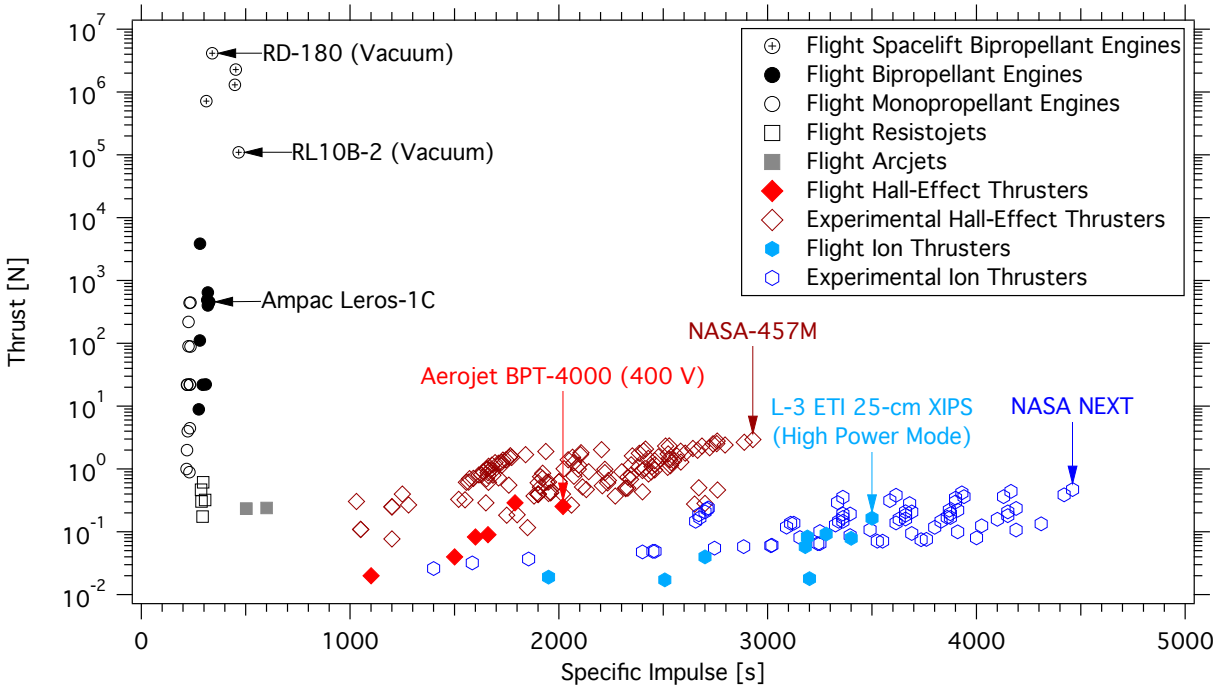

Figure 2.1: A plot of the thrust-versus-specific impulse trade space inhabited by various types of propulsion. Figure is reproduced from Reference [1].

energy in the chemical bonds of the propellant. This frees EP from the exhaust velocity limitations of chemical propulsion. It also typically frees EP systems from relying on the complex, hazardous, highly reactive, and difficult to store propellants often used in the chemical propulsion industry [13, 17, 20].

The large number of different types of electric thrusters are detailed below. In all cases they provide significant increases in thrust efficiency over chemical means. Specific impulses range from comparable to chemical thrusters (resistojets) to upwards of 8000 seconds (ion thrusters). Figure 2.1, which is reproduced from Reference [1], presents a large sample of measured performance data from various electric thrusters and chemical rocket engines. The figure shows that chemical engines can produce a wide range of thrusts but in all cases are limited to less than 500 seconds specific impulse. Resistojets are shown to provide essentially the same values of specific impulse as chemical engines with much smaller thrust values; arcjets produce specific impulses greater than chemical means with roughly the same

thrust values as resisotjets. Farther to the right along the abscissa are Hall thrusters and then ion thrusters, which in general provide much higher specific impulse values but with much lower thrusts than those provided by chemical propulsion. The significant ongoing development efforts to increase both specific impulse and thrust values for both technologies are displayed. Due to the lengthy and involved process of flight qualification [9, 21], flight thrusters typically lag behind experimental devices.

As the figure shows, EP is typically characterized by low thrust values, ranging from a few μN to 150 N for high-power arcjets [20]. These low total thrusts, as well as physical issues with operating outside of a vacuum environment, relegate EP to use in space, free from most gravity effects. Electric propulsion will never replace large chemical engines for space-launch applications. However, electric thrusters are typically designed to operate for burns orders of magnitude longer than chemical rockets. Whereas typical chemical rockets operate for minutes to hours, electric thrusters can operate anywhere from hours to years [17]. For instance, ion thrusters have demonstrated operational lifetimes of over 45,000 hours in-space [22] and over 50,000 hours in ground testing [23]. Electric thrusters have the capability to impart the same total energy onto a spacecraft as chemical thrusters. This is done over much longer firing times, but with substantially less propellant mass.

There are a number of other benefits of using EP systems for certain mission profiles. It can be said that there are four goals for human mission design [24]:

1. Short transfer time
2. Moderate mass
3. Short mission duration
4. High flexibility

However, using chemical propulsion or even nuclear thermal propulsion [25] only provides for two of these goals: maintaining moderate mass, which can be accomplished for

these systems by using ballistic trajectories and specially-designed orbital maneuvers such as Hohmann transfer orbits [16, 26], and minimizing transfer time.

Launch opportunities for these special trajectories are both short and relatively rare in occurrence. For instance, Mars and Earth only pass through a given alignment every 2.13 years [16]. This makes chemically-propelled missions very inflexible with respect to launch windows, and also makes the prospect of sending a rescue mission or providing an effective abort scenario challenging. Another complication for an abort is that chemical propulsion operates impulsively due to its short firing time: it imparts the necessary energy over a very short period of time and quickly expels all propellant. "Turning around" is all but impossible. One only has to recall the Apollo 13 mission [27] to recognize the need for abort options for human missions. The long firing times of EP, as well as their demonstrated ease at many start/stop cycles [22, 28], means that electric thrusters are much more well-suited for abort scenarios. The long-duration firing also lends itself to being much more forgiving for launch windows; a study by Williams showed that above a certain mission length, the departure window was of essentially no matter [29]. (This is true for deep-space missions—those going to the outer planets or beyond—regardless of the propulsion mechanism because of the long trip times).

There are a number of other benefits to using EP, as captured by NASA's Human Exploration Framework Team [30] and re-stated well by Brophy [5]. EP systems allow for significant power to be available to the spacecraft during non-thrust periods. Due to the long lifetimes of many EP technologies, there is the possibility of a spacecraft with EP to be usable across multiple missions (for instance, a cargo tug that transits back and forth between Mars multiple times [31]). Finally, the report noted that EP systems provide more "graceful" system failure modes than chemical propulsion systems. That is, instead of exploding or leaking dangerous, reactive, toxic chemicals in the event of a failure, typical EP failure mechanisms simply cause the propulsion system to shut off.

2.2.2 Types

Electric propulsion is typically divided into three categories depending on the particular thrust generation mechanism [2, 20]:

1. **Electrothermal propulsion**, in which propellant is electrically heated and then accelerated through a traditional converging-diverging nozzle. Examples of electrothermal propulsion include resistojets and arcjets. Because the acceleration mechanism still relies on traditional nozzles, these types of propulsion typically provide modest increases in specific impulse over chemical means. Resistojets provide specific impulses in the 250–350 second range. Arcjets have demonstrated values as high as 1500 seconds, though they are typically limited to closer to 700 seconds with non-exotic propellants. However, their efficiencies, in the 30–50% range, are significantly improved over chemical rockets [2].

2. **Electrostatic propulsion**, in which propellant is ionized and accelerated by the direct application of electric body forces. Examples of electrostatic propulsion include ion thrusters and Hall thrusters (though it should be noted that some authors such as Jahn categorize Hall thrusters as electromagnetic [20]). Performance of electrostatic propulsion devices can be very high, with specific impulses upwards of 8000 seconds and efficiencies above 70% for ion thrusters.

3. **Electromagnetic propulsion**, in which propellant is ionized and accelerated by interactions between currents driven through the propellant and either internally-generated or externally-imposed magnetic fields. Electromagnetic propulsion devices can be operated steady-state or in pulsed modes. Examples of electromagnetic propulsion include magnetoplasmadyanmic (MPD) thrusters, pulsed inductive thrusters (PITs), pulsed plasma thrusters (PPTs), ambipolar thrusters, and helicon thrusters. Electromagnetic propulsion devices can provide higher thrust densities than electrostatic devices, and they are not space-charge limited like ion thrusters. Though many technologies

are still in the early development phase, efficiencies vary from approximately 10% for solid-propellant pulsed plasma thrusters to as much as 50% for magnetoplasmadynamic thrusters. Specific impulses also cover a range from 600–2000 seconds for PPTs to values as high as 5000 seconds for MPDs.

2.3 Hall Thrusters

The Hall-effect thruster or Hall thruster appeared in the United States scientific literature in 1962 [32–34] in an early form called the Hall Current Accelerator. The concept first appeared in Russian literature in 1968 [35] and was further developed and first flown by Russian scientists during the 1960s and 1970s. The devices use crossed electric and magnetic fields to confine electrons while ejecting ions to produce thrust. Hall thruster operation is described below, followed by a short history of the device and a description of the scaling laws used to increase thruster power level.

2.3.1 Basics of Operation

There are two major variants of Hall thrusters: the "stationary plasma thruster" or SPT (sometimes referred to as a magnetic layer thruster), and the "thruster with anode layer" or TAL. The naming derives from translations of Russian work [36]. The fact that an SPT is steady-state and not pulsed seems to be where the "stationary" comes from; the compressed nature of the acceleration region of the TAL, very close to the anode, is the "anode layer". This work concerns magnetic-layer or SPT-style thrusters only, so for simplicity all references to "Hall thrusters" herein will be to this type of device unless clearly stated as otherwise.

A typical state-of-the-art Hall thruster features an annular discharge channel lined with dielectric material such as boron nitride, a combination anode/propellant distributor in the upstream end of the discharge channel, and a cathode mounted either in the center of the thruster ("centrally-mounted") or on the outer diameter of the thruster ("externally-

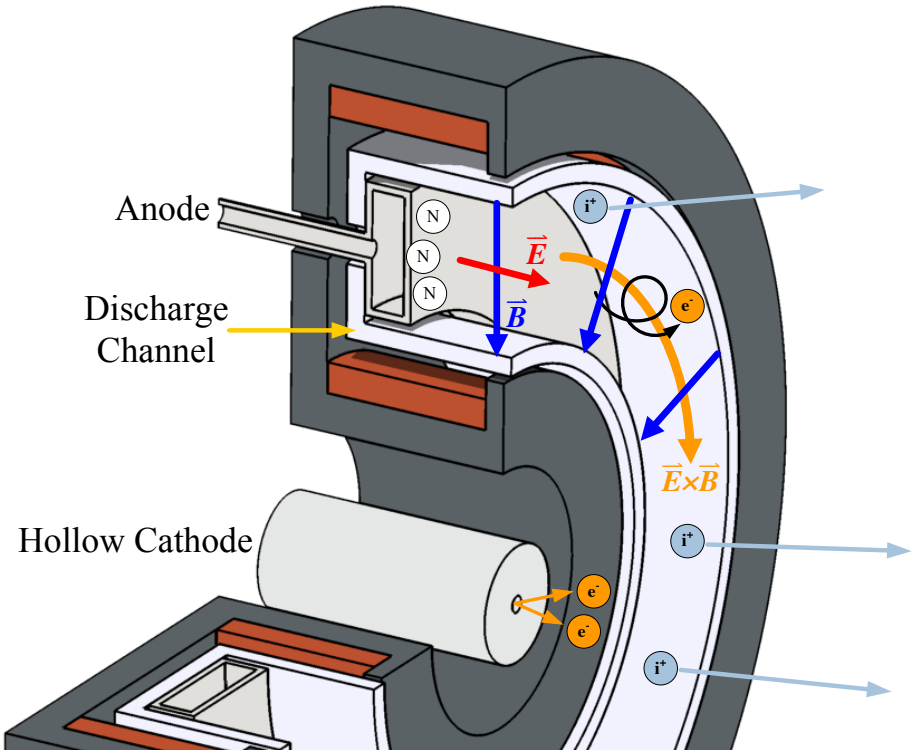

Figure 2.2: A schematic of a typical state of the art single-channel Hall thruster. Figure is reproduced from Reference [1].

mounted") that provides electrons. A schematic of a typical state of the art single-channel Hall thruster reproduced from Reference [1] is presented in Figure 2.2. The anode is biased to a voltage between 150 V and 800 V to provide the device's axial electric field. The magnetic field is produced by electromagnets, each of which can consist of one or multiple coils of wire depending on the device. This magnetic field is designed to be mostly radial. A Hall thruster's magnetic field is typically in the low-hundreds of Gauss, designed to confine electrons while allowing ions to be freely accelerated. Thrusters are typically optimized for a nominal throttling point but can operate across a range of conditions, characterized by changes in propellant flow rate, anode voltage, and magnetic field. Typically, the magnetic field shape is held constant in a plasma-lens topology [3, 37, 38] and only the field strength is changed between conditions.

Of course, many variations of Hall thrusters exist. Typically, lower-power thrusters,

including models that are currently flight-qualified such as the BPT-4000/XR-5 [21, 39], feature externally-mounted cathodes, and it is only larger thrusters (typically > 5 kW or so) that feature centrally-mounted cathodes due to physical limitations. One active area of current research involves the application of a technique called "magnetic shielding" to Hall thrusters to extend their operational lifetimes [40,41]. This technique exploits the equipotentialization along magnetic field lines inside Hall thrusters to maintain low ion velocities near the walls, significantly reducing channel wall erosion rates. This erosion has traditionally been the leading life-limiting mechanism of Hall thrusters. Other variations exist, such as conducting walls in magnetically-shielded thrusters [42] and permanent magnets in place of electromagnets [43].

The following is a simplified overview of Hall thruster operation. More detail can be found in Reference [2]. Seed electrons are emitted thermionically by the cathode, which is heated and ignited before the thruster discharge is started. A portion of these electrons make their way to the discharge channel where they are trapped in an azimuthal $\vec{E}\times\vec{B}$ drift [44,45] by the axial electric field imposed by the anode and the radial magnetic field imposed by the magnetic circuit of the thruster. Neutral propellant is introduced through the biased anode into the rear of the discharge channel. Through electron impact ionization [46], the neutral propellant is ionized.

The newly-liberated electrons are trapped in the discharge channel in the same $\vec{E}\times\vec{B}$ drift as the electrons from the cathode. The electrons create a region called the Hall-effect region in which the elevated local electron concentration creates a strong electric field potential. The newly created propellant ions experience the effect of this region and are accelerated. Because they are too massive to be trapped by the $\vec{E}\times\vec{B}$ drift that has confined the electrons, the ions are ejected from the device, creating thrust. Propellant is continuously introduced, and the electrons continuously ionize it; electrons from the cathode contribute to the ionization and to downstream neutralization of the ejected beam.

Although Hall thrusters are operated in a steady manner, high-speed interrogation of

these devices reveals a variety of instabilities and oscillations in the discharge, ranging from the low 10s of kHz to the MHz range [47]. These oscillations are an active area of research, and our understanding of their causes and effects is still evolving. Work has shown a correlation between thruster performance and oscillation strength [48], and others have suggested these oscillations may play a role in thruster processes such as anomalous transport [49, 50] and cathode erosion [51].

The dominant discharge current oscillation in Hall thrusters is typically in the low 10s of kHz and referred to as the breathing mode. In this mode, the thruster discharge oscillates in unison in a fashion that has been likened to a predator-prey effect and has been modeled in a number of studies [52–59]. In this framework, the breathing mode is treated as an ionization wave that oscillates axially within the discharge channel (the name thus deriving from the metaphor that the thruster is "inhaling" neutrals and "exhaling" plasma [60]). Fife's two-dimensional modeling, which was able to reasonably reproduce the breathing mode frequency of an SPT-100, derived an expression for the breathing mode frequency:

$$f_b = \frac{\sqrt{u_i u_n}}{2\pi L_i} \tag{2.3}$$

where f_b is the frequency of the breathing mode, u_i is the speed of the ions as they leave the thruster, u_n is the speed of the neutrals entering the ionization zone from the anode, and L_i is the characteristic length of the oscillation, typically taken as the length of the ionization zone. The breathing mode appears in high-speed discharge current traces and can be detected using a high-speed camera [61].

Similarly, another mode that is sometimes detected in Hall thruster discharges is the spoke mode or the local mode. Here, instead of a single axial ionization wave acting on the entire discharge at once, azimuthal ionization waves propagate around the discharge in the $\vec{E} \times \vec{B}$ direction of the thruster. These bright and dark regions of plasma, when captured on high-speed video, look like the spokes of a wheel. This type of oscillation was originally measured by Janes and Lowder [62] in 1966, but to date no coherent theory has yet

been developed to fully explain their origin and mechanism for propagation [60]. Work has shown that Hall thrusters can transition between breathing and spoke modes with changes in magnetic field strength [48, 60]. These oscillations have been measured with near-field ion saturation probes and captured using a high-speed camera in a similar fashion to the breathing mode [60,61]. Anywhere from a single spoke to six simultaneous spokes have been detected in Hall thrusters depending on operating condition. Though the work presented here involves a thruster without magnetic shielding, it is interesting to note that the spoke mode has to date only been detected in magnetically shielded thrusters at very high magnetic field conditions [63].

The four main controls available to the Hall thruster operator are discharge voltage, discharge current, magnetic field strength, and cathode flow rate. Typically, Hall thrusters are operated in a constant-voltage mode, where the discharge voltage is held steady by the power supply (or power processing unit in a flight-like system) and the discharge current is in turn dictated by the propellant mass flow rate. In laboratory operation, the propellant mass flow rate is typically controlled to a constant value, whereas for most flight units the mass flow is actually controlled to the discharge current. That is, the control unit will throttle the propellant flow rate to maintain a constant discharge current. Magnetic field strength and topology are set by the current to the electromagnets of the thruster. Cathode flow rate is typically set to be somewhere between 3–10% of the anode flow rate and is adjusted to minimize cathode coupling voltage and maximize thruster voltage utilization efficiency. More details on cathode flow fraction, especially as it relates to the X3 thruster, are provided below.

2.3.2 History

The essential makeup of Hall thrusters has not changed in the 45 years since they first flew on the Russian Meteor-18 spacecraft. That mission launched in late 1971, and the on-board Hall thrusters successfully fired in early 1972 [8]. Improvements have been made

to their construction and their operation, but the linage is clear. The generic Hall thruster description provided above could as easily describe the Russian SPT-100 [64, 65] from the 1970s as it could a recently-designed thruster such as NASA's HERMeS thruster [66, 67].

As mentioned in the introduction to this section, the initial concept of the Hall thruster first appeared in the U.S. literature [32–34], yet in the years between those initial tests and the end of the Cold War, only a small amount of development work occurred in the United States [68], where most of the attention was on the higher-specific impulse ion thruster. A majority of further Hall thruster development occurred in the Soviet Union [8].

When the Cold War ended and the Soviet Union collapsed, the improved Hall thruster came back to the United States. The Russian scientific literature spoke of a device that ran at 50% efficiency and 1500 seconds specific impulse, which put the device in an optimal operating regime for satellite station-keeping (maintaining the orbit and trajectory of Earth-orbiting satellites) in which ion thrusters did not efficiently function. This claim was first verified by a team of American scientists that traveled to Russia to test the device at the two Russian test facilities, NIITP and Fakel [69]. When their results validated the Russian performance claims, thrusters were brought to the United States for further testing at both NASA Glenn Research Center (GRC) (at the time, NASA Lewis Research Center) [64] and at the Jet Propulsion Laboratory (JPL) [65]. American interest in Hall thrusters grew through the early 1990s, and by the mid 2000s NASA and other institutions had developed their own single-channel Hall thrusters capable of 5 kW and above [70–72]. These devices were all quite similar to the SPT-100 but larger in size.

Hall thrusters began as devices for satellite station keeping, and before the SPT-100 even came to the United States the Soviets had flown more than fifty of the smaller SPT-70 thrusters [69]. However, as development continued, NASA and others began to see EP, and specifically Hall thrusters, as a viable option for efficient in-space propulsion at power levels well in excess of the 1–2 kW of a typical station-keeping thruster. With this in mind, GRC undertook an effort to develop high-power Hall thrusters starting in 1999. This culminated in

a series of 20–50 kW class Hall thrusters that demonstrated for the first time the performance capability of this technology at these power levels. The 50-kW class NASA-457Mv1 thruster, the highest-power thruster produced from this effort, was operated on xenon and krypton propellants through a range of operating conditions, demonstrating on xenon propellant a maximum total power of 97 kW, maximum discharge current of 112 A, maximum total efficiency of 0.58, and specific impulses from 1550-3560 seconds [10, 11, 73]. This thruster demonstrated scaling techniques and physical insight for creating high-power Hall thrusters. Leveraging insight from this work, NASA developed a higher-fidelity version of the thruster named the NASA-457Mv2 which demonstrated improved performance over the v1 thruster [74], though it was not tested beyond 50 kW discharge power. Additionally, the NASA-300M 20-kW thruster and NASA-400M 50-kW thruster were developed using similar scaling techniques, applying design lessons learned to continually improve performance [75,76]. This culminated in a demonstrated peak total efficiency of 0.67 at 500 V, 20 kW with the NASA-300M on xenon propellant, and a peak anode specific impulse of 4940 seconds with the NASA-400M on krypton propellant. A 150-kW single-channel Hall thruster was even designed using these techniques but never built [77]. This thruster, designated the NASA-1000M, would have been 1 meter in diameter, the largest Hall thruster ever built.

2.3.3 Scaling

NASA's development of the NASA-457Mv1, detailed by Manzella [11], derived a number of key scaling relationships that have been used in many thruster designs. Discharge voltages are held within the established range of approximately 150 V to 800 V, although thrusters have been fired in excess of 1000 V [78, 79]. Discharge current density, which is defined as the total discharge current divided by the annular exit area of the discharge channel, is held at established values. These values are export-controlled and will not be explicitly stated here.

As explained by Manzella, the proper discharge current density—fundamentally, propel-

lant density—balances counteractive effects from collisions: increasing propellant density increases the electron-neutral collision rate, thus increasing the ionization rate. Yet it simultaneously decreases electron temperature through elastic and inelastic collisions, and the lower electron temperature decreases the energy electrons have to ionize, which decreases the ionization rate. These fundamental plasma processes are independent of thruster size, and as such, effective values of propellant density hold regardless of thruster size or power level. This has required thruster size to increase as power levels increase: more channel area at the same current density will carry more total current.

Channel length also remains constant, regardless of thruster size, as does magnetic field strength and topology. Channel diameter and width are tuned in these larger thrusters to properly carry the target discharge current while maintaining magnetic field properties. This type of scaling worked well in going from the 10-cm diameter, 1.35-kW SPT-100 thruster to 5-6 kW thrusters with diameters of roughly 15–20 cm such as the NASA-173M [71] and the H6 [72, 80], and even up to the aforementioned 50-kW NASA-457Mv1. Although the GRC scaling effort demonstrated a road map toward 150-kW Hall thruster systems, one of the major challenges identified in this program (and exemplified by the NASA-1000M thruster design) was the excessively large footprint of higher-power thrusters. This is due to the fact that thruster diameter increases with power using these scaling techniques. Scaling a single-channel Hall thruster to the 100-300 kW targeted by NASA for its near-future manned and deep-space mission goals [30, 81] would require thruster diameters on the order of 3 meters [82], which becomes problematic from both a manufacturing perspective and a spacecraft-integration perspective. There are also questions regarding whether the discharge plasma would couple properly with the cathode on a thruster that size.

One technique to avoid this issue and scale Hall thrusters beyond 50-kW class devices while limiting diameter increase is to concentrically nest multiple discharge channels around a shared centrally-mounted cathode. This technique allows for improved packing density of the channels as compared to multiple single-channel thrusters [83, 84]. Nesting the dis-

charge channels of Hall thrusters was first explored in the open literature by Liang with the X2 thruster. Busek Co., Inc., has developed a similar nested Hall thruster (NHT), though nothing about the device other than a small mention and a photograph has been published in the open literature [85]. There is a patent by United Technologies for a NHT concept [86], and the Russian literature briefly mentions 3-channel NHT that operated each channel on a separate phase of 3-phase AC power [87]. Beyond these brief references, no rigorous thruster development work was documented openly before the X2. The X2 was a 6-kW thruster featuring two concentric discharge channels. Liang's research showed that the thruster performed comparably to single-channel thrusters of its class [88,89], with some interesting phenomena observed in both the ion current density profiles and in the recorded thrust while operating the two channels simultaneously [1]. A photo of the X2 is shown in Figure 2.3.

The success of the X2 proved a number of different concepts: that a multi-channel nested Hall thruster could be built; that the channels could be operated simultaneously; that the thruster would operate comparably to single-channel thrusters of its power class, both in single- and dual-channel mode; that the device could be run off of a single, centrally-mounted cathode; that a nested magnetic lens topology could be created; and that the magnetic circuit could share components between adjacent channels.

2.4 The X3

Building on the major contributions of the X2, the Plasmadynamics and Electric Propulsion Laboratory (PEPL) at the University of Michigan, in collaboration with the Air Force Office of Scientific Research, the Air Force Research Laboratory Rocket Lab, the Michigan/Air Force Center for Excellence in Electric Propulsion, NASA GRC, and JPL, developed a three-channel, 100-kW class NHT named the X3. Detailed in work by Florenz [4], the X3 demonstrated that the NHT concept could be used to substantially increase the discharge

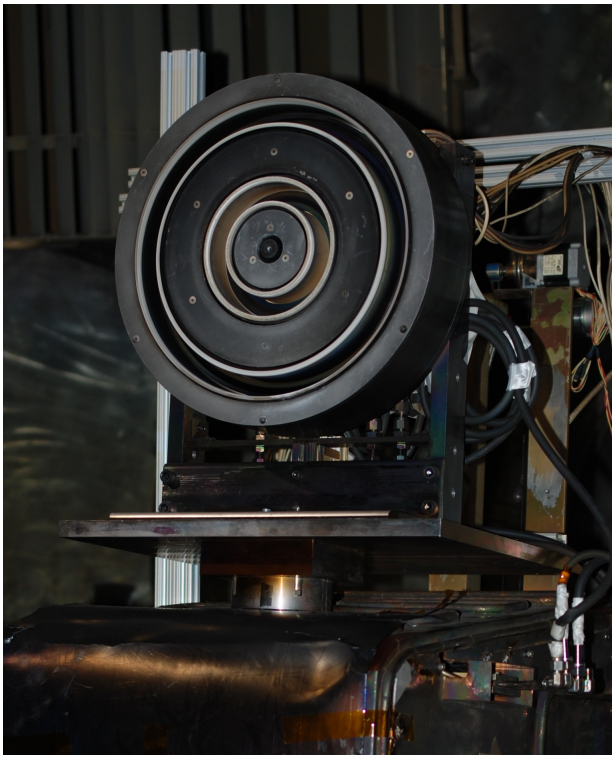

(a) Mounted in the Large Vacuum Test Facility at the University of Michigan.

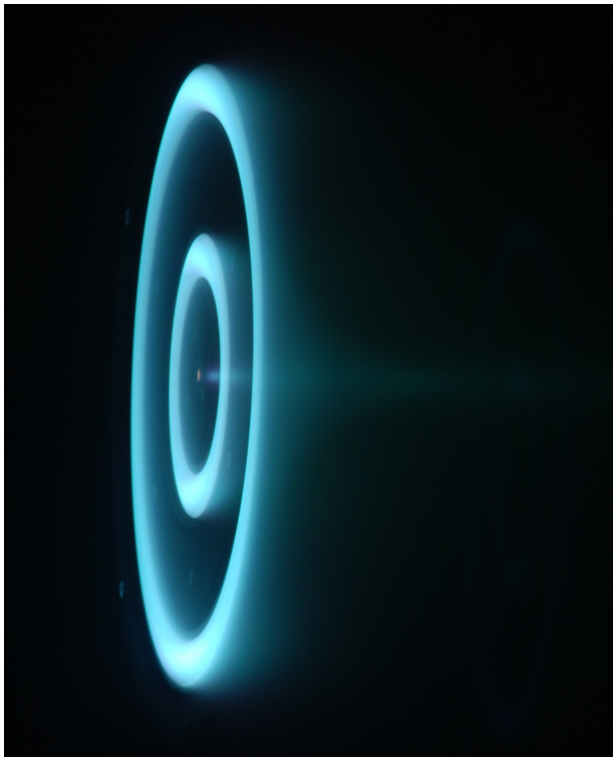

(b) Firing in dual-channel mode.

Figure 2.3: The X2, a 2-channel, 6-kW NHT that served as a proof-of-concept thruster for the nested configuration: (a) mounted in the Large Vacuum Test Facility at the University of Michigan and (b) firing in dual-channel mode.

power of laboratory Hall thrusters while maintaining a compact device footprint. The X3 is designed to be capable of operation up to 250 kW discharge power (at 1000 V discharge voltage), yet is only approximately 80 cm in diameter, a significant improvement over the 100 kW, 100 cm diameter NASA-1000M.

2.4.1 Description

The X3 was a laboratory NHT developed from 2010-2013 and first fired in September 2013. As noted above, it is approximately 80 cm in diameter, and weighs approximately 230 kg. The thruster can operate in seven different configurations: each channel alone, all three together, and any combination of two. The channels are referred to as "Inner" or "I", "Middle" or "M", and "Outer" or "O", where the Inner channel is the smallest diameter and the Outer channel is the largest. Combinations of channels are then referred to by their channel combinations, e.g., "IM" indicating the Inner and Middle channels operating together. The thruster is shown in Figure 2.4.

2.4.2 History: Design and Previous Work

The constraints that were chosen during the conceptual phase of the X3 design were in part:

- Maximum discharge power of 200 kW. This was based on estimates that suggested that this power level was an upper-bound for the capabilities of existing ground-test facilities [90]. As will be explored below, by clustering multiple thrusters together, a 200-kW thruster enables a wide variety of potential missions for total propulsion system powers in excess of 1 MW.

- Operation from 300 V to 800 V discharge voltage, with a stretch goal of 1000 V based on previous work with the NASA-457Mv1 [73] and the NASA-400M [76]. As noted, 1000 V operation would push the upper power limit of the X3 to 250 kW.

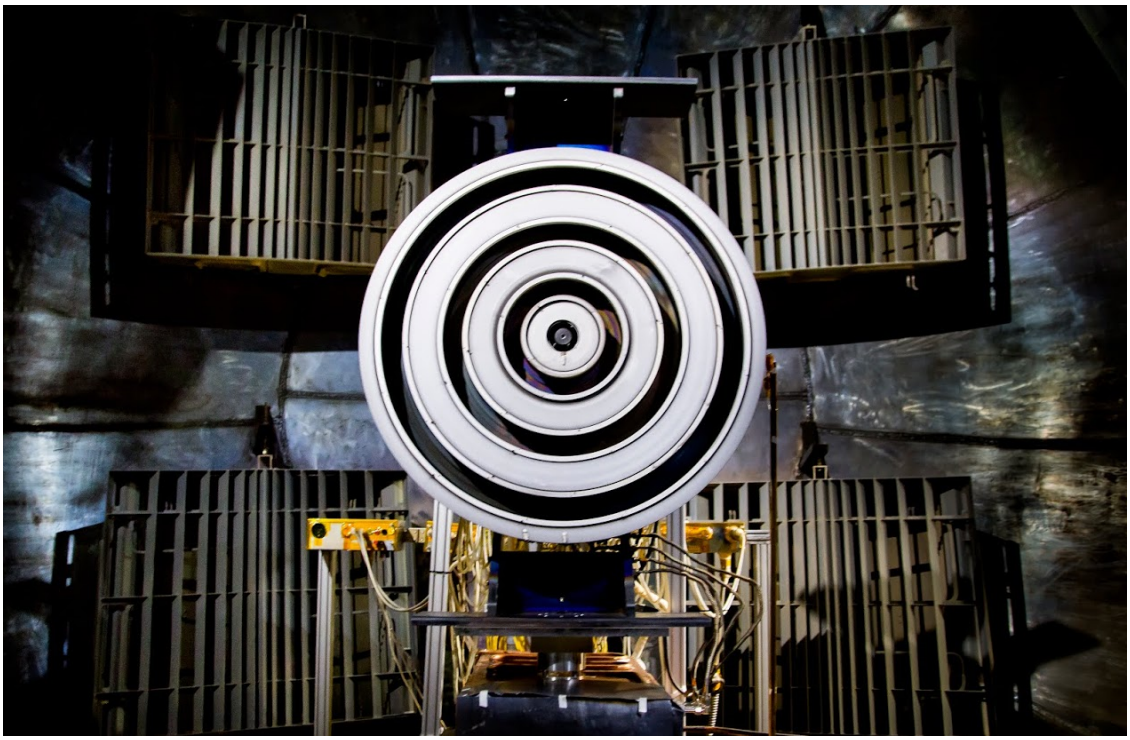

(a) Mounted in the vacuum chamber.

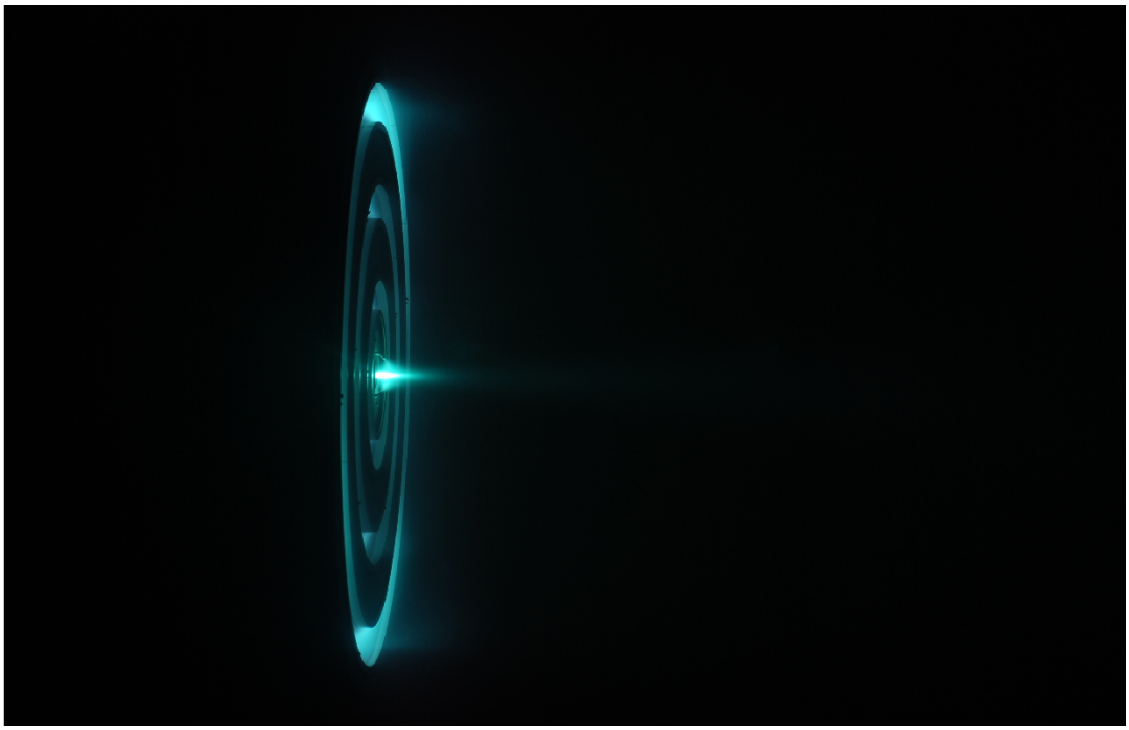

(b) Firing at 30 kW in the IMO configuration.

Figure 2.4: The X3, a 100-kW class, three-channel nested Hall thruster: (a) mounted in the vacuum chamber and (b) firing at 30 kW in the IMO configuration.

- Capability of operation on both krypton and xenon propellant. As will be discussed below, krypton provides higher specific impulses at similar discharge voltages as compared to xenon due to its smaller mass, which expands the possible operating envelope of the thruster.
- Operation spanning discharge current densities tested with the H6 [91].
- Three concentric channels, capable of operating alone or simultaneously in any combination.
- Nested magnetic lens topology and shared pole pieces between adjacent channels.
- Operation with a single centrally-mounted cathode.

Based on these criteria, a lengthy design process was undertaken, as summarized in References [4] and [92] and detailed more thoroughly in Reference [91]. The initial step of the X3's design was to use the scaling laws developed by Manzella [11] described above to determine the basic dimensions of the thruster. The channels feature very similar cross-sections: each channel is the same length (from anode exit plane to channel exit plane), and the channels are similar in width. Using the selected range of current densities, the exit areas of the three channels were selected.

With the basic geometry of the thruster set, the details of the propellant distribution scheme and magnetic field topology could be designed. Simulations of many design iterations were performed of both the propellant distributors and the magnetic circuit. Three propellant distribution schemes were explored thoroughly: a porous design, which used porous media and a baffle chamber; a drop-in design, which used small holes and a series of baffles and manifolds in a separate unit that was inserted into the channel cup; and an integrated design, which featured flow inlets integrated into the channel cup. Schematics of these design concepts can be found in References [4] and [92]. Ultimately, a drop-in design was selected, which had the most heritage to other thruster designs. The anodes of all three channels

feature similar cross-sections and were designed to provide uniform flow across the projected range of propellant flow rates.

The X3 features an inside and an outside electromagnet for each discharge channel, resulting in a total of six, typically referred to as "Magnets 1–6" or "M1–6", where M1 and M2 are the inside and outside magnets for the Inner channel, M3 and M4 for the Middle channel, and M5 and M6 for the Outer channel. Each consists of a coil of wire around a single bobbin to minimize footprint. The X3 does not feature any trim coils like those used in thrusters such as the NASA-173M [79] and the X2 [89]. The shared magnetic circuit causes the radial magnetic field direction to be the same for the Inner and Outer channels but to switch for the Middle. The topography of the magnetic field of the X3 is a plasma lens, as was discussed above and described in detail in References [37], [38], and [79]. The X3 is an "un-shielded" thruster, meaning that the magnetic field lines intersect the channel walls and cause erosion of the boron nitride.

Thermal modeling during the design phase of the thruster suggested that the thruster's ability to radiate heat was critical to maintaining reasonable temperatures at high discharge powers. To facilitate this, the outer surface of the X3 is coated in an alumina spray coating to increase its emissivity. The alumina spray coating provides a higher emissivity by nearly a factor of two over the bare magnetic iron. Thermal data from initial firings of the X3 suggested that the thruster was operating cooler than the model suggested [4], and work is ongoing to improve the fidelity of this modeling effort based on collected thruster data [93].

A series of photographs from the first experimental campaign of the X3 are presented in Figure 2.5. The top-left photograph shows the X3 installed in the Large Vacuum Test Facility at the University of Michigan prior to first light. Prior to the work described in this document, the thruster was fired in all seven of its configurations through a brief burn-in on krypton propellant, followed by a limited performance characterization on xenon propellant. A photograph from the first light of the X3, which was of the Inner channel on krypton propellant in September 2013, is presented in the top-right of the figure. During that initial

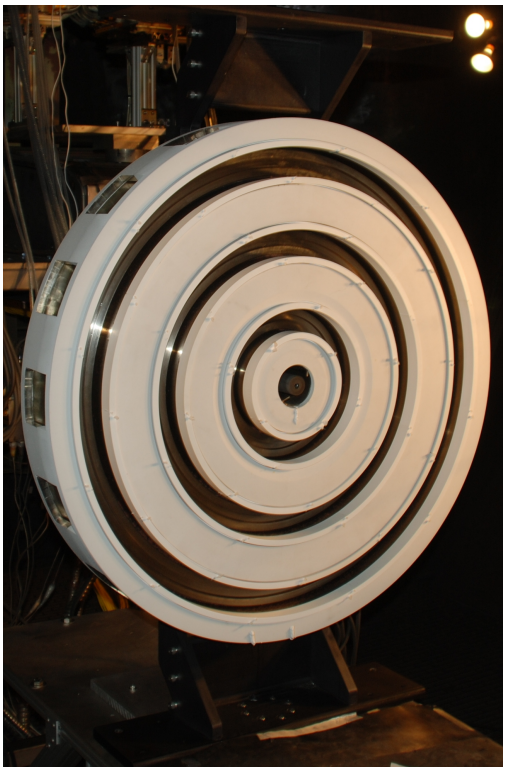

(a) Prior to first light

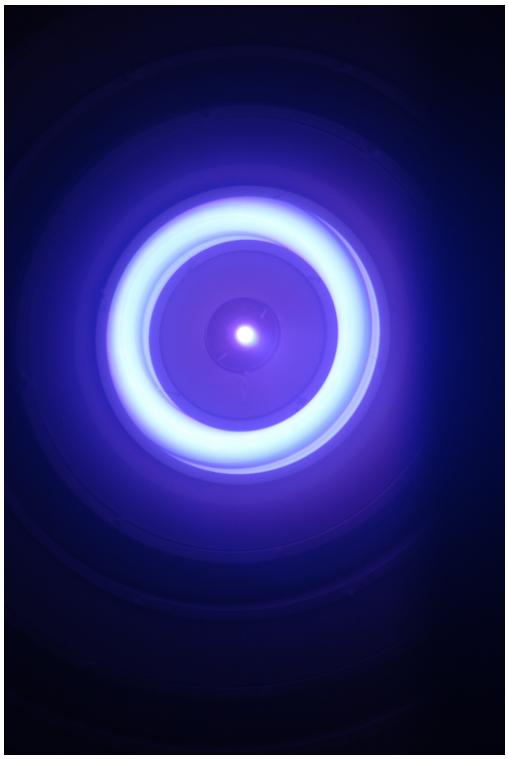

(b) First light (Inner channel), krypton

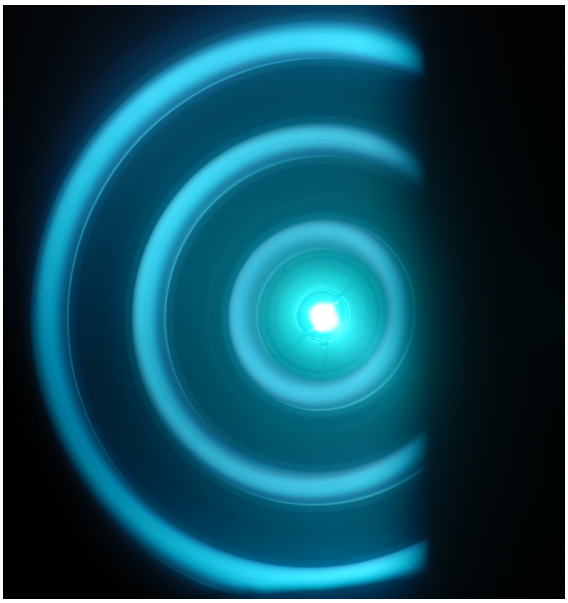

(c) IMO operation, xenon. The right hand side of the thruster is blocked by the facility beam dump.

Figure 2.5: Photographs from the first test campaign of the X3 in late 2013: (a) prior to first light, (b) during first light of the thruster, which was of the Inner channel on krypton propellant, and (c) during IMO operation on xenon propellant. In (c) the right hand side of the thruster is blocked by the facility beam dump.

firing, steady operation of the Inner channel was demonstrated at 300 V, 4.2 kW for 2 hours.

During the krypton burn-in process, the thruster went through initial outgassing and thermal cycling, and wear bands [79] were established on the boron nitride discharge channel rings. The original throttling table for the test called for operation of each channel individually at 300 V, 550 V, and 800 V for durations of approximately four hours at each voltage, but intermittent facility issues plagued operation at higher voltages and testing proceeded to the xenon-propellant portion of the test having only achieved a de-scoped test matrix. More details of this test, including thruster telemetry and high-speed discharge current traces, can be found in Reference [4]. Following the krypton burn-in, a preliminary xenon performance characterization was undertaken. The test used a modified inverted-pendulum thrust stand to measure the thrust of the X3 across all seven configurations at 300 V and a low current density. A photograph from this test of the X3 firing on xenon in the IMO configuration is presented in the bottom of Figure 2.5. A summary of the results of this test are presented below.

2.4.3 NextSTEP

NASA considers NHTs a promising technology and is funding continued development through the Next Space Technologies for Exploration Partnerships (NextSTEP) program, which is investing in the technologies that will be necessary for future crewed missions to Mars [81]. In total, three electric propulsion concepts are being funded by NextSTEP: the Variable Specific Impulse Magnetoplasma Rocket (VASIMR) [94, 95]; the electrodeless Lorentz-force thruster [96]; and the XR-100 nested Hall thruster system. The overall goal of the these projects is to demonstrate 100 continuous hours of 100-kW operation of the system operating at a total system efficiency in excess of 60%. The XR-100 system is being developed by a team led by Aerojet Rocketdyne (AR) and including UM, GRC, and JPL. The system consists of the X3 NHT, a JPL-developed high-current hollow cathode [97], and a power processing unit and xenon flow controller being developed by AR. Additional con-

tributions include plasma and thermal modeling by JPL and facility and test infrastructure by NASA GRC. A system-level overview of the project can be found in Reference [98].

2.5 Performance Metrics

The performance of propulsion devices is evaluated using a number of metrics. We will apply and reference these metrics, which apply generally to other devices but here are presented in the context of Hall thrusters, throughout this work. They characterize how efficiently and to what speed a device accelerates propellant particles. Arguably the most important parameter for any propulsion system is the thrust it produces. Thrust is defined as the reaction force that a propulsion device creates. We express this mathematically as:

$$T = \dot{m} u_e. \tag{2.4}$$

The mass flow rate $\dot{m}$ is a readily-known parameter in Hall thruster testing, as it is a prescribed parameter set by highly-accurate, calibrated mass flow controllers. The exhaust velocity u_e is much more difficult to measure for Hall thrusters. However, because thrust is a reaction force, it is directly measurable in the laboratory. Measurement of the low thrust that characterizes most electric propulsion devices is more complex and delicate than the simple load-cell methods used for high-thrust chemical systems [99], but it is nevertheless a mature and widely-adopted diagnostic technique.

As mentioned above and introduced in Equation 2.2, specific impulse is another important metric in propulsion. It indicates how effectively propellant is being converted into thrust. Without the ability to easily characterize the exhaust velocity of a Hall thruster, an alternate method to calculate specific impulse in the laboratory is required. Rearranging Equation 2.4 to solve for u_e and substituting that into Equation 2.2, we find a definition of specific impulse that uses quantities that are measurable in the laboratory:

$$I_{sp} = \frac{T}{\dot{m}g}. \tag{2.5}$$

Another important performance metric for Hall thrusters is the efficiency. First, we define jet power, which is the time rate of expenditure of the kinetic energy of the beam, as:

$$P_{jet} = \frac{1}{2}\dot{m}u_e^2. \tag{2.6}$$

The total efficiency of the thruster is then defined as the ratio of jet power (that is, effective thrusting power) to input power (for electric propulsion devices, the electrical power):

$$\eta = \frac{\text{jet power}}{\text{electrical power}}. \tag{2.7}$$

We can then express this using existing definitions as:

$$\eta = \frac{\frac{1}{2}\dot{m}u_e^2}{P_{elec}}. \tag{2.8}$$

Using the definition of u_e extracted from Equation 2.4, we re-write total efficiency in terms of measurable quantities as:

$$\eta = \frac{T^2}{2\dot{m}P_{elec}} \tag{2.9}$$

Thus, the thrust, specific impulse, and efficiency of a thruster can be known using parameters measurable in the laboratory.

2.5.1 Anode Quantities

To calculate the total efficiency and specific impulse of a device, the total electrical power and mass flow rate are used in the calculation of Equations 2.5 and 2.9. For a Hall thruster, the total power includes the power dissipated by the electromagnets (typically a few percent of the discharge power or less) and by the cathode keeper if it is not floating (typically a

small fraction of discharge power), and the total mass flow rate includes the cathode flow (typically somewhere between 5% and 10% of the discharge flow rate [100], though some thrusters require significantly higher values [76]). In laboratory thrusters, often little effort is made to minimize these values. To isolate the discharge behavior from the unoptimized cathode and magnet contributions, across the Hall thruster performance literature so-called anode, discharge, or thrust versions (all are names for the same quantities) of specific impulse and efficiency are often used.

Anode specific impulse is calculated as:

$$I_{sp,a} = \frac{T}{\dot{m}_a g}, \tag{2.10}$$

and anode efficiency as:

$$\eta_a = \frac{T^2}{2\dot{m}_a P_d}, \tag{2.11}$$

where anode mass flow $\dot{m}_a$ is used in place of the total mass flow rate $\dot{m}$ and the discharge power P_d is used in place of the total electrical power P_{elec}.

2.5.2 Phenomenological Efficiency Model

A significant body of work exists exploring ways to further decompose Hall thruster efficiency in an effort to isolate physical processes and better understand the specific loss mechanisms [79, 101–105]. Dissecting thruster performance in this manner provides insight into how the plasma processes change with changes in thruster operating parameters such as discharge voltage and propellant flow rate. This type of analysis can aid with targeted thruster design improvements, as was the case with the development of the plasma lens magnetic field topography [3, 79].

Hofer developed a four-component model [102] that later gained a fifth component [104]. This model has been used in many works across the literature [2, 72, 105–108]. The five components of this model are:

1. **Charge utilization efficiency** (η_q), which captures loss effects due to the creation of multiply-charged ions.

2. **Voltage utilization efficiency** (η_v), which is related to the conversion of voltage into axially-directed velocity. Voltage utilization efficiency can also be thought of as related to the spread in ion velocities, which can be a result of factors such as the spatial variation in the ionization location and the overlap of the ionization and acceleration regions of the discharge, as well as the magnitude of the cathode coupling voltage.

3. **Current utilization efficiency** (η_b), which is the fraction of ion current in the discharge current. Discharge current in a Hall thruster is the sum of the ion and electron currents, yet only ion current contributes to thrust generation.

4. **Mass utilization efficiency** (η_m), which is related to the conversion of neutral mass flux into ion mass flux.

5. **Plume divergence utilization efficiency** (η_d), which accounts for the finite beam divergence in Hall thrusters.

The product of these five efficiency components is the anode efficiency:

$$\eta_a = \eta_q \eta_v \eta_b \eta_m \eta_d, \tag{2.12}$$

and each of the components are calculated as follows:

$$\eta_q = \frac{\left(\sum \frac{\Omega_i}{\sqrt{Z_i}}\right)^2}{\sum \frac{\Omega_i}{Z_i}} = \frac{\left(\sum \frac{\Omega_i}{\sqrt{Z_i}}\right)^2}{\alpha_m} \tag{2.13}$$

$$\eta_v = \frac{V_a}{V_d}, \tag{2.14}$$

$$\eta_b = \frac{I_b}{I_d}, \tag{2.15}$$

$$\eta_m = \frac{\dot{m}_b}{\dot{m}_a} = \alpha_m \left(\frac{m_{Xe} I_b}{\dot{m}_a e} \right), \tag{2.16}$$

$$\eta_d = (\cos\theta)^2 , \tag{2.17}$$

where Ω_i is the current fraction of the i^{th} species, Z_i is the charge state of the i^{th} species, V_a is the acceleration voltage, V_d is the discharge voltage, I_b is the beam current, I_d is the discharge current, $\dot{m}_b$ is the ion beam mass flow rate, m_{Xe} is the mass of xenon, e is the elementary charge, and θ is the divergence angle of the beam.

These efficiency parameters are typically studied in the laboratory using an array of plasma diagnostics in the thruster plume [1, 71, 72, 79, 106] that includes Langmuir probes, retarding potential analyzers, Wien Filter Spectrometers (also called ExB probes), and Faraday probes. More about these plasma diagnostics and the computations necessary to calculate the efficiency parameters is presented in Chapter 4 below. As will be discussed, we were able to utilize only a limited plasma diagnostics array on the X3 during high-power testing. Thus, small modifications were made to the phenomenological model to accommodate the collected data. We detail these modifications in the following section.

2.5.3 Limited-Diagnostics Phenomenological Efficiency Model

Early Hall thruster phenomenological efficiency analysis work such as that by Hofer [79] did not include η_d as in Equation 2.17. Instead, diagnostics were coupled with thrust-stand-based performance measurements which provided anode efficiency η_a, and the beam current I_b was solved for as the only unknown in the system. It was only later that η_d was explicitly added as a separate term [104].

As will be discussed in more detail below in Chapter 4, the plasma diagnostic arm for the high-power X3 testing was limited to spatially-static (non-moving) probes, and as such, we were unable to include a Faraday probe for ion beam measurements. Faraday probes provide I_b and θ measurements. To provide phenomenological efficiency analysis with this setup, we revert to a model more like that of Hofer's early work in [79], but with the beam

divergence explicitly called out (though not solved for).

If we substitute Equations 2.13–2.17 into Equation 2.12 and re-arrange for the two unknowns (I_b and θ), we find:

$$I_b \cos\theta = \sqrt{\frac{\eta_a}{\eta_q \eta_v \alpha_m} \frac{I_d e \dot{m}_a}{m_{Xe}}}, \tag{2.18}$$

which can be solved for explicitly using probe diagnostic data and thruster telemetry information. Using $I_b \cos\theta$, the mass utilization efficiency η_m and current utilization efficiency η_b can be re-written into what we call their divergence-weighted values:

$$\eta_{md} = \left(\frac{m_{Xe} I_b \cos\theta}{\dot{m}_a e} \right) \alpha_m \tag{2.19}$$

and

$$\eta_{bd} = \frac{I_b \cos\theta}{I_d}, \tag{2.20}$$

which are actually identical to those developed by Hofer in [79] but which explicitly note that the beam current I_b is weighted by the beam divergence angle θ. These modified efficiency values will be calculated for the X3 using probe data collected during high-power testing.

2.5.4 Mass Ingestion Efficiency Correction

The thrust produced by a Hall thruster is a function of the actual or total mass flow rate being processed by the discharge, summed over all channels:

$$T = u_e \sum_i \dot{m}_{ch,act,i}, \tag{2.21}$$

where $\dot{m}_{ch,i}$ is the total flow rate being processed by channel i. This flow rate has two major sources in Hall thruster ground testing:

$$\dot{m}_{ch,i} = \dot{m}_{a,i} + \dot{m}_{ingest,i}, \tag{2.22}$$

where $\dot{m}_{a,i}$ is the prescribed anode flow rate to channel i and $\dot{m}_{ingest,i}$ is flow rate ingested by that channel by separate mechanisms, for instance by flux to the channel from a ground test facility with finite background pressure. The discharge power of a Hall thruster is similarly dependent on the total flow rate being processed by the channel:

$$P_d = \sum_i V_{d,i} I_{d,i} = \sum_i V_{d,i} \frac{e}{m_{prop}} \dot{m}_{ch,i}, \tag{2.23}$$

where P_d is the discharge power, $V_{d,i}$ and $I_{d,i}$ are the discharge voltage and current of channel i, e is the elementary charge and m_{prop} is the propellant atomic mass. This is due to the fact that the discharge current is dictated by the total mass being processed by the discharge from all sources.

Both T and P_d as described in Equations 2.21 and 2.23 are quantities that are measured in the laboratory. Thus, it is apparent that to properly calculate anode or total efficiency values, the proper mass flow rate to use is not simply $\dot{m}_a$ but $\dot{m}_{ch}$, which takes into account the "free" propellant not accounted for by the anode mass flow rates alone. A correction must be made to the so-called "lab" anode efficiency, which only uses the anode mass flow rates in the calculation, to account for this effect. We can re-write Equation 2.11 to accommodate this updated mass flow rate as:

$$\eta_{a,corrected} = (\eta_{a,lab}) \left(\frac{\dot{m}_a}{\dot{m}_{ch}} \right). \tag{2.24}$$

We find that this correction factor is less than unity when there is additional propellant being ingested, which makes sense as the corrected anode efficiency value is now accounting for all propellant the thruster is processing.

2.5.4.1 Test Facilty Ingestion

Traditionally, $\dot{m}_{ingest}$ due to facility ingestion ($\dot{m}_{facility}$) is calculated as the flux to the discharge channel using kinetic theory [67, 109–111]:

$$\dot{m}_{facility} = \Phi m_{prop} A_{ingest} = \sqrt{\frac{m_{prop}}{2\pi k_b T_g}} p_b A_{ingest}, \tag{2.25}$$

where Φ is the particle flux, A_{ingest} is the ingestion area (here, the thruster discharge annular exit area), k_b is Boltzmann's constant, T_g is the gas temperature (typically taken as the anode temperature, around 500 K), and p_b is the facility background pressure.

However, tests with multiple thrusters in multiple facilities has demonstrated that Equation 2.25 does not fully account for the changes seen in thruster operation with changing facility background pressure [67, 110, 111]. Recent work by Frieman has demonstrated that the so-called thermal model, as Equation 2.25 is referred to, is not accurate because it fails to account for the presence of bulk motion in the background gas typical to most EP test facilities [112, 113]. In place of the thermal model, Frieman proposes a model called the background flow model that captures these bulk motion effects. Frieman's work generalizes a model originally developed by Cai [114] to accommodate a variety of facility configurations. The model was validated against a variety of experimental data and generally showed significantly improved agreement over the thermal model results. The background flow model calculates an expected flux at the plane of the thruster based on facility factors such as the cryopump temperature, surface area, sticking coefficient, and location, as well as facility cross-section. The model also accounts for collisional scattering between the neutrals flowing toward the thruster and those flowing out of the thruster. From this expected flux, the ingested mass flow rate $\dot{m}_{facility}$ is calculated. The mathematics of the model are left to the reader in Reference [113].

We applied the background flow model here to calculate a more accurate value of $\dot{m}_{facility}$. We used specific coefficients to represent GRC's Vacuum Facility 5 (where the high-power

X3 testing occurred) in the model that matched validation work done by Frieman [1]. These values are based on the pump configuration (pump area, temperature, and location) of the facility and can also be found in NASA facility literature [2]. Using this value of $\dot{m}_{facility}$ in Equation 2.24, we find an expression for anode efficiency corrected for facility ingestion:

$$\eta_{a,facility} = \eta_{a,lab} \frac{\dot{m}_a}{\dot{m}_a + \dot{m}_{facility}} = \eta_{a,lab} \eta_{facility}, \tag{2.26}$$

where we define the mass flow rate fraction as $\eta_{facility}$. This new efficiency term will be unity when no background gas is ingested and less than unity for cases with ingestion, thus correcting the anode efficiency to reflect the additional propellant mass.

2.5.4.2 NHT Cross-Channel Ingestion

In addition to facility ingestion, each channel of a nested Hall thruster can also ingest neutral gas from other operating channels of the thruster. This is a separate effect from facility ingestion because it will still exist on-orbit and is not a by-product of testing in a facility with finite pumping speed. In its most general form, this NHT ingestion term $\dot{m}_{NHT}$ can be expressed as:

$$\dot{m}_{NHT,i} = \sum_j \gamma_{ij} \dot{m}_{a,j}. \tag{2.27}$$

Each channel i ingests some amount of neutral flow from the other firing channels j. This ingestion from each channel is a function of that channel's flow rate ($\dot{m}_{a,j}$) and a geometry factor (γ_{ij}), which itself is a function of the relative ingestion areas and proximity of channel j to channel i, as well as a weighting factor accounting for whether the channel is operating or simply flowing cold gas. The summation over j should also include the cathode, which is also a source of neutral particles. This is of increased importance on thrusters such as the X3 that feature external neutral injection at the cathode.

[1] Personal correspondence, October 2017.

[2] For instance: https://www1.grc.nasa.gov/facilities/epl/. Accessed January 1, 2018.

Thus, following the development above, we can expect a similar mass flow rate fraction relating the effect of the ingestion from other channels on the anode efficiency of the thruster. However, unlike the $\eta_{facility}$ term, which should be removed from $\eta_{a,lab}$ to correct for effects not seen on-orbit, the effects due to propellant-sharing between channels is an effect that will be seen on-orbit and thus does not need to be removed from the ground-test thruster efficiency analysis. However, we can still separate the value and calculate it as a separate efficiency term in the phenomenological model. In this manner we find:

$$\eta_{a,lab} = \eta_{a,space} = \left(\frac{\eta_{a,lab}}{\eta_{facility}}\right)\left(\frac{\sum_i \dot{m}_{a,i} + \dot{m}_{facility,i} + \dot{m}_{NHT,i}}{\sum_i \dot{m}_{a,i} + \dot{m}_{facility,i}}\right) = \frac{\eta_{a,lab}}{\eta_{facility}}\eta_{NHT}, \tag{2.28}$$

where we find that η_{NHT}, as it is defined, boosts the effective in-space anode efficiency $\eta_{a,space}$ when the channels are ingesting from each other and is unity when there is no cross-channel neutral sharing.

Determining $\dot{m}_{NHT}$ in the lab can be difficult, as a number of competing phenomena are at work. The ideal way to do so would be to operate each of the channel combinations at a fixed facility background pressure and magnetic field and observe how the ingestion (by way of changes in the discharge current) varies between cases for a given set of discharge voltages and currents. Variations could be made in cathode flow and channel flow separately to make quantitative assessments of γ_{ij} for each combination. However, this manner of highly-controlled experiment may be difficult, especially for large thrusters where the background pressures vary considerably between conditions. It must be understood that a η_{NHT} value established for a set of channel combinations at fixed discharge voltage and current but for varying background pressure and magnetic field settings likely contains effects from multiple phenomena, not all of which are exclusive to NHTs. For instance, work has shown that the discharge current oscillation level varies with facility pressure [111], and that these levels are correlated with changes in thruster performance [48]. Changes in background pressure across conditions may cause a channel to change oscillation modes and thus change performance, resulting in an anomalous value of η_{NHT} from an apparent difference in required propellant

flow rate that is not a product of cross-channel ingestion. Background pressure has also been shown to shift the location of the acceleration region within the Hall thruster discharge channel, again resulting in a change in performance and operation of the thruster [115]. These effects are difficult enough to uncouple from one another for single-channel thrusters, and the addition of multiple channels interacting with one another only increases the complexity. However, a preliminary attempt can still be made to estimate the value of η_{NHT} from thruster performance data.

2.6 X3 Performance

2.6.1 Expected

The X3 was developed using the Hall thruster scaling laws described above and developed by Manzella [11]. As such, we can derive the expected performance of the X3. These figures will be based on experimental results from a number of thrusters also developed using similar scaling procedures: the NASA-300M [75], the NASA-400M [76], the NASA-457Mv1 [10, 11, 73], the NASA-457Mv2 [74], and the H6 [72]. The X3 is designed to operate at similar current densities and discharge voltages as these thrusters. As such, we can infer a number of expected performance parameters for the X3 from the performance data reported from these thrusters. A number of these thrusters, including the NASA-300M, NASA-400M, and NASA-457Mv1, were also operated on krypton, providing ample data from which to draw expected krypton performance of the X3.

All of these thrusters are un-shielded: though tests have shown that performance is generally comparable when comparing the un-shielded and shielded version of the same thruster (e.g., for the H6MS [41] and for the NASA-300MS [116]) and when comparing a natively-shielded thruster to unshielded thrusters [67], the selected un-shielded thrusters are more directly related to the design of the X3, and thus to its expected performance.

It should be noted that none of these expected performance values are drawn from the

performance of any nested thrusters. The X2's performance values have been left out due to that thruster's differences in channel scaling. The work of the X2 suggested that there may be additional processes present during dual-channel mode [1]. If those processes extend to the X3, they can be expected to change the performance values during multi-channel operation. This analysis does not make an attempt to account or predict for these possible differences.

2.6.1.1 Specific Impulse

As discussed above for Equation 2.2, the specific impulse of a thruster is directly related to the exhaust velocity of its propellant. In an electrostatic device such as a Hall thruster, the electric potential through which an ion is accelerated dictates its velocity. This potential is directly related to the discharge voltage in a Hall thruster, and as such, it is generally expected that specific impulses should increase with increasing discharge voltage.

We can create an expression for theoretical specific impulse by performing an electrostatic energy balance:

$$I_{sp,theory} = \sqrt{\frac{2eV_d}{m_{prop}g^2}}, \tag{2.29}$$

where all variables hold their previous definitions. We can use this expression to derive theoretical specific impulse values at various discharge voltages. As can be seen, higher discharge voltages are expected to yield higher specific impulses. Also, the smaller mass of krypton provides a higher specific impulse value at the same discharge voltage. Values predicted by this expression for discharge voltages of 300 V and 800 V are 2140 seconds and 3490 seconds for xenon and 2680 seconds and 4370 seconds for krypton, respectively.

The actual exhaust velocity produced by a thruster, and thus the actual specific impulse, can be affected by a number of factors in a Hall thruster. Most notably, Hall thruster ion beams have a non-negligible divergence angle. This indicates that a non-negligible population of ions have a portion of their velocity off-axis, which will not contribute to thrust (from which specific impulse is calculated in the laboratory). Other efficiency losses, which

will be discussed below in more detail, can limit the velocity achieved by the ions. Finally, the presence of multiply-charged ions will actually cause an increase in the effective specific impulse because those ions are accelerated to higher velocity than their singly-charged counterparts. Experimental data from the selected thrusters firing on xenon propellant shows that total specific impulse values tend to be lower than the theoretical value by about 10% at 300 V but approach and even exceed the theoretical value around 600 V. Peterson's work on the NASA-400M suggested that this was due to an increase in the multiply-charged ion population at higher discharge voltages [76].

When the total specific impulse for a given discharge voltage is plotted as a function of discharge power (which is effectively plotting against mass flow rate or current density), all of the selected thrusters demonstrated generally flat behavior except at low power levels or flow rates, where specific impulse values typically decreased. This indicates that these low flow rates may provide too low a current density for optimal Hall thruster operation.

Based on these results, the X3 can be expected to produce total specific impulses:

- On the order of 1800–2000 seconds at 300 V discharge voltage
- On the order of 2600–2800 seconds at 600 V discharge voltage
- On the order of 2900–3300 seconds at 800 V discharge voltage

with the possibility of larger values if the population of multiply-charged ions increases with discharge voltage.

The use of krypton as a propellant is intended to increase the specific impulse capability of the X3. Previous results on the NASA-400M [76] and NASA-457Mv1 [73] indicated that, if the discharge voltage stretch goals for the X3 are realized and the thruster is operated at 1000 V on krypton, specific impulses around 4500 seconds may be possible. The NASA-400M achieved 4943 seconds anode specific impulse at 1050 V and an anode efficiency of 0.68, and the NASA-457Mv1 achieved 4495 seconds total specific impulse at 1000 V and 0.63 anode efficiency.

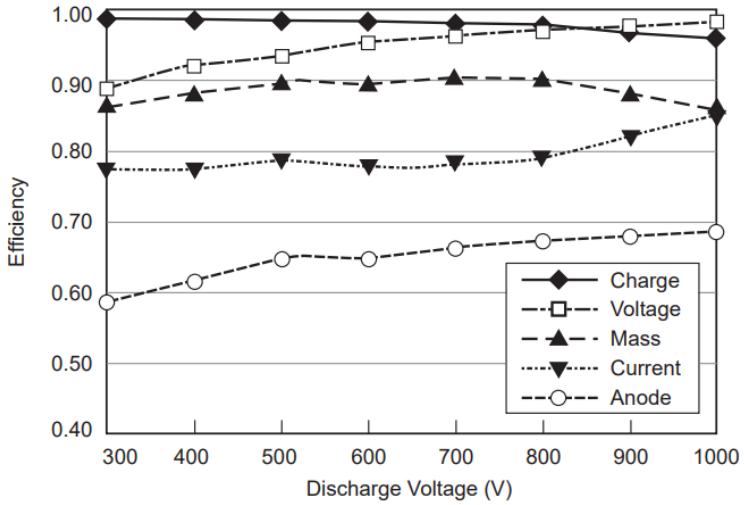

Figure 2.6: Trends in the five phenomenological efficiencies as function of discharge voltage for the NASA-173Mv2 Hall thruster. Figure reproduced from [2], which itself reproduces data from [3].

2.6.1.2 Efficiency

As noted above, the anode efficiency can be decomposed into charge, voltage, current, mass, and beam divergence utilization partial efficiencies. The trends in these then dictate the trend in anode efficiency. Hofer's work with the NASA-173M thruster provided clear trends for the various efficiencies as a function of discharge voltage at constant flow rate [3]. Figure 2.6 shows the trends of these various efficiencies as a function of discharge voltage for the NASA-173Mv2 thruster [2, 3]. As can be seen, the anode efficiency increases monotonically with discharge voltage. This is driven by the increases in voltage utilization efficiency and current utilization efficiency. These results, which pioneered the efficient operation of Hall thrusters at high discharge voltage/specific impulse, emphasize the importance of the magnetic field in optimizing electron dynamics. This trend of increasing efficiency with discharge voltage is seen to varying degrees with all of the thrusters noted in this section. Based on these data, including the subsequent improvements to efficiency seen in other thrusters such as the NASA-300M, the X3 can be expected to operate at total efficiencies on the order of 60% at 300 V discharge voltage and in excess of 65% at 800 V.

The penalty for a switch to krypton propellant is typically a drop in efficiency. Work by Linnell [117, 118] highlighted the efficiency defects present in krypton performance, namely that krypton produces propellant utilization and beam divergence efficiencies that are lower by several percentage points than those produced by xenon. Krypton also has a higher first ionization energy (14.0 eV, as compared to xenon's 12.1 eV [119]), which makes it more costly to ionize. Based on the performance tests of the NASA-300M, NASA-400M, and NASA-457Mv1, the X3 can be expected to provide total efficiencies on krypton propellant of approximately 50-55% at 300 V and closer to 60% around 800–1000 V.

2.6.1.3 Thrust

The expected thrust at a given discharge voltage can be predicted by the expected thrust to power ratio (T/P) at that discharge voltage. Data suggest that T/P is comparable across this family of thrusters at a given discharge voltage. If we re-arrange Equation 2.9 to solve for T/P, then re-arrange Equation 2.5 to solve for thrust and substitute it in, we find the following:

$$T/P = \frac{2\eta}{I_{sp}g}, \tag{2.30}$$

which indicates that T/P goes like the inverse of specific impulse. This is found to be the case with Hall thrusters, which have previously demonstrated T/P values on the order of 65-69 mN/kW at 300 V, around 50 mN/kW at 500 V, and around 35-40 mN/kW at 800 V. Based on these values, we calculated expected thrusts at various operating conditions. We present these in Table 2.1.

2.6.2 Previous X3 Results

These results, originally presented by Florenz [4], were collected during the first firing of the X3. Further details of the work can be found in References [120], [4], and [92]. These

Table 2.1: Expected thrust values for the X3 at various operating conditions.

V_d	P_d	$T_{expected}$
300 V	30 kW	2.0 N
400 V	50 kW	2.8 N
400 V	80 kW	4.5 N
400 V	100 kW	5.6 N
500 V	100 kW	4.1 N
800 V	200 kW	8.0 N

tests were performed in the University of Michigan's Large Vacuum Test Facility, which is described in more detail in Chapter 4. The thruster was operated on xenon propellant off of the same infrastructure (power supplies, flow controllers, etc.) as described below in Chapter 4 as LVTF Configuration 1. The X3 was placed on a modified version of the existing PEPL thrust stand, described elsewhere in the literature [121]. These modifications included the insertion of torsional bearing arms similar to those used in the HMT stand described in Chapter 4. The thrust stand was operated in displacement mode for the performance characterization. Florenz characterized the uncertainty of the thrust measurements at approximately ± 100 mN based on noise in the displacement signal. This rather large uncertainty on the thrust measurement propagates to uncertainties of approximately ± 300 seconds and ± 0.18 on anode specific impulse and anode efficiency, respectively. These error bars were not included in the original work and are suppressed on the plots here for clarity.

The performance parameters, as reported by Florenz [4], are reproduced in Table 2.2. New plots were created of these data and are presented here to better facilitate comparisons between these data and those presented throughout the remainder of this work. First, we present thrust versus discharge power in Figure 2.7. It can be seen that, in general, as discharge power increases, thrust increases. This is to be expected. The only point of note is that it appears as though the IO condition is performing slightly lower than the trendline from the other data would predict.

Anode specific impulse versus discharge power is shown in Figure 2.8. Anode specific impulse should trend upwards with discharge voltage, not power, and thus it would be

Table 2.2: X3 initial performance on xenon propellant at 300 V, as reproduced from Florenz [4].

CHANNEL	$I_{D,total}$	T	I_{sp}	η_a
I	13.1 A	0.25 N	1955 s	0.604
M	31.7 A	0.5 N	1710 s	0.438
O	55.2 A	0.88 N	1850 s	0.484
IM	45.1 A	0.78 N	2277 s	0.642
MO	86.9 A	1.18 N	1919 s	0.426
IO	68.9 A	1.16 N	2220 s	0.612
IMO	100.3 A	1.54 N	2347 s	0.590

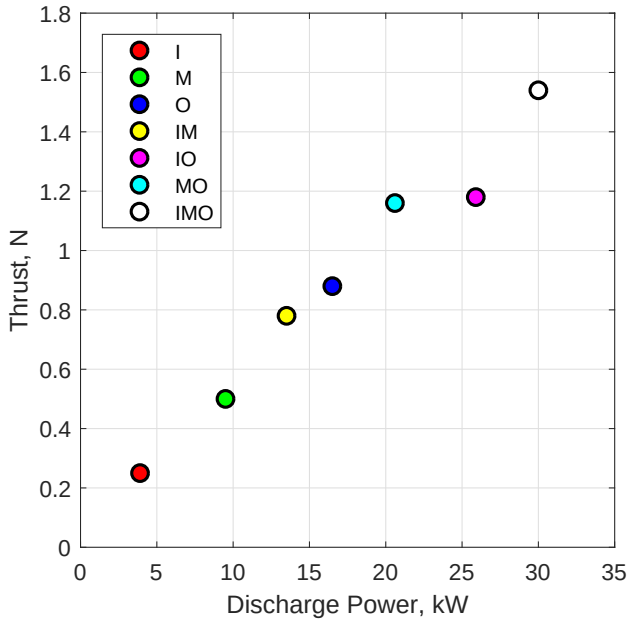

Figure 2.7: Thrust versus discharge power from the X3 initial performance test campaign on xenon. These data are reproduced from Reference [4] but are replotted here.

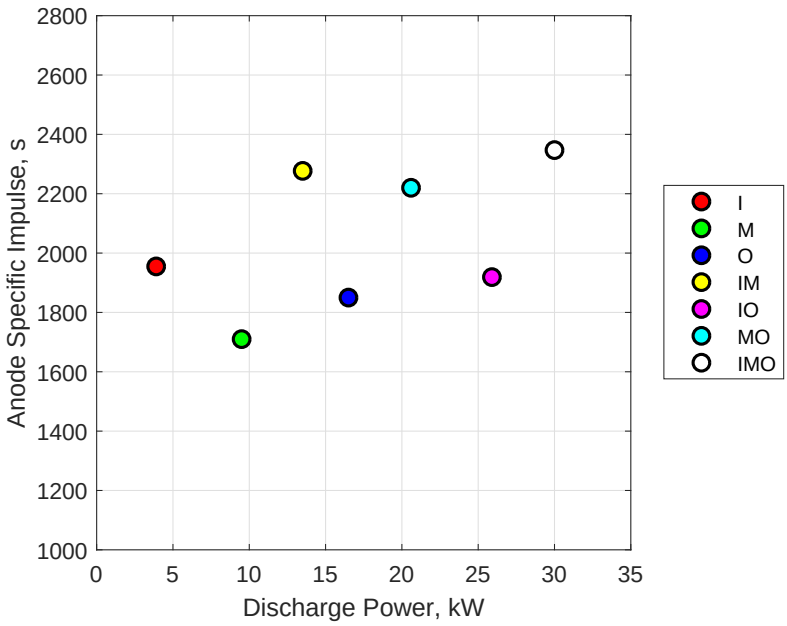

Figure 2.8: Anode specific impulse versus discharge power from the X3 initial performance test campaign on xenon. These data are reproduced from Reference [4] but are replotted here.

expected that the trend on this plot would be flat. In general, this is true: there is no obvious trend with increasing discharge power. However, the spread of these data is rather large, with values ranging over 600 s from the minimum to the maximum. In particular, the IM, MO, and IMO conditions appear to be performing at elevated anode specific impulse as compared to the other four conditions. In fact, these three higher conditions appear to be producing specific impulses that are atypically high for 300 V operation, though the uncertainty in the measurements is large enough that conclusions are difficult to draw.

Finally, anode efficiency is presented in Figure 2.9. Once again, both theory and results from other thrusters [10, 75, 76] suggest that efficiency should in general increase with increasing discharge voltage but have no dependence on discharge power. And once again, no clear trend is seen with increasing discharge power, yet there appear to be three points separated from the remaining four. Here, the I condition has joined with the IM, MO, and IMO conditions at approximately 60% anode efficiency, whereas the M, O, and IO conditions

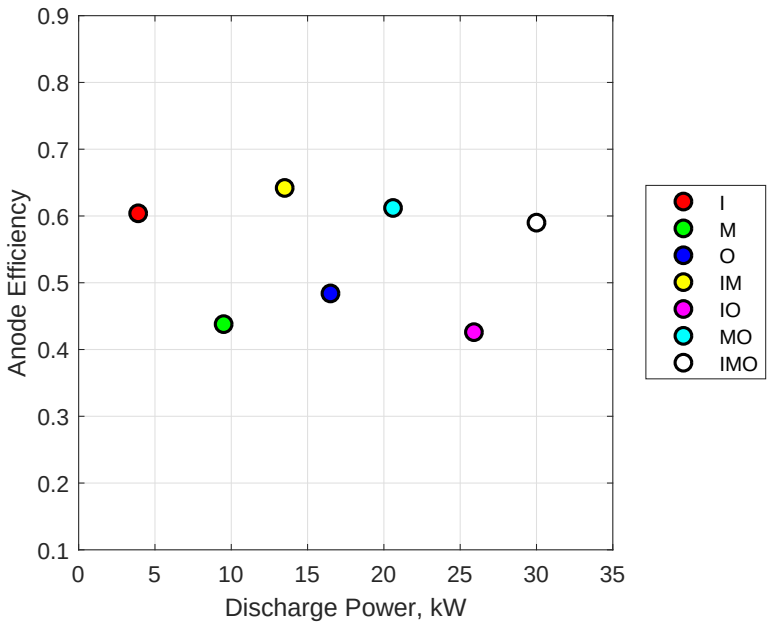

Figure 2.9: Anode efficiency versus discharge power from the X3 initial performance test campaign on xenon. These data are reproduced from Reference [4] but are replotted here.

each perform below 50%. Reasons for this trend are also not clear, along with the reason for the I condition exhibiting relatively low anode specific impulse yet relatively high efficiency. It is also worth noting that these measurements do not vary outside their uncertainty (approximately ± 0.18).

Because of the large uncertainty in the thrust measurements, it is difficult to draw any specific conclusions, especially regarding the performance of certain channel combinations versus others except at the exact throttle points tested. Additionally, the thruster was being fired at a relatively low current density due to facility pumping speed limitations. This test showed the functionality of the X3 in all seven configurations and established that in general it was producing a stable, uniform discharge with discharge current oscillations within the expected range. Further photographs, telemetry, and performance data from these xenon tests are presented in References [4] and [92].

The limited results from this work indicate the possibility that the X3 is, for unidentified reason(s), not performing as expected. However, the data were limited in power level and

both current density and discharge voltage throttling and were collected at relatively high facility background pressures. All of these limitations, plus the indications of off-nominal performance, make the need apparent for a detailed performance study of the X3 with both thrust measurements and plasma plume diagnostics. The results from such a study have implications for high-power Hall thrusters, NHTs, and Hall thruster scaling methodologies.

2.7 Summary

This chapter provided background information for this work. Electric propulsion is a means of attaining propulsion exhaust velocities an order of magnitude larger than those provided by chemical propulsion techniques and more relevant for interplanetary and deep-space exploration. This work specifically focuses on Hall thrusters, which provide attractive performance but are typically considered 1–10 kW devices. By concentrically nesting multiple discharge channels in a single device, higher power can be attained while maintaining relatively compact device footprint. The X3 nested Hall thruster, a 100-kW class, three-channel NHT, has been developed to test the viability of such a scaling technique. This thruster is the focus of this work. Hall thruster performance is typically assessed with an efficiency model that uses plasma diagnostic measurements to study various physical mechanisms within the device. We will use a similar technique in this work, with some key modifications to accommodate the X3. Based on theory and historical values, we expect the X3 to be capable of producing up to 8 N of thrust at 200 kW at an efficiency in excess of 65%.

CHAPTER 3

Motivation

"Critical to deep space exploration will be the development of breakthrough propulsion systems and other advanced technologies."

– President Barack Obama [122]

3.1 Introduction

The X3 represents a substantial advance in the capabilities of Hall thrusters. It is important to understand the motivation for creating such a device and advancing these capabilities beyond state of the art. This chapter explores this motivation in detail. First, we identify the capabilities of >300-kW EP systems through a survey of mission-study literature in Section 3.2. Next in Section 3.3 we detail the benefits that Hall thrusters offer as a propulsion technology for these types of missions, showcasing their superior performance and specific impulse capability. Finally, in Section 3.4 we explore the unique benefits of the NHT configuration for designing multi-hundred-kW missions. We do so by expanding an existing mass and cost model for EP system design to include the traits of NHTs in mass and power density.

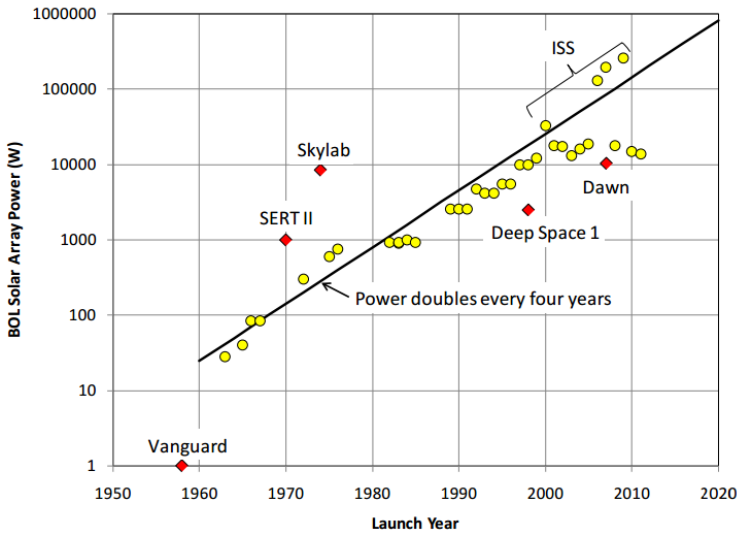

Figure 3.1: The highest-power spacecraft flown every year from 1960 to 2011. This illustrates that on average available on-board power doubles every four years. Reproduced from Reference [5].

3.2 Capabilities of EP Systems in Excess of 300 kW

3.2.1 Progress in On-Board Power

A 500-kW or 1-MW solar array would be very large using current state-of-the-art technology. The International Space Station features solar arrays with an active area of approximately 1680 m^2 that provide on-board power of approximately 260 kW [5]. Based on these values, a 500 kW array of the same technology would be approximately 3230 m^2 and a 1 MW array would be 6460 m^2. These values are only for on-board powers at 1 AU; power drops off as $1/r_{sun}^2$, where r_{sun} is the distance from the sun.

However, solar array technology is continually improving at a rapid pace. Estimates based on the highest-powered spacecraft launched in each year since 1959 (the year the Vanguard spacecraft launched with 1 W of on-board solar power) show that the available power from solar arrays approximately doubles every 4 years [5]. Figure 3.1 reproduces a figure from Reference [5] illustrating this trend. In that reference, Brophy points out that the 1-kW SERT II spacecraft launched in 1970 was well above the trendline, yet the more recent Deep

Space 1 and Dawn missions fell well below the trendline. He argues that this illustrates that the capabilities of solar arrays have caught up to the demands of electric propulsion. With available power doubling every four years and solar array efficiencies improving as well, discussing a 500 kW or even 1 MW solar array is not unreasonable when discussing missions in the next few decades. However, because of the $1/r_{sun}$ dependence of solar power, these values decrease significantly as a spacecraft flies toward destinations such as Mars or Jupiter.

Nuclear reactors are often used in conceptual mission studies looking farther into the future [123, 124]. Estimates suggest that nuclear reactors would likely be able to provide power levels in the range of 1–5 MW. However, as of this writing, development of nuclear reactors for space applications has been limited, and only a single nuclear-powered spacecraft has been launched by the United States. That reactor, the SNAP-10A, ended after producing 590 W of power due to the failure of a voltage regulator after its launch in 1965 [125]. A number of Russian satellites have been launched with nuclear power with greater success [124]. More recently, NASA had a short-lived program aimed at advancing the state of the art of nuclear-powered electric propulsion. This program, called Project Prometheus [126], was funded from 2003 to 2006. It sought to develop a 100-kW flagship science mission and progressed as far as developing the Jupiter Ice Moon Orbiter concept before funding was cut. Very recently, the Kilopower program, a collaboration between NASA GRC and Los Alamos National Laboratory, has been working to advance the state of the art of fission reactors for space applications. They are developing 1–10 kW reactors using both thermoelectric and Stirling technology [127, 128] and are preparing for a 28-hour full-power demonstration of an engineering unit [1].

Thus, both solar panel and nuclear reactor technology are continuing to advance to higher power capability. A set of 10-kW Kilopower-like reactors could provide constant power for a high-power EP mission regardless of the destination. High-power, state of the art solar panels could also be used, though they would require throttleability of the propulsion system as

[1] See for instance: https://www.nasa.gov/directorates/spacetech/feature/Powering_Up_NASA_Human_Reach_for_the_Red_Planet. Accessed January 1, 2018.

power decreased. Both of these technologies show promising progress toward being capable of powering electric propulsion missions in excess of 300 kW.

3.2.2 Possible Destinations

A number of mission studies that have been produced for EP systems with power levels at or exceeding 500 kW are explored below. These types of missions illustrate the various ways that EP systems can directly impact mission design. The tenants of these missions can be applied, in some fashion, to a wide variety of other destinations and mission types not discussed here or explored widely in the literature.

A number of these mission studies follow what is called a hybrid approach as part of NASA's Electric Path architecture [129]. This architecture, originally developed by Strange at JPL, combines a chemical propulsion system and a solar electric propulsion (SEP) system on both cargo and crewed vehicles. This hybrid approach uses the SEP system where possible to maximize the mass savings but relies on the chemical system when necessary to enact large orbital maneuvers quickly. As we will show below, the hybrid approach has been demonstrated as applicable for many destinations and to provide significant mass savings with little to no increase in trip times over a chemical-only system.

3.2.2.1 Near-earth Asteroids

Near-earth asteroids (NEAs) are of interest for robotic or human missions for a number of reasons. These bodies vary widely in composition and can be of significant scientific merit. A Discovery-class mission was recently awarded to a team looking to use a string of 5-kW Hall thrusters to explore the large asteroid Psyche, which is suspected to be the core of a small planet that failed to form during the early development of the solar system [130]. This mission looks to use unmanned exploration of Psyche to understand metal cores and explore a type of body unique within our solar system. Psyche is just one specific example of the variety of NEAs and the potential science questions that could be answered through

exploration of these bodies. Additionally, there is strong interest in the potential of asteroid mining as a means of capturing the resources found within NEAs [131].

A study by Brophy [5] explores using a 300-kW solar electric propulsion system for the human exploration of NEAs. The specific mission studied begins with the launch of the Deep Space Habitat (DSH), the Space Exploration Vehicle (SEV), and a 300-kW electric propulsion system called the SEP Freighter. The Freighter moves the DSH and SEV out from the low-Earth orbit (LEO) in which they were launched via a typical low-thrust spiral. Once they are clear of the Van Allen radiation belt, the human crew launches to join with the vehicle. A chemical propulsion stage provides an initial impulsive Earth-escape burn. This chemical stage is then jettisoned, and the Freighter provides a continuous burn for the remainder of the 160 day trip to the NEA (a generic "hard-to-reach" NEA in the study). After a 30 day stay at the NEA, the crew returns to Earth in a 200 day trip.

This study found that an EP system on the order of 300 kW would reduce the number of heavy-lift launches needed for this mission by a factor of two. The mission was made capable by a system of 40-kW Hall thrusters operating at a specific impulse of 2000 seconds. This mission study was designed to minimize the cost of the system, and the authors prioritized cost reduction over mass reduction. This is part of the reason why a 300-kW system of 40-kW thrusters was selected: it represented a balance between development costs and mission benefits. The details of these types of trades are presented below in Section 3.4. As the development of high-power propulsion devices and solar arrays continues, it can easily be seen that the benefits to a mission of this type would increase with increasing system power. A higher system power at similar specific impulses and efficiencies would result in either reduced trip times, increased payload capabilities, or both.

3.2.2.2 Phobos and Deimos

Phobos and Deimos, Mars's two moons, provide enticing targets for possible human missions, especially as first steps in a larger Mars mission program. A 2009 report by the Advisory

Committee on the Future of the U.S. Space Program [132] concluded that "Mars is the ultimate destination for human exploration; but it is not the best first destination." Phobos and Deimos provide two potential better first destinations. Mars's moons provide a number of benefits to a larger human Mars mission architecture. They potentially house valuable resources such as water, and they provide a representative intermediate step for a Mars mission program, allowing a spacecraft to travel the full distance to Mars without traveling into, and requiring a trip out of, Mars's deep gravity well.

A base established on one of the moons could serve a number of purposes. It could function as a mission control for missions to Mars's surface, which would circumvent the long communication delay with a mission control team on Earth's surface, providing faster response in the event of a problem and even providing means for a quick human rescue team if needed. It could also serve as a quarantine way station, which would help prevent contamination from equipment incoming from Earth as well as from samples returning to Earth.

A study by Strange [133] investigated the benefits that a high-power solar electric propulsion system could provide to a human mission to Phobos or Deimos. This study applied the hybrid approach detailed above to two mission scenarios. In the first, a SEP system is used to pre-position cargo and the crew follows behind in a deep space habitat propelled by a chemical system. In the second, a SEP system is used for both cargo and crew transport. In both scenarios, a 300-kW system was used for cargo transport, firing at 3000 seconds specific impulse. For the crew system, where reducing trip time is of increased importance, a 700-kW system firing at 1800 seconds specific impulse is used. Both mission scenarios are designed for high total mass transport (69.2 t of cargo) to the moons. The mission takes two Deep Space Habitat units and leaves one behind on the surface of the moon for future missions (the beginning of a base). There is capacity for crew supplies for a mission duration of three years.

The study found that using 300-kW SEP for cargo and chemical propulsion for the crew

provided a mass savings of 44% to 46% over an all-chemical system. And by employing the 700-kW SEP system in place of the chemical system for crew transport, another 13% to 26% was saved (the values vary depending on which launch window was selected). There is a trade-off, however, in time on the surface: whereas the chemical-propelled crew vehicle allowed for stays of 1–1.5 years, the longer trip times of the SEP-propelled crew vehicle drop those stay times to 4–6 months. Depending on mission goals, the mass savings may outweigh the decreased stay time. The total mission times remain similar.

3.2.2.3 Mars

The same 2009 report by the Advisory Committee on the Future of the U.S. Space Program [132] also noted that "Mars is unquestionably the most scientifically interesting destination in the inner solar system," an opinion prevalent long before the report was published. Substantive Mars mission concepts have been published since at least 1953. It was in that year, 16 years before humans first landed on Earth's moon, that Werner Von Braun published his vision of a many-crewed, all-chemical, multi-ship mission to Mars in *The Mars Project* [134].

Since then, Mars has been a perpetual goal of human spaceflight [135]. Much of the buildup of infrastructure for electric propulsion testing—infrastructure still relied on today, even for the testing detailed below—was completed in the late 1950s and early 1960s because at that time it was expected that ships propelled by nuclear electric propulsion systems would be an integral part of human missions to Mars in the 1970s [136]. There have been book-length mission plans published [137], as well as a science fiction novel based around a nuclear-electric mission [7].

Thus, a large body of work exists on missions to Mars. For our purposes, these concepts can be simply split into two categories: those with all-chemical propulsion systems and those that incorporate, in some fashion, electric propulsion. These latter concepts vary widely in power level and scope, from low-power concepts involving as few as 8 kW of EP [29] to

massive multi-megawatt missions [123].

A number of studies investigate the benefits of using SEP systems at power levels on the order of a couple hundred kilowatts and below [29, 138–140]. Many of these lower-power studies investigate the use of an SEP system for cargo transport only, leaving the crew transport to a separate chemically-propelled vehicle. The reason for this is that at lower powers, the mass benefits of a SEP system are outweighed by the longer trip times that a lower-power system necessitates. A number of these studies, including those by Williams [29] and by Price [139], look to utilize established technology (for Williams, NASA's NSTAR ion thruster [141], and for Price, NASA's HERMeS Hall thruster [142]). Even at these lower power levels, the benefits of using SEP are clear. In Price's study, two 100-kW SEP cargo tugs firing at 3000 seconds specific impulse are used to pre-position a habitat and other equipment at Phobos, which then allows for a crew to follow behind for a 950-day mission. These two 100-kW tugs can easily be combined into a single 200-kW unit, and scaling beyond those power levels only increases the available payload mass.

A study by Myers illustrates the benefits of using higher-power SEP systems for cargo transport [31]. The architecture proposed in that study uses two cargo tugs: one with 300 kW of SEP propulsion and one with 600 kW. Using a SEP system for cargo provides an improvement in initial mass launched to low-Earth orbit of 46% for direct return missions and 55% for missions returning to LEO. Doubling the power level of the propulsion system doubles the payload capability of the ship, as expected. The 600-kW cargo tug envisioned by Myers would utilize five 150-kW thrusters firing at 3000 seconds specific impulse. This system would use approximately 65 t of xenon (though the article points out that krypton would serve as well) to ferry 70 t-class payloads to Mars. Payloads of this class include Mars descent and ascent vehicles, 4-crew Mars habitats, and surface exploration equipment: the type of equipment necessary for long-duration human missions. Around power levels of 300 kW, application of SEP systems to crewed capsules becomes reasonable. One such study by Strange [129] showed that an SEP system on the order of 250 kW, operating at 3000 seconds

specific impulse and an efficiency of about 60% (for a total thrust of approximately 15 N) would enable missions to Mars lasting a total time of 1 year. By raising the system power to 500 kW, the mission time can extend to 3 years.

A study by Gefert [143] from 1999 illustrated how improvements in areas outside propulsion allowed for substantial reductions in necessary propulsion system power. Many previous studies [144–147] had concluded that power levels on the order of 8 MW were necessary to provide reasonable trip times for all-electric piloted missions from low-Earth orbit to Mars capture orbit. However, Gefert concluded that a combination of a hybrid chemical/electric propulsion architecture (identical in nature to the hybrid approach recently proposed by Strange and others at JPL that was mentioned above), aerobrake technology, and cryogenic propellant storage would drop a mission's power needs closer to 0.5-1 MW. The study further concluded that the mission performance of a 800-kW system operating at 2000 seconds specific impulse and that of a 1.9-MW system operating at 4000 seconds specific impulse were very similar, showcasing that for Mars missions, lower specific impulse is desired to reduce trip times.

Another study by Chiravalle [124] investigated a Mars mission architecture and found that, in terms of trip times and mass fraction, a 3800-second specific impulse system always outperformed an 8000-second system. To obtain transfer times less than one year, a 3800-second system at 1.5 MW was necessary. That same system, scaled down to a power level of only 500 kW, would provide 18-month transfer times. The 1.5 MW system would be necessary to keep transfer times low enough for crewed missions, but an 18 month transit time for a large cargo tug is completely reasonable within certain mission architectures.

Finally, a study by Schmidt [24] studied a mission architecture that involved a 145-day outbound trip to Mars and stays ranging from 30 to 120 days. The study concluded that an electric propulsion system operating at 3000 seconds specific impulse and providing 100 N of thrust would enable the mission. Schmidt assumed a system efficiency of 37%, which put the approximate power necessary for the mission at 4 MW. Improving the system efficiency to

60%, commensurate with EP technologies such as ion and Hall thrusters, drops the necessary power level to 2.5 MW.

A number of other studies have investigated power levels beyond those discussed here. One by Braun [123] traded between four approaches, including an all-chemical system, a nuclear-thermal system, a single capsule with both chemical and electric propulsion systems (once again, the hybrid approach), and two separate vehicles, one with a chemical system and one with an electric system. This study used a 4-MW, 3000-second specific impulse argon ion thruster system. It found that no system clearly dominated the others but that the best system depended on exact mission structure. Another study by Keaton [148] analyzed the benefits of using a 150-N EP system at 3000-seconds specific impulse and found that it could complete a given Mars mission profile with 40% less mass delivered to LEO than required of an all-chemical system.

Thus, the literature shows that electric propulsion systems benefit Mars missions at powers as low as 8 kW, with those benefits increasing with power. Both cargo tugs and crew capsules with electric propulsion systems on the order of 600 kW or higher provide substantial benefits to Mars mission architectures over all-chemical alternatives.

3.2.2.4 Deep Space

Deep space is a realm mostly left to lower-power spacecraft at the moment. The success of Deep Space 1 [149] and Dawn [150], which both used electric propulsion systems of less than 10 kW, has demonstrated that for Discovery-class missions and small probes, the use of low-power, high-specific impulse propulsion provides substantial mass savings. At the moment, the focus of NASA's Human Exploration and Operations Mission Directorate is on large-payload human or cargo missions to places mentioned above and for continued robotic (i.e., small-payload) exploration of targets beyond Mars.

However, human missions aside, high-power EP systems would provide the capability of large sample return missions to deep-space objects. There are many bodies beyond Mars of

considerable scientific interest, including comets, asteroids, planets, and moons. A number of the moons of other planets in the Solar System, including Jupiter's Europa [151] and Saturn's Titan [152] and Enceladus [153] show promising signs of potential water oceans under crusts of ice. These types of features hold the possibility of containing life and thus are of significant scientific interest.

The same study by Chiravalle mentioned above [124] also investigated two deep space missions for high-power EP. One mission studied transferring a 10 t payload from Earth to Jupiter's moon Europa, and the other studied a similar-sized mission to Saturn's moon Titan. These studies investigated only 8000-second specific impulse systems and did not do trades on possible benefits of lower specific impulses. However, both missions displayed similar power requirements. For the Europa mission, 1.1 MW of power would provide a trip time of 54 months, and dropping the power to 232 kW extended the trip time to 72 months. For the mission to Titan, similar results showed that a system power level of 417 kW produced a trip time of 84 months, and reducing the power level to 145 kW extended the trip to 114 months. Nuclear power would be required to make these power levels practical in deep space.

Though studies are currently limited, these results demonstrate the potential of high-power EP systems as applied to deep space missions. As the technology, and the goals of space exploration, continue to develop, these missions will be treated with greater variety and attention, and further advantages of high-power EP systems will likely be uncovered.

3.3 Benefits of Hall Thrusters

We have demonstrated the capabilities and benefits of EP systems at power levels on the order of several-hundred kilowatts and above. Next, we turn to the question of which EP technology is best suited for these types of missions. A wide range of factors drive this selection, and depending on the relative importance given to various factors (e.g., putting

emphasis on cost savings at the expense of mass savings [5]), different technologies may be selected. Here, we detail the various factors that affect the selection of an EP technology, then analyze leading candidate technologies for these high-power, high-specific impulse missions.

Factors that can drive EP technology selection include, in no particular order:

- **Technology maturity** or technology readiness level. This encompasses flight history and relative development of not only the thruster technology but of the supporting technology needed in flight such as power processing units, propellant flow systems, gimbals, etc. This has a direct effect on both mission risk and mission cost, as technologies with lower maturity require increased development costs to bring them to flight qualification.

- **Performance**, including efficiency and specific impulse range. This directly impacts the mass savings and trip times achievable for a given mission profile.

- **Scalability** to relevant powers while maintaining performance.

- **Specific mass** (typically given in units of kg/kW). Payload mass savings can be adversely effected by overly heavy propulsion systems.

- **Lifetime capability**. Thrusters must be able to operate for long enough to complete a mission. Some electric propulsion mission architectures call for system re-useability to limit costs, which extends the lifetime demand beyond a single trip.

- **Footprint or thrust density**. Although spacecraft concepts for high-power missions typically have ample space for propulsion systems [6], minimizing propulsion system footprint is often still desirable to limit plasma plume size and possible spacecraft impingement.

- **System complexity**. Some electric propulsion concepts rely on exotic technologies such as super-conducting magnets. Others rely on high currents or voltages for operation. Complexity can increase risk and possible failure modes for a system.

- **Ground test capabilities**. Ground testing certain technologies in a relative environment can be difficult or impossible. Proper ground testing is key to proper flight qualification.

Of the available electric propulsion technologies, a number can be eliminated immediately. Resistojets and arcjets simply do not provide specific impulses high enough for the missions detailed above. Resistojets typically provide specific impulses of 200-350 seconds and arcjets have demonstrated a maximum of 1500 seconds on hydrogen propellant [17].

There are a number of technologies that show promise but that are underdeveloped when compared to other options. These include the Variable-Specific Impulse Magnetoplasma Rocket (VASIMR) [94], which has gone through a significant amount of theoretical development but relies on superconducting magnets [154] and has been plagued by ground test issues that limit its operation to "shots" less than 30 seconds in duration followed by substantial cooling periods [95], leaving significant questions regarding its feasibility and lifetime. Also in this category are field-reversed configuration (FRC) thrusters [155,156]. These thrusters rely on the development of plasmoids that are ejected at high-frequency from the device to generate thrust. Early experiments have been promising, but a significant amount of work is necessary to advance the technology readiness level of FRC thrusters.

Magnetoplasmadynamic (MPD) thrusters have a number of attractive qualities. They can operate at specific impulses of 2000-5000 seconds [157], well-suited for the missions detailed above. They have excellent thrust density [17], which translates to a small footprint for a MW-level system, and they have been demonstrated in the laboratory at multi-megawatt power levels [158]. However, MPDs do not operate at the efficiencies that technologies such as Hall and ion thrusters do, typically operating at efficiencies less than 0.25 for argon and less than 0.10 for xenon, though efficiencies up to 0.69 have been measured on lithium propellant [159]. The multi-megawatt MPD in Reference [158] demonstrated 0.38 efficiency at 4000 seconds specific impulse on argon. Additionally, MPD thrusters lack the lifetime capability necessary for months-long or years-long mission durations. Lifetime assessments indicated

that MPD thrusters were two to three orders of magnitude below what was necessary for high-power missions to Mars [160].

Pulsed Inductive Thrusters (PITs) offer another potential EP technology. PITs as large as 1 meter in diameter have been demonstrated [161], and TRW, Inc., owns a patent for a mutli-megawatt PIT design [162]. However, PITs lack both the development and performance necessary for these types of missions. Thrusters have demonstrated large thrust values and specific impulses in the 2000-7000 second range, but have only achieved efficiencies up to 0.50 [163, 164].

Unlike most of the device types detailed above, ion thrusters have a significant, successful flight heritage and performance that lends itself well to the success of the missions detailed above. Beginning with the flight of the SERT I spacecraft in 1964 (an in-depth summary of which is presented in Reference [20]), which demonstrated 20 mN of thrust over 31 minutes of operation to become the first in-space demonstration of an EP device, ion thrusters have reliably demonstrated specific impulses from 3000–8000 seconds and efficiencies in excess of 0.70 [2]. NASA's NSTAR thruster has demonstrated lifetimes of 30,000 hours in ground tests [165], 16,000 hours on the Deep Space 1 mission [28], and 46,000 hours on the Dawn mission [22]. The next-generation NEXT ion thruster demonstrated 50,170 hours of operation during a long-duration ground test that spanned a number of years [23]. In 1967, GRC (at the time, Lewis Research Center) demonstrated a 200-kW ion thruster on mercury propellant at specific impulses of 7000–8600 seconds [166] [2].

One mark against the ion thruster in the applications explored here is that it actually operates at too high a specific impulse for some of the larger-payload human mission concepts: as discussed, specific impulses closer to 2000 seconds help to minimize transit time, which is critical for human missions. Ion thrusters do not perform as efficiently at these low specific impulses [167]. Additionally, ion thrusters are typically operated at low thrust densities to extend their lifetimes. The NEXT thruster is a 7 kW device capable of producing 237 mN

[2] This thruster is still on display at GRC and is shown next to the X3 in the front matter of this document.

at its highest-thrust throttling point [23] and is 40 cm in diameter [168]. As a point of comparison, the H6MS, a 6-kW long-life Hall thruster, is approximately 20 cm in diameter and produces 385 mN at its nominal condition of 300 V, 20 A [41].

This then brings the discussion to Hall thrusters. As discussed in Chapter 2, Hall thrusters are capable of specific impulses between 1500 and 3500 seconds on xenon propellant, with values closer to 5000 seconds attainable with krypton. They are capable of 0.60 total efficiency or higher across this range of specific impulse. Traditionally, Hall thrusters demonstrate discharge channel erosion due to plasma impingement that limits lifetimes to less than 10,000 hours of operation [2]. However, magnetic shielding has shown promise in extending these lifetimes significantly, with a target of 10^5 hours [40, 41, 169].

The Hall thruster's high efficiency and high thrust-to-power ratio at moderate specific impulses (typically on the order of 65 mN/kW at 300 V discharge voltage), coupled with its substantial flight heritage, long lifetime capabilities, and demonstrated scalability to high power levels, makes it an ideal candidate for the types of missions outlined above.

3.4 Benefits of Nested Hall Thrusters

We have shown that the Hall thruster is a promising technology for a wide variety of missions. We next turn to the design and makeup of the propulsion system. That is: how many thrusters, of what power level, should comprise the propulsion system? Fault tolerance is an important driving factor of the propulsion string composition. One downside to electric propulsion systems is that the high-voltage nature of their operation gives rise to a number of potential failure states. Redundancy for fault protection must therefore extend beyond the thruster itself to the entire "string" (power processing unit, thruster, and cathode). Typically, robotic missions are designed to be single-fault tolerant, whereas human missions need double- or triple-fault tolerance [6].

This means that a single 500-kW thruster is not enough for a 500-kW mission because

it is not fault tolerant, and that in all likelihood adding a second or third 500-kW string would be overly massive. But what offers a better solution: three 250-kW thrusters? Eleven 50-kW thrusters? Five hundred and one 1-kW thrusters? This type of decision is a complex one that, much like the decision process for the thruster technology itself, depends strongly on the relative importance given to various factors. A model by Hofer and Randolph [6] was developed to analyze this problem in great detail. Their work, described briefly below, focused on single-channel Hall thrusters (SHTs). Subsequently, we expanded their model to capture NHTs to explore their impact on the full propulsion system design optimization.

3.4.1 Description of Electric Propulsion System Mass and Cost Model

The Hofer-Randolph model was designed to study the optimum number of thrusters for a given total mission power requirement. To do so, the model consisted of two main equations: one for system mass and one for system cost. Though the model was tuned to focus on SHTs, it was designed to be applicable to any type of EP technology with appropriate input data. The model captured the entire EP string, including the thruster, gimbal, xenon flow system (XFS), propellant management assembly (PMA), power processing unit (PPU), xenon tank, and cabling. An image of the system, reproduced from Hofer and Randolph [6], is shown in Figure 3.2. For a given total mission power requirement, the electric propulsion string can consist of a single high-power thruster or an abundance of lower-power thrusters. The model predicted the mass of a BPT-4000 flight string designed for a NASA science mission to within 3%, demonstrating its fidelity.

As discussed by Hofer and Randolph, the nature of electric propulsion and its typical failure mechanisms (high-voltage issues including grounding and arcing, as well as thermal issues) have led the community to adopt a fault tolerance scheme involving the inclusion of one, two, or three redundant thruster strings corresponding to single-, double-, or triple-fault tolerance. For example, the ion propulsion system for the Dawn mission was in a "2 plus 1"

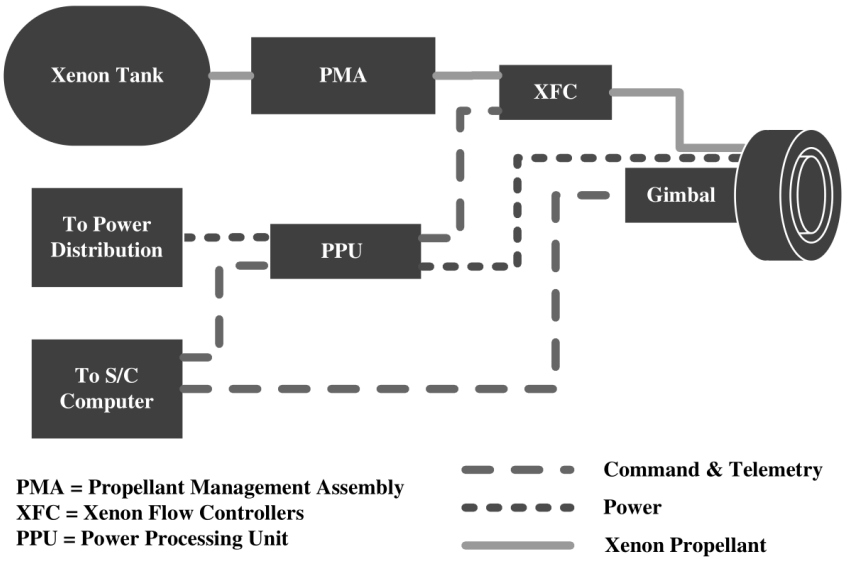

Figure 3.2: The major elements of a Hall thruster propulsion system. Figure reproduced from Reference [6].

configuration, in which there were two active ion thrusters and one redundant thruster [150]. The Hofer-Randolph model can capture an arbitrary fault-tolerance level, which allows for study of the effects of carrying redundant thrusters on system optimization.

The mass equation is a combination of mass estimations of the various system components. Each of these components has a scaling equation that incorporates (typically linear) coefficients established by the authors. These coefficients were gathered from a number of sources, including scientific literature, flight programs, and industry sources. The global mass equation takes the form:

$$M_{sys} = (1 + f_s)\left(M_{tg} + M_{PPU} + M_{XFS} + M_{cab} + M_{tank}\right), \tag{3.1}$$

where M_{sys} is system mass, f_s is the structural weighting factor, M_{tg} is the mass of the thruster and gimbal, M_{PPU} is the mass of the power processing unit, M_{XFS} is the mass of the xenon flow system, M_{cab} is the mass of cabling, and M_{tank} is the mass of the propellant tankage. This can be expanded to:

$$M_{sys} = (1 + f_s)\left\{(N_{tot})\left[\left(A_{tg} + A_{PPU} + A_{cab}\right)\left(\frac{P_{sys}}{N_{ac}}\right) + B_{tg} + B_{PPU} + B_{XFS} + B_{cab}\right] + C_{XFS} + (1 + f_{tank})\, I_{Xe} P_{sys}\right\}, \tag{3.2}$$

where A, B, C, and I are all linear regression coefficients for the various sub-systems, N_{tot} is the total number of thrusters, P_{sys} is the total system power, and N_{ac} is the number of active thrusters. The various sub-system equations take the following forms. Thruster and gimbal:

$$\begin{aligned} M_{tg} &= M_{th} + M_{gim} \\ &= \left(1 + f_{gim} D_{gim}\right)(N_{ac} + N_{rd})\left[A_{th}\left(\frac{P_{sys}}{N_{ac}}\right) + B_{th}\right] \\ &= (N_{ac} + N_{rd})\left[A_{tg}\left(\frac{P_{sys}}{N_{ac}}\right) + B_{tg}\right], \end{aligned} \tag{3.3}$$

where M_{th} is the mass of the thruster, M_{gim} is the mass of the gimbal, D is a regression coefficient, and N_{rd} is the number of redundant thrusters. Power processing unit:

$$M_{PPU} = (N_{ac} + N_{rd})\left[A_{PPU}\left(\frac{P_{sys}}{N_{ac}}\right) + B_{PPU}\right], \tag{3.4}$$

cabling:

$$M_{cab} = (N_{ac} + N_{rd})\left[A_{cab}\left(\frac{P_{sys}}{N_{ac}}\right) + B_{cab}\right], \tag{3.5}$$

xenon flow system:

$$M_{XFS} = (N_{ac} + N_{rd})\, B_{XFS} + C_{XFS}, \tag{3.6}$$

xenon propellant:

Table 3.1: Mass model coefficients used in this work. These values come from the work of Hofer and Randolph [6].

Name (subscript)	A [kW/kg/string]	B [kg/string]	C [kg]	D [-]	f [-]	I [kg/kW]
Thruster (th)	Variable	0	—	—	—	—
Gimbal (gim)	—	—	—	0.5	—	—
PPU (ppu)	1.7419	4.654	—	—	—	—
XFS (xfs)	—	3.2412	4.5189	—	—	—
Cabling (cab)	0.06778	0.7301	—	—	—	—
Structural (s)	—	—	—	—	0.26	—
Tankage (tank)	—	—	—	—	0.04	—
Xenon (Xe)	—	—	—	—	—	100

$$M_{Xe} = I_{Xe} P_{sys}, \tag{3.7}$$

and tankage:

$$\begin{aligned} M_{tank} &= f_{tank} M_{Xe} \\ &= f_{tank} I_{Xe} P_{sys}. \end{aligned} \tag{3.8}$$

For each of the subsystem components above, linear scaling is used as a first order approximation. Subsystems either scale with the power of each thruster (thruster, PPU), the number of strings (XFS, cabling), the total system power (xenon), or the total propellant mass (tankage). Except for the changes detailed below, our work here used the same nominal coefficients used by Hofer and Randolph. For reference, these coefficients are presented in Table 3.1.

It is important to note that Hofer and Randolph's thruster specific mass coefficient, A_{th}, was determined by fitting a line to a set of masses and nominal powers for single-channel Hall thrusters. By doing so, the authors implicitly impose certain thruster characteristics: namely, nominal Hall thruster current densities and discharge voltages of 300-400 V. To

appropriately compare results, these characteristics must be imposed on the NHT designs we add here as well. This puts arbitrary limits on thrusters and does not capture the entire performance envelope of Hall thrusters. To use the X3 as an example, at 300 V and a nominal discharge current density the X3 is an 80 kW thruster, implying that three X3s would be necessary for a 200-kW mission. Yet the X3 has a 200-kW throttle point at 800 V discharge voltage, meaning that a single X3 has the potential to serve the needs of a 200-kW mission.

3.4.2 Modifications to Mass Model to Capture NHTs

There are three terms in the Hofer-Randolph model that we change to accommodate NHTs: the cabling, xenon flow system, and thruster specific mass. The cabling and XFS terms are modified such that they scale per channel, not per thruster, and the thruster specific mass is modified to vary with the number of channels. The PPU equation is not changed because the discharge units for high-power SHTs (the source of the bulk of the mass and cost in the PPU) are likely to be made up of lower-power (10-20 kW) sub-units. For a NHT at the same power, these would simply be distributed across the N_{ch} channels. The cabling required by this is captured in the M_{cab} term. It is expected that there will be a small increase in PPU mass with increasing number of channels due to the fact that each channel requires its own magnet power supplies, but these will remain less than 1 kW each and thus a small fraction of the discharge power. Additionally, because NHTs are capable of operating off of a single centrally-mounted cathode [1, 4], the cathode heater and keeper requirements will be identical for a given thruster power level, regardless of the number of channels. Thus, we do not modify the PPU mass scaling.

3.4.2.1 Cabling and XFS

The equations for cabling and the xenon flow system are modified such that they are multiplied by the number of channels in each thruster (N_{ch}):

$$M_{cab} = N_{ch}\left(N_{ac} + N_{rd}\right)\left[A_{cab}\left(\frac{P_{sys}}{N_{ac}}\right) + B_{cab}\right], \tag{3.9}$$

and

$$M_{XFS} = N_{ch}\left(N_{ac} + N_{rd}\right)B_{XFS} + C_{XFS}. \tag{3.10}$$

This accounts for the fact that each active channel requires cabling and xenon mass flow control. With these modifications, the global mass model equation becomes:

$$\begin{aligned} M_{sys} = (1 + f_s)\Bigg\{(N_{tot})\Bigg[\Big(A_{tg} + A_{PPU} + N_{ch}A_{cab}\Big)\left(\frac{P_{sys}}{N_{ac}}\right) + B_{tg} + B_{PPU} \\ + N_{ch}B_{XFS} + B_{cab}\Bigg] + C_{XFS} + (1 + f_{tank})\, I_{Xe}P_{sys}\Bigg\}. \end{aligned} \tag{3.11}$$

3.4.2.2 Thruster Specific Mass: General Expression

For an NHT, we denote total thruster power as the sum of the discharge powers of each channel:

$$P_{th} = \sum_{i=1}^{N_{ch}} V_{D,i}\mathcal{I}_i, \tag{3.12}$$

where $V_{D,i}$ is the discharge voltage, and $\mathcal{I}_i$ is the current for the i^{th} channel. The current for each channel can be expressed as

$$\mathcal{I}_i = j_i\mathcal{A}_i, \tag{3.13}$$

where j_i is current density and $\mathcal{A}_i$ the exit area of the i^{th} channel. A Hall thruster channel is an annulus, the area of which is

$$\mathcal{A}_i = \pi d_i w_i, \tag{3.14}$$

where the characteristic dimensions of the channel are the mean diameter d_i and the channel width w_i. Equations 3.14 and 3.13 then can be substituted into Equation 3.12 to yield

$$P_{th} = \pi \sum_{i=1}^{N_{ch}} V_{D,i} j_i d_i w_i. \tag{3.15}$$

Thruster mass can be expressed as:

$$m = \rho\pi \left(\frac{d_{N_{ch}}}{2} + \frac{w_{N_{ch}} + w_{pole,N_{ch}}}{2} \right)^2 \Delta z, \tag{3.16}$$

where ρ is the thruster volumetric mass density, $d_{N_{ch}}$ is the diameter of the outermost channel, $w_{N_{ch}}$ is the width of the outermost channel, $w_{pole,N_{ch}}$ is the width of the outer channel's poles, and Δz is the thickness of the thruster. This expression becomes

$$m = \frac{\pi}{4} \rho \Delta z \left(d_{N_{ch}} + w_{N_{ch}} + w_{pole,N_{ch}} \right)^2. \tag{3.17}$$

Equations 3.15 and 3.17 then can be combined to create an expression for the NHT specific mass:

$$A_{th,NHT} = \frac{\rho \Delta z \left(d_{N_{ch}} + w_{N_{ch}} + w_{pole,N_{ch}} \right)^2}{4 \sum_{i=1}^{N_{ch}} V_{D,i} j_i d_i w_i}. \tag{3.18}$$

Equation 3.18 serves as a general analytical expression for the NHT geometry with an arbitrary number of channels, where no particular thruster design assumptions have been made.

3.4.2.3 Thruster Specific Mass: Example NHT Nesting Technique

A set of thruster design assumptions can be applied to Equation 3.18 to provide quantitative values of $A_{th,NHT}$ for use in the mass model. As discussed above, a discharge voltage of 300 V and a typical Hall thruster current density must be selected such that comparisons to the results of the original work are valid. Discharge voltage and current density is assumed to be constant across all channels (which matches the vast majority of NHT operation to date)

and can be pulled outside of the summation in Equation 3.18. Next, we design an example NHT channel nesting technique. First, we assume that the thruster front magnetic pole pieces, which must be accounted for to scale the channels such that they do not overlap, are approximately one-half the width of the channel to which they are adjacent. This is seen as a reasonable approximation: Hall thruster scaling work has demonstrated that maintaining appropriate magnetic field strengths on channel centerline is key to proper thruster operation [11], a task that will require more magnetic circuit material as channel width increases. Thus, as a first-order approximation, channel width and pole piece size scale together. Additionally, it is assumed that pole pieces between nested channels are shared but are twice as wide as those associated with a single channel. This approximately agrees with the designs of both the X2 and the X3.

There are many possible schemes for nesting the channels of a NHT. Thruster power goes with discharge channel exit area, so a design goal to maximize the benefit of the NHT configuration is to maximize the amount of exit area contained within a given thruster outer diameter. However, the magnetic pole pieces inhibit this. Ultimately the nesting technique selected must balance engineering practicality with optimized geometry. The X3 provides one such solution. The nesting technique used in the X3 features channel widths that increase slightly with increasing channel diameter, of the general form:

$$w_i = w_1 + \alpha\,(i-1)\,, \tag{3.19}$$

where w_i is the width of the i^{th} channel, w_1 is width of the innermost (i=1) channel, and α is a scale factor. For this example nesting technique, α is set to unity.

The remaining free variables in Equation 3.18 are the thruster density ρ and the thruster thickness Δz. Hall thruster scaling work has shown that the discharge channel length (and thus thruster thickness) are primarily dictated by the discharge voltage [11], which is held constant at 300 V here. Thus, Δz is simply combined with thruster density ρ, and a value for $\rho\Delta z$ is chosen such that the Hofer and Randolph single-channel A_{th} value of 2.425 is recovered

Table 3.2: Values for $A_{th,NHT}$ from Equation 3.18 for the example NHT channel nesting technique.

N_{ch}	1	2	3	4	5
$A_{th,NHT}$	2.425	1.941	1.851	1.841	1.854

for single-channel thrusters. This assumes that the value of $\rho\Delta z$ remains constant no matter how the channels are scaled.

The values for $A_{th,NHT}$ produced by this nesting technique in Equation 3.18 are presented in Table 3.2. It is worth noting that this particular nesting technique demonstrates asymptotic behavior for $N_{ch} \geq 3$. This indicates that for this method the benefits from nesting more than three channels may be limited. The implications of this at the system level are explored below.

It can be argued that the NHT configuration only becomes useful at large enough thruster sizes, i.e., high enough thruster powers, that significant diameter savings can be achieved by nesting discharge channels. Single-channel Hall thrusters have been successfully demonstrated up to power levels of 50 kW [10], so this power level is chosen as an arbitrary lower limit for NHT application for the purposes of this model.

3.4.3 NHT Mass Model Results

We now use the values of $A_{th,NHT}$ from Table 3.2 in Equation 3.11, and apply that equation to four high-system-power mission profiles: 500 kW, 750 kW, 1.0 MW, and 1.25 MW. These results are intended to demonstrate the basic trends of NHTs as compared to single-channel thrusters and the potential advantages as system power level increases. The plots presented below are designed to provide direct comparison to the results found with the original Hofer-Randolph model. For a given total system power level, the model calculates the expected system mass for number of active thrusters (N_{ac}), which we plot along the abscissa in the following figures. Thus, for a 1-MW system power, "1 active thruster" corresponds to a system of one active and one redundant thruster at input powers of 1 MW each, "2 active

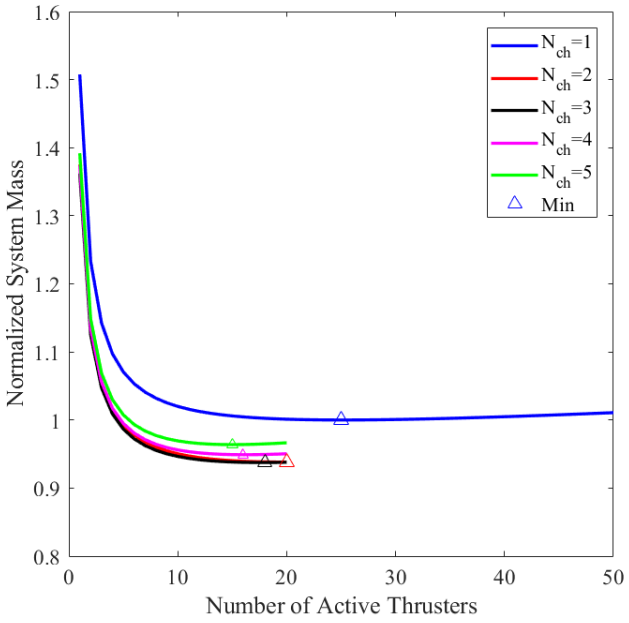

Nch	Minimum Value	# of Thrusters at minimum
1	1.00	25
2	0.938	20
3	0.937	18
4	0.949	16
5	0.964	15

Figure 3.3: System mass profile for X3-like NHT scaling for values of N_{ch}=1–5 and 1 redundant thruster at 1-MW system power. The minimum of each curve is denoted with a triangle. The minimum values and their corresponding number of thrusters are listed in the table to the right of the figure.

thrusters" corresponds to a system of two active and one redundant thrusters at input powers of 500 kW each, and so on. To facilitate direct comparisons, the results are normalized against the minimum value of the N_{ch} = 1 case. Thus, values from the ordinate of the plots correspond to the relative mass as compared to the single-channel minimum value.

3.4.3.1 Variation with System Power Level

Figure 3.3 presents the results for X3-like scaling for a 1-MW system with one redundant thruster. The profiles are plotted as the normalized system mass versus the number of active thrusters. All curves display a broad plateau near the minimum, similar to the original Hofer-Randolph model results. The original work notes that there were wide ranges of thruster powers where the system mass is within 5 to 10% of the minimum value depending on system power, giving system designers a large design space in which to work. This feature appears

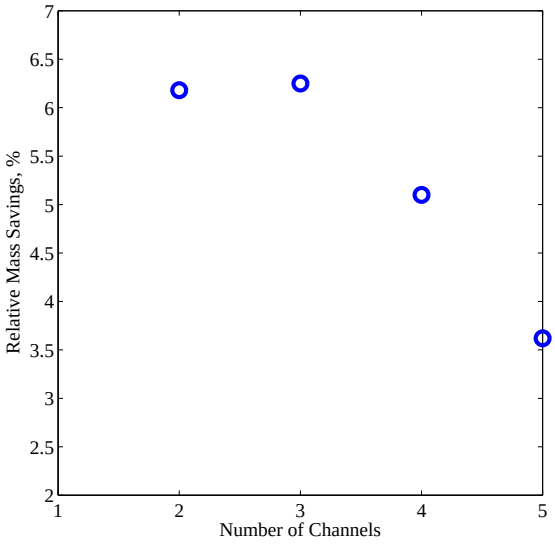

Figure 3.4: The maximum mass savings for values of N_{ch} between 2 and 5 for a 1-MW mission with X3-like channel nesting.

in the NHT results as well.

Figure 3.4 plots the savings at mass minimum versus number of nested channels. We find that the mass savings are nearly identical for N_{ch} = 2 and 3. This result seems odd because Table 3.2 shows three-channel NHTs to have a lower specific mass than two-channel ones with this type of nesting. However, these results can be explained by Figure 3.5. Here, the thruster mass and cabling/XFS mass are plotted separately along with their sum. It can be seen that the mass savings provided by the lighter thruster in the three-channel case are almost exactly canceled out by the increase in cabling and XFS required by the additional channel. This explains the heavier four- and five-channel systems as well: as shown in Table 3.2, thruster specific mass is nearly constant for N_{ch}=3–5, so the additional cabling and XFS cause these thruster systems to be heavier than the three-channel systems.

Next, we look at results for system powers from 500 kW to 1.25 MW. Figure 3.6 presents the mass savings provided by nesting 2–5 channels with the X3-like technique. The trends seen for the 1-MW case are present at all power levels studied. For this nesting technique

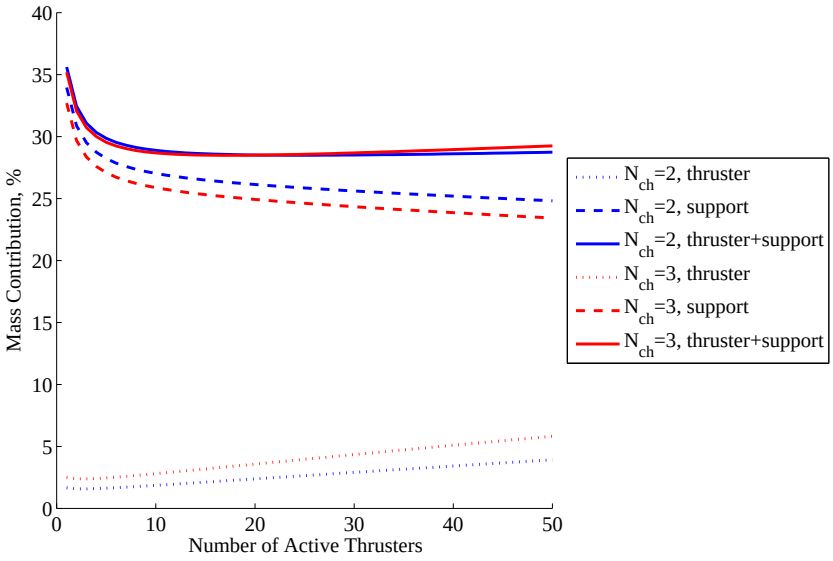

Figure 3.5: The NHT-modified terms in the mass model plotted separately and summed.

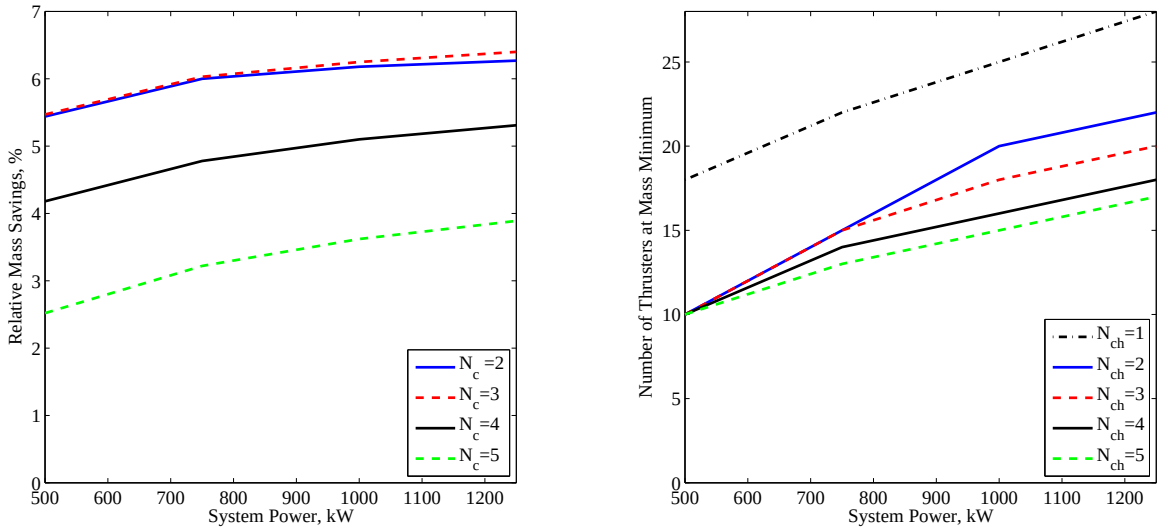

(a) Maximum mass savings

(b) Number of thrusters at mass minimum

Figure 3.6: The savings provided by NHTs scaled like the X3 for total mission powers of 500 kW–1.25 MW: (a) maximum mass savings and (b) number of thrusters at mass minimum. These savings manifest both as decreased mass and as fewer number of thrusters.

and a 500 kW system, 2- and 3-channel thrusters provide approximately 5.5% savings, and when system power is increased to 1.25 MW, mass savings increase to about 6.3%. One outcome of these results that is relevant to the cost model is the fact that the NHT mass minimum typically occurs at smaller number of active thrusters than for SHTs. This result is demonstrated in the plot in Figure 3.6b. Here, it can be seen that NHTs optimize to anywhere between 5 and 11 fewer thrusters than SHTs. The implications of this result on system cost will be discussed below.

3.4.3.2 Variation with Number of Redundant Thrusters

The results of the original mass model showed that the mass-optimum number of thrusters for a given system power grew as the number of redundant thrusters increased. This is to be expected: for the extreme case of a single active thruster, a triple-fault-tolerance requirement would require flying one active and three redundant thrusters. This would double the system mass over the single-fault-tolerance (one active and one redundant thruster) case.

Our results showed similar trends, as presented in Figure 3.7 for a 1-MW system and up to five redundant thrusters. These results demonstrate a decrease of mass savings with more redundant thrusters. This is due to the fact that increased redundancy pushes the mass minimum to lower power thrusters (and thus larger numbers of active thrusters, away from the realm of NHTs). As N_{rd} increases from 1 to 3, the mass savings provided by the NHT geometry decreases from 6% to 4%. However, the minimum for single-channel thrusters shifts to higher numbers of active thrusters, occurring at 44 active thrusters for $N_{rd} = 3$. This shift to more thrusters has an impact on the cost side of the model, as will be discussed below.

3.4.3.3 Variation with Number of Active Channels

Complexity is a factor not captured explicitly by this model, yet complexity adds system risk and thus cost potential due to failures or timeline delays. Though it is difficult to quantify

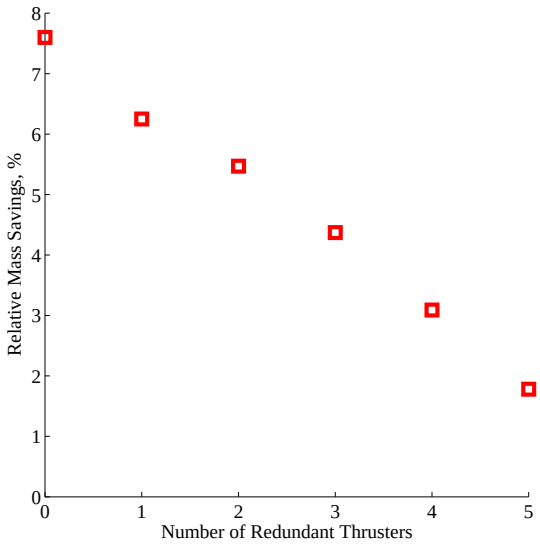

Figure 3.7: The relative mass savings versus number of redundant thrusters for a 1-MW system with 3 nested channels.

complexity or risk for the model, one way to explore this is to present the results not as functions of the number of active thrusters but as functions of the number of active channels. This will not fully address the issue, as the question of whether, for instance, a system of five 1-channel thrusters is more complex (more thrusters to manufacture and mount; likely a higher total part count; five cathodes) or less complex (lower power levels lend to easier ground testing; larger thrusters are likely more difficult to manufacture and assemble) than a 5-channel NHT is best answered only with specific thruster designs. However, it presents the results in a unique way and further illustrates the benefits of the NHT configuration.

Results for a 1-MW mission are re-plotted in Figure 3.8 against the number of active channels instead of the number of active thrusters. Two cases are plotted. First, to isolate the NHT benefit, the relative system masses are plotted for zero redundant thrusters in Figure 3.8-a. If number of channels is viewed as an approximation of system complexity, it can be seen that for a given system complexity, the NHT configuration offers notable mass savings, practically regardless of the number of channels being nested. Plotting in this

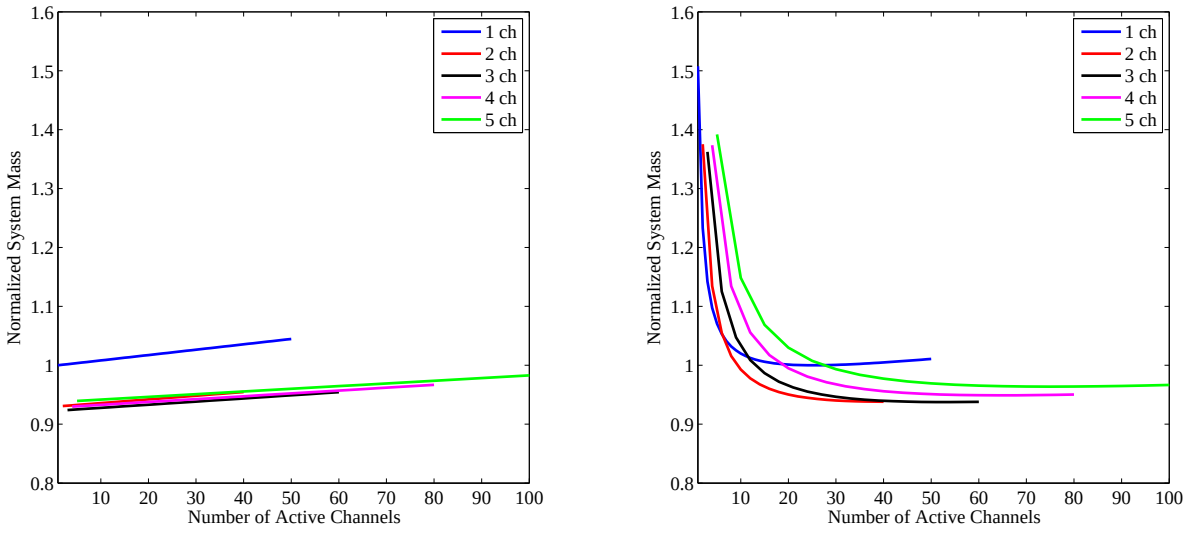

(a) Zero redundant thrusters

(b) One redundant thruster

Figure 3.8: Mass model results for X3-like NHT scaling applied to a 1.0-MW system power-level mission for: (a) 0 and (b) 1 redundant thrusters, plotted as a function of the total number of active channels instead of the total number of active thrusters. These results illustrate the packing efficiency of the NHT geometry.

manner removes the effects of the cabling and xenon flow system (because these items contribute equally to equal numbers of channels) and isolates the thruster-based mass savings. Ultimately the mass savings for a given number of active channels is driven by the value of A_{th}, and for $N_{ch} = 2 - 5$ these values are very close together.

Next, the results are presented for one redundant thruster to provide a practical result, as shown in Figure 3.8-b. Because each redundant thruster is equivalent to N_{ch} redundant channels, NHTs (especially those with many channels) see a sharp increase in system mass at very small numbers of channels. At large enough values of active channels, the redundant channels no longer contribute significantly to total system mass, and the increased packing efficiency offered by the NHT configuration becomes clear as in Figure 3.8-a.

3.4.4 Cost Model

3.4.4.1 Modifications to Capture NHTs

The cost portion of the original model splits the system cost into two portions: recurring expenses (which occur for every thruster manufactured in the system) and non-recurring expenses (which occur once for the entire system, regardless of the number of thrusters). Recurring expenses include the cost of raw materials and manufacturing of all of the various string components (thruster, gimbal, PPU, cabling, XFS). Non-recurring expenses include the string development costs (such as the flight qualification of the thruster). In general, the cost model is not focused on specific dollar amounts but is intended to provide rough-order-of-magnitude comparisons between system configurations.

A major driver of the recurring expenses come from the PPU. For the same reasons that the mass is not expected to change significantly, we do not expect the cost to vary significantly between SHTs and NHTs. The manufacturing costs of a single NHT is likely higher than a single-channel thruster of equal power. However, as demonstrated in the mass portion above, a NHT will be lighter. The thruster mass is a good proxy for the cost of raw materials, and as such, it can then be expected that the NHT geometry will offer a savings

in raw material cost over single-channel thrusters. To the level of accuracy and specificity of the model, these details are considered likely to balance each other. Finally, though the costs involved with flight-qualifying a high-power thruster are likely quite high due to the increased xenon throughput and thermal loading to the vacuum facility (causing increased usage of liquid nitrogen in cryogenic pumps, for instance), there is no reason to expect these costs to differ between a SHT and a NHT of the same power. The Hofer-Randolph cost model already factors in these expenses when scaling to high power, and no change is thus necessary for NHTs.

Due to the fact that the PPU and thruster costs are likely to not vary much between a single-channel and multi-channel Hall thruster at a fixed thruster power level and that flight-qualification costs will be the same per kW between SHTs and NHTs, no attempt is made to modify the cost model here to account for NHTs. However, because the mass and cost are coupled via launch costs, the cost model will produce different curves for NHTs than for SHTs. The differences are explored below.

3.4.4.2 Cost Model Results

Example results from the cost model are presented in Figure 3.9 for a 1-MW mission with a single redundant thruster. For clarity, we only present results for N_{ch}=1 and 3. The location of the mass minimum of each curve is denoted with a triangle. Just as was shown with the original Hofer-Randolph model, these cost curves feature sharp increases on either side of their minima, which are located at values less than 10 active thrusters. The savings in mass provided by NHTs only contributes to a portion of the cost savings, which is represented by the offset between the two lines. This comes from the fact that system mass is directly tied to cost through launch costs; lighter systems are less expensive to launch into orbit, creating the offset. The other contribution comes from the mass minimum shifting to fewer active thrusters for NHTs. For the 1-MW example mission, the mass minimum shifts from 25 active 1-channel thrusters to 18 active 3-channel thrusters, and this contributes to a further

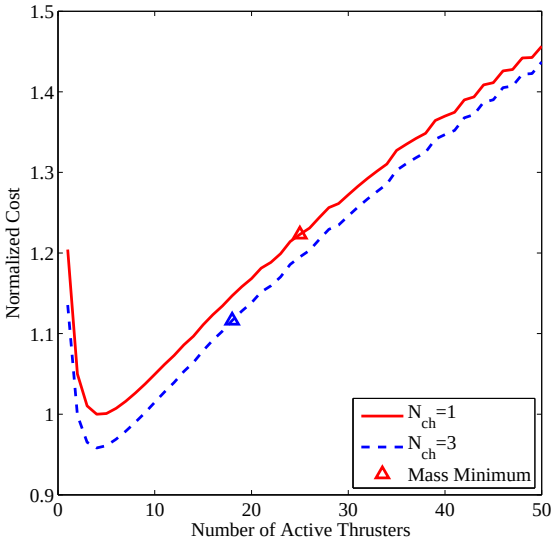

Figure 3.9: Cost model results for a 1-MW mission with one redundant thruster. The location of the mass minimum is indicated on each curve with a triangle.

decrease in system cost. In total, for the 1-MW mission, the mass-optimized 3-channel NHT system offers cost savings of approximately 9% over the mass-optimized 1-channel thruster system.

Next, we calculate the relative cost savings of a 3-channel NHT system as compared to a 1-channel system for mission powers of 500 kW to 1.25 MW. The cost savings offered by the NHT configuration are actually higher for smaller system power levels, even though the relative mass savings are lower. This is due to the fact that at lower powers, NHTs shift the mass minimum farther than for higher power systems, contributing to larger cost savings. For X3-like channel scaling, potential system cost savings predicted by the model range from 9–13%. Savings of this magnitude have the potential of equaling several $10s of millions or more.

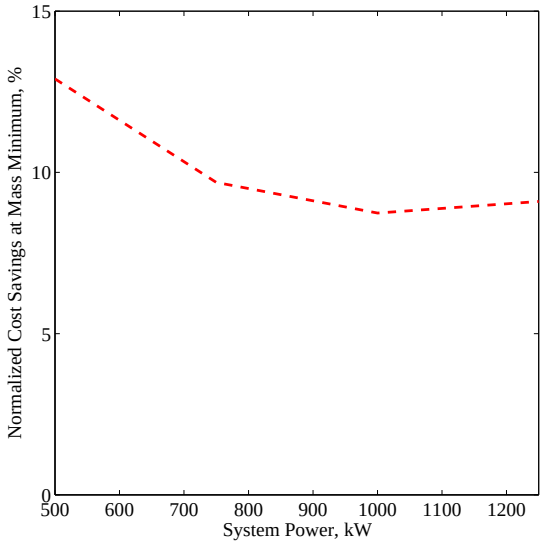

Figure 3.10: Relative cost savings of a 3-channel NHT system as compared to a 1-channel system versus total system power.

3.4.5 Discussion of Results

Our results indicate that the impact on overall propulsion system mass provided by the reduced thruster specific mass of an NHT geometry similar to the X3 NHT is small but not negligible. This can be attributed to the fact that the thruster mass itself makes up less than 40% of the total system mass in the Hofer-Randolph model. Placing the potential impact of thruster mass savings in the context of the entire propulsion system is important from both a mission planning perspective and a thruster development perspective. A 3-channel NHT nested like the X3 produces a thruster that is approximately 24% lighter than a SHT of the same power level. Yet when compared to an optimized system of SHTs (for a 1-MW system, 28 SHTs), the same 3-channel NHT only offers total propulsion system mass savings on the order of 6%. The modest mass savings correspond to slightly higher savings in cost, driven by the combination of the improved mass and the shift in mass-optimized number of thrusters to a smaller value. X3-like nesting provides cost savings ranging from 9-13% for 3-channel thrusters at mass minimum.

However, this X3-like nesting technique is not the only method to nest channels. Other potential scaling options include constant-width channels and channel widths that scale at constant w/d. The limiting case, where there are infinitely-thin magnetic pole pieces, is not a physical thruster design solution but provides an upper bound of the available mass savings. To explore these options, we calculate $A_{th,NHT}$ for the following nesting techniques:

1. **Constant w:** For this example geometry, the innermost channel is designed using a fixed w/d ratio, and the remaining channels are also w_1 wide:

$$w_i = w_1 \tag{3.20}$$

   This results in channels that are equally spaced with equal-sized magnetic pole pieces between them.

2. **Constant w/d:** Here, each successive channel is wider at a constant ratio of w/d. This scaling agrees with some literature on single-channel thrusters [170]. Here, the same w/d value from above is used for all 5 thruster channels, which results in channels that get successively farther apart:

$$w_i = \left(\frac{w}{d}\right) d_i \tag{3.21}$$

3. **Infinitely Thin Pole Pieces:** This is the limiting case where the entire thruster face consists of exit area and there are, in essence, no distinct channels. It should be noted that this is not realizable in practice, but provides an upper bound for the NHT configuration.

The values of $A_{th,NHT}$ produced by these scaling methods are presented in Table 3.3.

In Figure 3.11 we compare the mass savings offered by these alternative nesting techniques. These results match what would be expected from the values given in Table 3.2. The Infinitely Thin Pole scaling provides the greatest mass savings because it provides the lowest values of $A_{th,NHT}$; the Constant w scaling provides minimal savings due to its modest

Table 3.3: Values of $A_{th,NHT}$ as calculated by Equation 3.18 for three alternative channel nesting techniques. All are for 300 V operation at nominal Hall thruster current densities.

N_{ch}	Constant w	Constant w/d	Infinitely Thin Poles
1	2.425	2.425	2.425
2	2.160	1.828	1.676
3	2.110	1.692	1.440
4	2.100	1.652	1.329
5	2.101	1.639	1.266

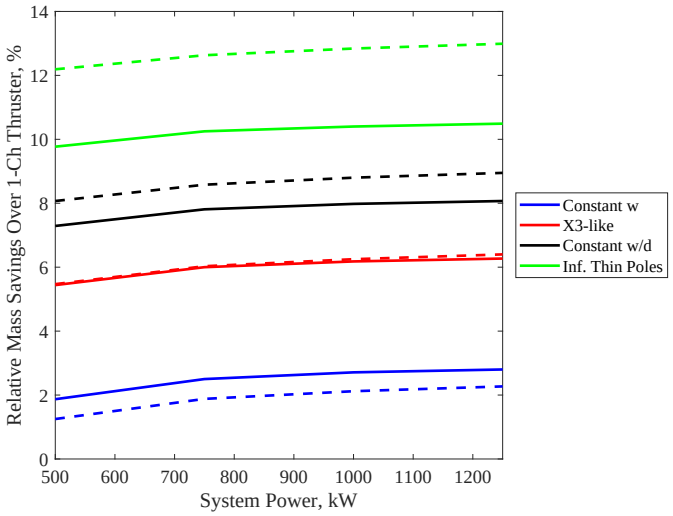

Figure 3.11: The relative mass savings provide by both 2-channel (solid line) and 3-channel (dashed line) configurations for all four example NHT nesting techniques.

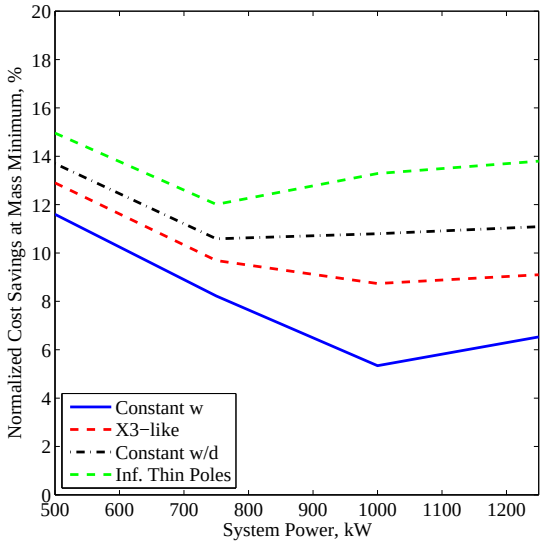

Figure 3.12: Cost savings at mass minimum versus system power level for all four example NHT nesting techniques.

decreases in $A_{th,NHT}$. The curves also help to illustrate the effects of the asymptotic nature of the $A_{th,NHT}$ values: for the Constant w scaling, where $A_{th,NHT}$ values for $N_{ch} = 3 - 5$ are all approximately equal, the three-channel geometry actually provides smaller system masses than the two-channel; for the Constant w/d method, where the values of $A_{th,NHT}$ have not plateaued fully even at $N_{ch} = 5$, the three-channel geometry still offers mass savings in excess of 1.5% over the two-channel.

These mass savings correspond to cost savings in the same manner as for X3-like nesting. Cost model results are presented in Figure 3.12. All of these nesting techniques display similar behavior: maxima at 500 kW system power and a slight drop in cost savings as system power increases. These results indicate that cost savings in excess of 12% may be possible with alternative channel nesting techniques, which could translate to savings of many tens of millions of dollars.

Though system mass savings of 6% are much less than what the early NHT work implied may be possible, these savings still have the potential for important mission impact. A 6%

savings in mass for a 1-MW mission corresponds to approximately 800 kg. This mass can either be left off, thus realizing the full drop in system cost predicted by the model, or re-purposed for more supplies or scientific instruments at the same launch cost. And because the cost savings associated with NHTs come from both launch costs and the shift to fewer thrusters, even if the mass savings is re-allocated to supplies or instruments, NHTs still provide a cost savings to the mission. This provides designers with increased flexibility as the full mission is being developed.

3.4.6 Other NHT Benefits

There are a number of features of NHTs that are not explicitly captured in this model, features whose addition would allow model users further insight when designing the propulsion string [82, 90]. One is the partial redundancy offered in NHTs: certain failure states may remove a single channel from operation, not the entire thruster, and a NHT with a wide throttling range such as the X3 would likely be able to throttle to a different condition to make up for the missing channel. Because the existing model only looks at redundancy from a string perspective, adding this feature with limited data on the types and relative frequencies of failure states seen on orbit would be difficult to do. However, as EP systems continue to proceed through flight qualification to flight, more of this information will become available. There is also footprint/area savings offered by NHTs over systems of many single-channel thrusters, a feature that would have an impact on spacecraft design. Additionally, further refinements to the cost model will be possible as more development work is done for NHTs and for high-power Hall thruster systems in general. These additional data may reveal scaling that differs from that proposed by the original work that may change the conclusions regarding propulsion string design.

3.5 Summary

The purpose of this chapter was to demonstrate the pressing need for a 100-kW class, 60% efficient, 2000-3500-second specific impulse electric propulsion device. The wide variety of applications for this device detailed in the literature make this point clear. Such a device can be used for cargo or crew transport to a wide variety of scientifically interesting and economically beneficial bodies, including near-Earth asteroids and Mars. By updating a mass and cost model to incorporate NHTs, we demonstrated that there are benefits on the order of 10% for both metrics by using NHTs like the X3 in place of single-channel Hall thrusters for high-power missions. There are also other factors, including the increased throttling range of a NHT over a SHT, that are not captured in the model but that may provide increased mission benefits. This is not to say that the only answer for affordable mission architectures for Mars exploration or near-Earth asteroids is the X3, or even a Hall thruster. However, based on the current state of electric propulsion technology, this path certainly provides benefits with respect to cost, mass, and overall effectiveness.

CHAPTER 4

Experimental Apparatus and Analysis

"In the name of the cathode, the anode, and the holy grid?"

– *Gravity's Rainbow* by Thomas Pynchon [171]

"Nevertheless, the logistical requirements for a large, elaborate expedition to Mars are no greater than those for a minor military operation extending over a limited theater of war."

– Werner Von Braun [134]

4.1 Introduction

The characterization of the X3 took place in multiple state of the art vacuum facilities and utilized a number of thruster and plasma diagnostics. This chapter provides information about the various pieces of test equipment used during this work. It also details the data analysis techniques used. In Section 4.2, we describe the configuration of the X3 as it was tested here, including the electrical and propellant details and the hollow cathode used. Section 4.3 describes the test setup at the University of Michigan's Plasmadynamics and Electric Propulsion Laboratory, where the low-power characterization of the X3 took place. Following that, Section 4.4 details the equipment used during high-power testing at NASA Glenn Research Center, including a description of the various plasma diagnostics and data analysis techniques. Example data from each diagnostic and a description of error analysis

are also provided. Section 4.5 describes the X3-dedicated thrust stand that was originally developed at UM, versions of which are now at both UM and GRC. Finally, Sections 4.6 and 4.7 detail the test setup for magnetic field and propellant uniformity mapping of the X3, respectively.

4.2 The X3 Nested Hall Thruster

The X3 is a 100-kW class, three-channel NHT. The design and history of the thruster is provided above in Chapter 2. Below, we provide details of the operational configuration of the X3 during the testing described here. This includes a description of the power and propellant configuration as well as a description of the cathodes used during testing.

4.2.1 Laboratory Operation

The three channels of the X3 are operated off of three separate laboratory power supplies. However, the X3 has always been operated with all channels firing at the same anode potential. A small amount of mixed-voltage operation was done with the X2 [1], but those tests provided no results to suggest that this was beneficial, so it has not been pursued with the X3. The X3 is capable of mixed-voltage operation should further investigation be desired. The anode potential was always set manually (i.e., no circuitry or autonomous control was used to ensure matched potentials), so channels were typically within 2 V of each other.

Because the thruster runs off of a single centrally-mounted cathode, the three discharge channels all share a cathode return line (the negative side of the discharge circuit). Each of the six electromagnets, the cathode keeper, and the cathode heater all operate off of separate laboratory power supplies as well. A notional electrical schematic of thruster operation is shown in Figure 4.1.

A simplified schematic of the propellant feed setup is presented in Figure 4.2. Xenon propellant is fed to each channel and the cathode through individual 0.25-inch stainless steel

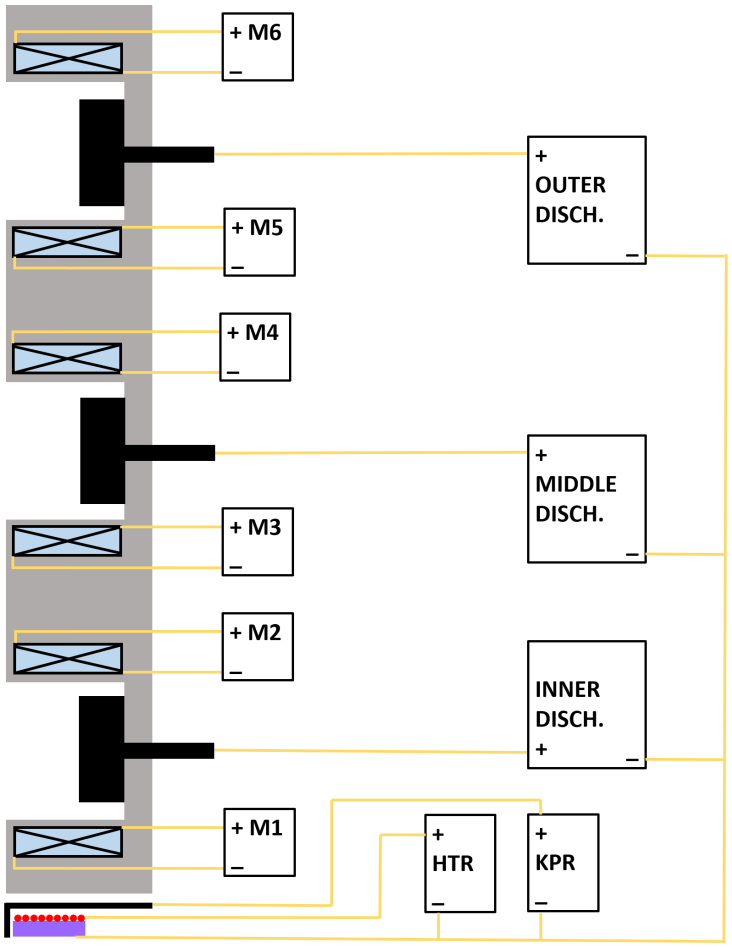

Figure 4.1: A notional electrical schematic for laboratory operation of the X3.

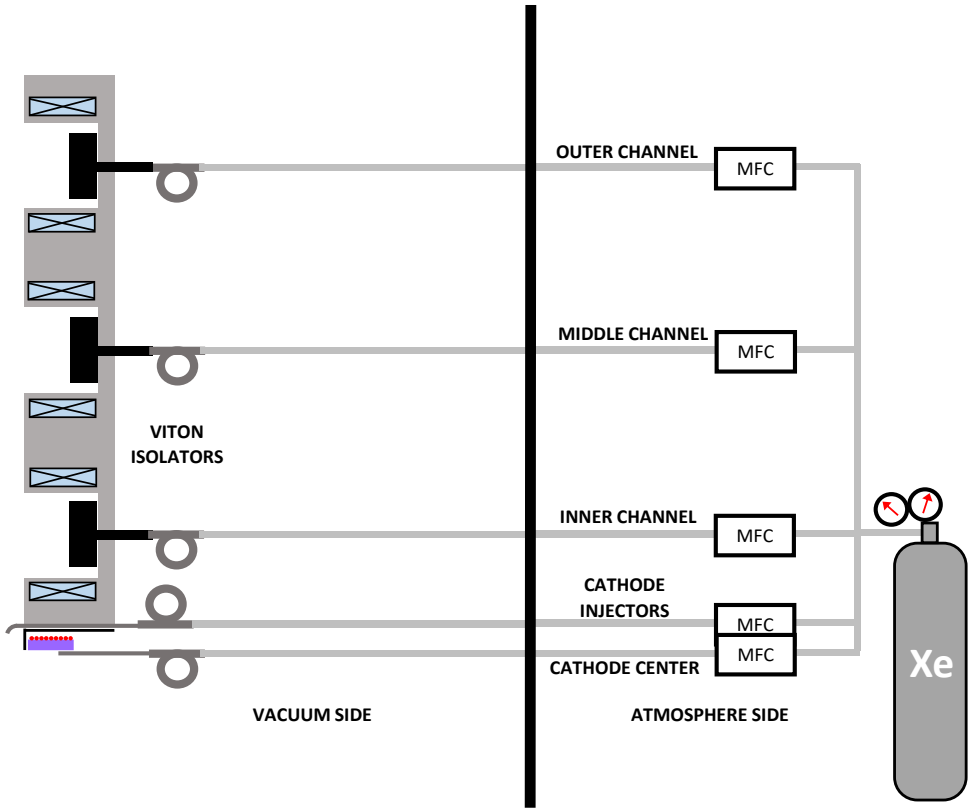

Figure 4.2: A simplified schematic of the propellant feed setup for the X3. Each of the five lines is plumbed from a single bottle of xenon kept on the atmosphere side of the test setup and is controlled by a separate mass flow controller (MFC). The stainless steel lines are isolated from the thruster discharge by lengths of Viton in each line located at the thruster.

lines from the single xenon bottle (kept on the atmosphere side of the setup), through a flow control manifold, and into the chamber. The lines are routed through the thrust stand (omitted from the schematic for clarity) and electrically isolated from the thruster discharge via lengths of Viton tubing located immediately upstream of the thruster. This setup was identical between the UM and GRC setups described below except for the exact mass flow controllers used.

4.2.2 Hollow Cathodes

Both GRC and JPL built high-current hollow cathodes for the X3. The GRC cathode is based on the barium oxide (BaO) emitter technology being developed there [172, 173],

whereas the JPL cathode is derived from a lineage of lanthanum-hexaboride (LaB_6) emitter design from that center [97]. The JPL cathode features neutral gas injectors at its tip to help reduce energetic ion bombardment of the cathode (thus extending cathode lifetime) [174].

We used the JPL cathode (Figure 4.3) throughout this work. The cathode went through a number of small design iterations but remained fundamentally the same in geometry and operation. We set the cathode flow for X3 operation as a fixed percentage of total anode flow rate, following typical Hall thruster procedures [100]. However, unlike traditional hollow cathodes, the JPL cathode features the aforementioned external injectors. Thus, the total cathode flow fraction (TCFF) for this cathode is split between the cathode center and the injectors:

$$\text{TCFF} = \frac{\dot{m}_{c,t}}{\dot{m}_{a,t}} = \frac{\dot{m}_c + \dot{m}_{inj}}{\dot{m}_{a,I} + \dot{m}_{a,M} + \dot{m}_{a,O}}. \tag{4.1}$$

For thrusters without downstream injectors on the cathode (i.e., most Hall thrusters), TCFF $\rightarrow$ CFF (the traditional "cathode flow fraction"), so for simplicity we use TCFF as the nomenclature throughout this work when referring to the cathode flow fraction of any thruster.

Work at JPL has identified 16 sccm as the optimized flow rate for the cathode center ($\dot{m}_c$), so any remaining flow of the TCFF is sent through the external injectors [97, 174]. The TCFF varied from 10% in initial work (following the precedent of the X2 and early testing with the X3) to 7% in later low-power testing and the high-power characterization at GRC. Additionally, as described in Appendix A, at one operating condition we swept the TCFF from 7% to 3% in 1% steps and characterized the thruster performance and plume properties at each step. The thruster demonstrated stable operation and performance that was relatively unchanged from the 7% case. However, as detailed in the Appendix, further work is needed to fully study this low-TCFF operation.

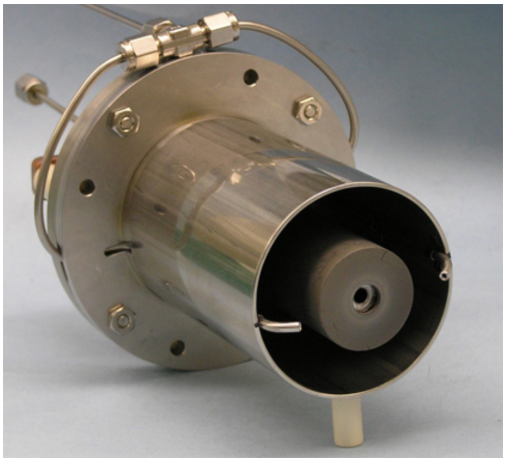

Figure 4.3: The JPL-designed high-current LaB_6 hollow cathode used during this work. The tubes at the 3 and 9 o'clock positions near the exit orifice of the cathode are the external gas injectors.

4.3 Testing at PEPL

This section describes the infrastructure and setup for testing that took place in the University of Michigan's Large Vacuum Test Facility at the Plasmadynamics and Electric Propulsion Laboratory.

4.3.1 The Large Vacuum Test Facility

The Large Vacuum Test Facility (LVTF) is a 9 meter long, 6 meter diameter stainless-steel-clad cylindrical vacuum chamber. The chamber was evacuated to rough vacuum (100-mTorr range) using four 400-cfm mechanical pumps backed by two 2000-cfm mechanical blowers. Seven CVI TM-1200 re-entrant cryopumps were then used to bring the chamber down to a base pressure on the order of 10^{-7} Torr. The cryopumps featured liquid-nitrogen-cooled shrouds and liquid-helium-cooled cold heads that operated between 13–20 K. In this configuration (without additional pump thermal shields as described below), the chamber had a pumping speed of approximately 240,000 l/s on xenon.

This testing occurred over a number of years during which the exact configuration of the LVTF testing apparatus changed slightly. These different configurations will be denoted as

Configuration 1 and Configuration 2 and are detailed below. Differences include changes to the waterfall setup, the pressure measurement setup, and the cryopump protection scheme.

4.3.1.1 LVTF Configuration 1

LVTF Configuration 1 was mostly similar to the setup used during Florenz's work with a few key differences [4]. In an attempt to protect the cryopumps from radiant heat from the thruster and plasma plume, we hung mylar blankets over the front face of the cryopumps. These blankets limited cryopump crashing (a phenomenon that plagued the first firing of the X3 in the LVTF, in which the cryopumps will begin very rapidly warming, then shed their gas loads, causing a surge in chamber pressure) at the expense of chamber pumping speed: the blankets caused the chamber pressure at similar conditions to increase over those achieved during the first firing of the thruster by a factor of approximately 2.5.

Pressure in the LVTF was measured in Configuration 1 using Bayard-Alpert hot-cathode ionization gauges mounted in various locations both on the wall of the LVTF and internal to the chamber near the thruster. The internal gauge was a Varian 563 model, mounted with the entrance pointed downstream and in the thruster exit plane. The gauge was located approximately 20 cm outside the outer diameter of the thruster. All gauges used in Configuration 1 were calibrated on nitrogen and were thus corrected for xenon using the traditional technique described by Dushman [175].

Thruster telemetry was output from the breakout box as described by Florenz [4]. That telemetry, and the suite of thermocouples mounted to the X3, were read by a pair of Agilent 34970A data loggers and recorded by a computer running Agilent Benchlink software. We measured thruster discharge current in Configuration 1 using magnetoresistive sensors sized for the expected current of each channel. These were placed in-line with the discharge circuit downstream of the filter capacitors.

4.3.1.2 LVTF Configuration 2

We made a series of small setup upgrades after testing in Configuration 1 was completed. A new dedicated aluminum-extrusion table was designed and fabricated for the LVTF. This table consisted of modular pieces to allow for modification of the height of the thruster mounting surface, depending on the size of the thruster being fired. With the table in its shortest configuration and the second-generation HMT stand installed, the centerline of the X3 coincided with the centerline of the LVTF. This is an improvement over Configuration 1, where the centerline of the X3 was approximately 1 m below the centerline of the LVTF.

We redesigned the mylar shielding of the cryopumps for increased robustness in Configuration 2. Instead of being loosely hung in front of the pumps as in Configuration 1, a setup that would often be disrupted during chamber venting, aluminum frames were designed and manufactured. The mylar blankets were mounted to these frames, which then were affixed to the liquid nitrogen shrouds of the cryopumps.

Initially, only six shields were deployed, with the pump at the top of the LVTF not receiving a shield due to the difficult nature of reaching it and mounting a shield. However, during the initial bakeout of the X3, the unshielded pump ran approximately 4 K hotter than the other six pumps (which were within 1 K of each other) before crashing approximately 3 hours into operation of the outer channel at 16 kW discharge power. We subsequently vented the chamber and installed a shield on this pump. Afterward, it behaved comparably to the others, indicating the effectiveness of the shields.

However, as with the shield implementation in Configuration 1, these gains in pump operating temperature were at the expense of pumping speed, due to the reduction in open area fraction on the pumps. The effect on pumping speed was calculated during the testing that occurred in Configuration 2. We achieved an average effective pumping speed of approximately 165,000 L/s on xenon with these shields in place, a decrease of approximately 35% over the un-shielded pumping speed of the LVTF. A new pressure gauge was installed in the LVTF in Configuration 2; matching with what other electric propulsion facilities in the

country have implemented, an MKS Stabil ion gauge was purchased and installed [111,176]. The gauge was placed in approximately the same location as the Varian 563 in Configuration 1. The new gauge was calibrated on nitrogen, so its measurements were corrected for xenon in the same manner as described above for Configuration 1.

A new data acquisition system was implemented in Configuration 2. The system was a National Instruments PXIe-1078 chassis with one PXIe-8840 embedded controller; one TB-4300B high-voltage, 30:1 attenuation card; two PXIe-4353 thermocouple cards; one PXI-6225 multifunction DAQ card; one PXI-6722 8-channel analog output card; and one PXI-8430/16 RS-232 communication card. This system was only capable of reading voltages up to ± 10 V. As such, modifications to the breakout box were necessary to reduce all telemetry voltages to acceptable levels for the PXI system.

Voltage dividers were added to the discharge, heater, keeper, and electromagnet voltage sense lines. The currents of the electromagnets, heater, and keeper were measured using Deltec precision shunt resistors, which caused these sense lines to have a high common-mode voltage. To properly condition these for the PXI system, Analog Devices AD202KN isolation amplifiers were added to these sense lines to both filter out the common-mode voltage and amplify the voltage drop across the shunt resistor. The isolation amplifiers were mounted on custom-designed printed circuit boards.

Additionally, the method of measuring discharge current was modified in Configuration 2. The magnetoresistive sensors were removed from the thruster telemetry setup due to their susceptibility to damage during thruster transient events. In their place, we used precision shunt resistors (like those used for the electromagnets, heater, and keeper but sized for the expected current of each line) on the high side of each discharge power line, upstream of the filter capacitors. The signal from each of these resistors passed through an isolation amplifier designed to amplify the voltage drop across the resistor to a maximum of 2 V with respect to ground. All of these modifications were housed inside the existing breakout box structure.

4.3.2 Other Test Apparatus

A number of pieces of equipment at PEPL were unchanged between LVTF Configurations, including the propellant manifold, the power apparatus, and the discharge current oscillation measurement technique. We detail these here.

We provided electric-propulsion grade xenon to the X3 using five electro-polished stainless steel lines. The flow for each of the five thruster channels (Inner, Middle, and Outer anodes, cathode center, and cathode external injectors) was prescribed by a separate Alicat-brand MFCs. In both LVTF configurations, a 400-sccm controller was used for the Inner channel, an 800-sccm controller for the Middle channel, a 3000-sccm controller for the Outer channel, a 50-sccm controller for the cathode center flow, and a 1000-sccm controller for the cathode external injectors. These flow controllers were used in a manifold like what is described in Figure 4.2 above and were calibrated in situ using a Bios Definer 220L Dry-Cal system, which had an accuracy of 1%. Before and after each experiment, each MFC was calibrated at a number of flow rates and a linear fit was created to calculate the xenon flow rate at any given set point.

We ran the X3 from three separate laboratory discharge power supplies: a 60-kW supply for the Inner channel, a 100-kW supply for the Middle channel, and a 150-kW supply for the Outer channel. Each supply was connected to two 160-μF capacitors in series to filter the plasma oscillations from the discharge supplies. The common for all three discharge supplies was shared with the single centrally-mounted cathode. Power to the six electromagnets and the cathode heater and keeper was provided using commercially-available rack-mounted power supplies. We non-invasively monitored discharge current oscillations using three commercially-available current sensors each attached to a commercially-available amplifier. AC-coupled discharge current was measured with precision down to 5 mA.

4.4 Testing at NASA GRC

We performed the high-power characterization of the X3 in Vacuum Facility 5 (VF5) at NASA Glenn Research Center (GRC) in Cleveland, OH. This section describes the testing apparatus used there.

4.4.1 Vacuum Facility 5

Vacuum Facility 5 is a 4.6 m diameter, 18.3 m long cylindrical vacuum chamber that features 33 m^2 of cryogenic pump surfaces, providing a pumping speed of 700,000 L/s on xenon. The facility walls and cryogenic panels are lined with graphite plates to minimize backsputter during thruster operation. Pressure inside the facility was monitored using four MKS Stabil Ion Gauges mounted in various locations. The pressures reported here are from the gauge mounted in the exit plane of the thruster approximately 1.5 meters from thruster centerline, pointed downstream. This gauge was calibrated on xenon and was corrected for orientation using techniques by Yim [177]. Facility base pressures were typically on the order of $1x10^{-7}$ Torr during the GRC test campaign. Orientation-corrected background pressures while firing the thruster ranged from $4.3{\times}10^{-6}$ to $4.2{\times}10^{-5}$ Torr.

A new thruster support cart was designed specifically for the NextSTEP program. The cart rolls into VF5 through the endcap along the rail system inside the chamber and sits in front of the existing thruster testing location typically used for smaller thrusters such as the HERMeS thruster [67, 178, 179]. A schematic of the location of the X3 inside VF5, as well as the location of the plasma diagnostics package, is shown in Figure 4.4, and a photograph of the X3 NHT installed in VF-5 is shown in Figure 4.5.

4.4.2 Propellant

Electric propulsion-grade xenon propellant was provided to the thruster via five electropolished stainless steel feed lines. Each line featured a precision flow controller to prescribe the

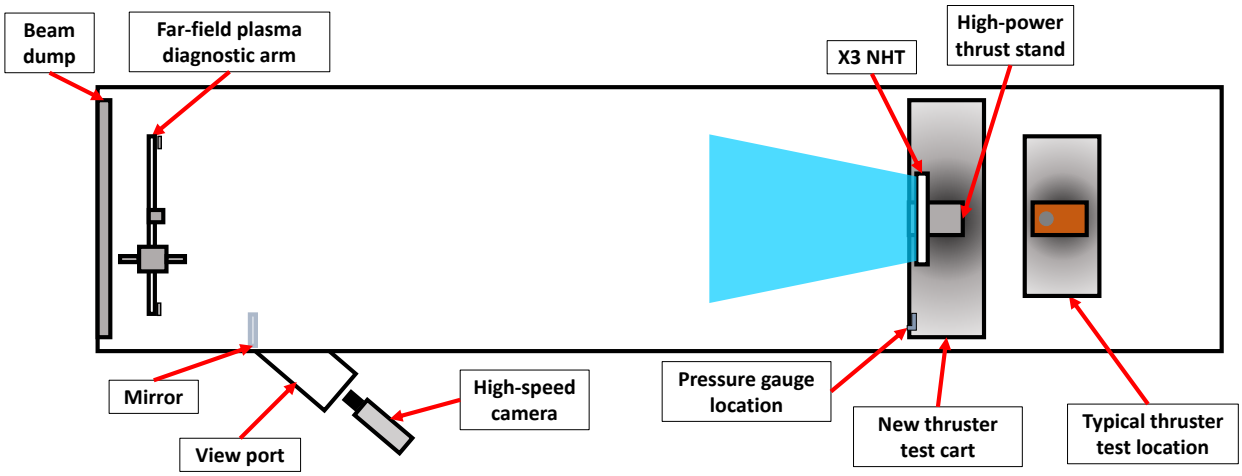

Figure 4.4: A schematic showing the location of the X3 as well as diagnostic equipment inside VF5. Schematic is not to scale.

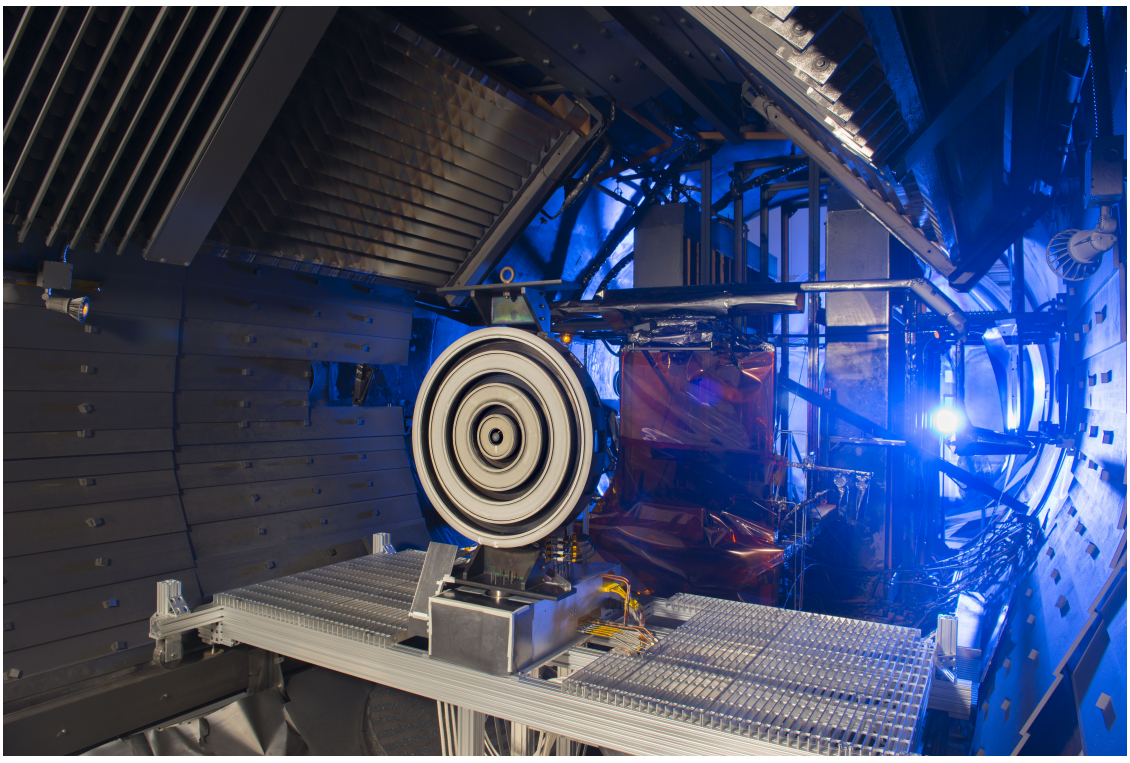

Figure 4.5: The X3 NHT installed inside VF5 at NASA GRC. Photo courtesy NASA.

xenon: a 500-sccm controller for the inner channel, a 1000-sccm controller for the middle channel, a 2000-sccm controller for the outer channel, a 200-sccm controller for the cathode, and a 200-sccm controller for the cathode external injectors. The plumbing matched the setup used during both configurations of LVTF testing as shown schematically above in Figure 4.2. The controllers were calibrated using a dry-cal system before and after the experiment in a similar fashion as the LVTF controllers described above.

4.4.3 Power

We operated the thruster off of a set of laboratory power supplies, which included six separate supplies for the electromagnets, a cathode heater supply, and a cathode keeper supply as described above. Each of the three channels was operated from a separate high-voltage, high-current power supply. The inner was operated using a set of three 1000 V, 15 A supplies that were connected in a master/slave configuration, the middle was operated using a 2000 V, 100 A supply, and the outer channel was operated using a 1000 V, 150 A supply. Additionally, during a small subset of operation, all three anodes were operated with the 150-A power supply. Each discharge channel featured a 100 μF capacitor across the anode and cathode lines. These capacitors helped protect the power supplies from the thruster and allowed the thruster to experience high-current transients without extinguishing.

4.4.4 Telemetry

We measured thruster telemetry in a new breakout box developed at GRC for the NextSTEP program that leverages recent work done for HERMeS thruster development. This breakout box contained precision shunts, voltage dividers, and isolation amplifiers that allowed us to measure discharge, magnet, and cathode telemetry. This telemetry was collected by an Agilent data logger controlled by LabView. Telemetry was recorded at a rate of approximately 0.3 Hz. In addition to the low-speed measurements taken in the breakout box, high-speed measurements of the discharge currents were taken using a set of Tektronix TCP0150 current

guns read by two Tektronix DPO4034B oscilloscopes. The discharge current oscillations were characterized using peak to peak (P2P) and root mean square (RMS) values that were calculated by the oscilloscopes and read by the Agilent data logger. Additionally, synchronized with acquisition of high-speed video as described below, we collected high-speed measurements of the discharge current of each channel, which were used for power spectral density analysis of the dominant oscillation frequencies.

4.4.5 Plasma Plume Diagnostics

We used a stationary far-field plasma diagnostics package to evaluate basic plasma plume properties of the X3. The plasma diagnostics package used here was provided by UM but relied on GRC diagnostic infrastructure and was operated similar to previous GRC plasma diagnostic work [108, 180]. The package was positioned 8.7 meters from the exit plane of the thruster, near the beam dump of VF5. In an effort to limit setup complexity and possible failure modes, the diagnostics package was fixed to the chamber floor instead of mounted to motion stages. Figure 4.6 shows an annotated photograph of the probe array indicating the location of each diagnostic. Below is a description of each of the probes and the data reduction techniques we used to process the results. These techniques generally follow those described by Huang [108, 181], as do the uncertainty analyses.

4.4.5.1 Retarding Potential Analyzer

We used a retarding potential analyzer (RPA) to measure the ion energy distribution function of the plume ions. The RPA is a plasma diagnostic that uses a series of biased and swept grids to selectively collect ions of a certain energy [182]. The RPA used in this experiment is an Air Force Research Laboratory design [108] that features four grids and a collector: a floating grid at the entrance, an electron repelling grid that was biased to 30 V below ground, the ion selector grid that is swept from 0 V to twice the thruster discharge voltage, and the secondary electron emission repression grid that was also biased to 30 V below ground. The

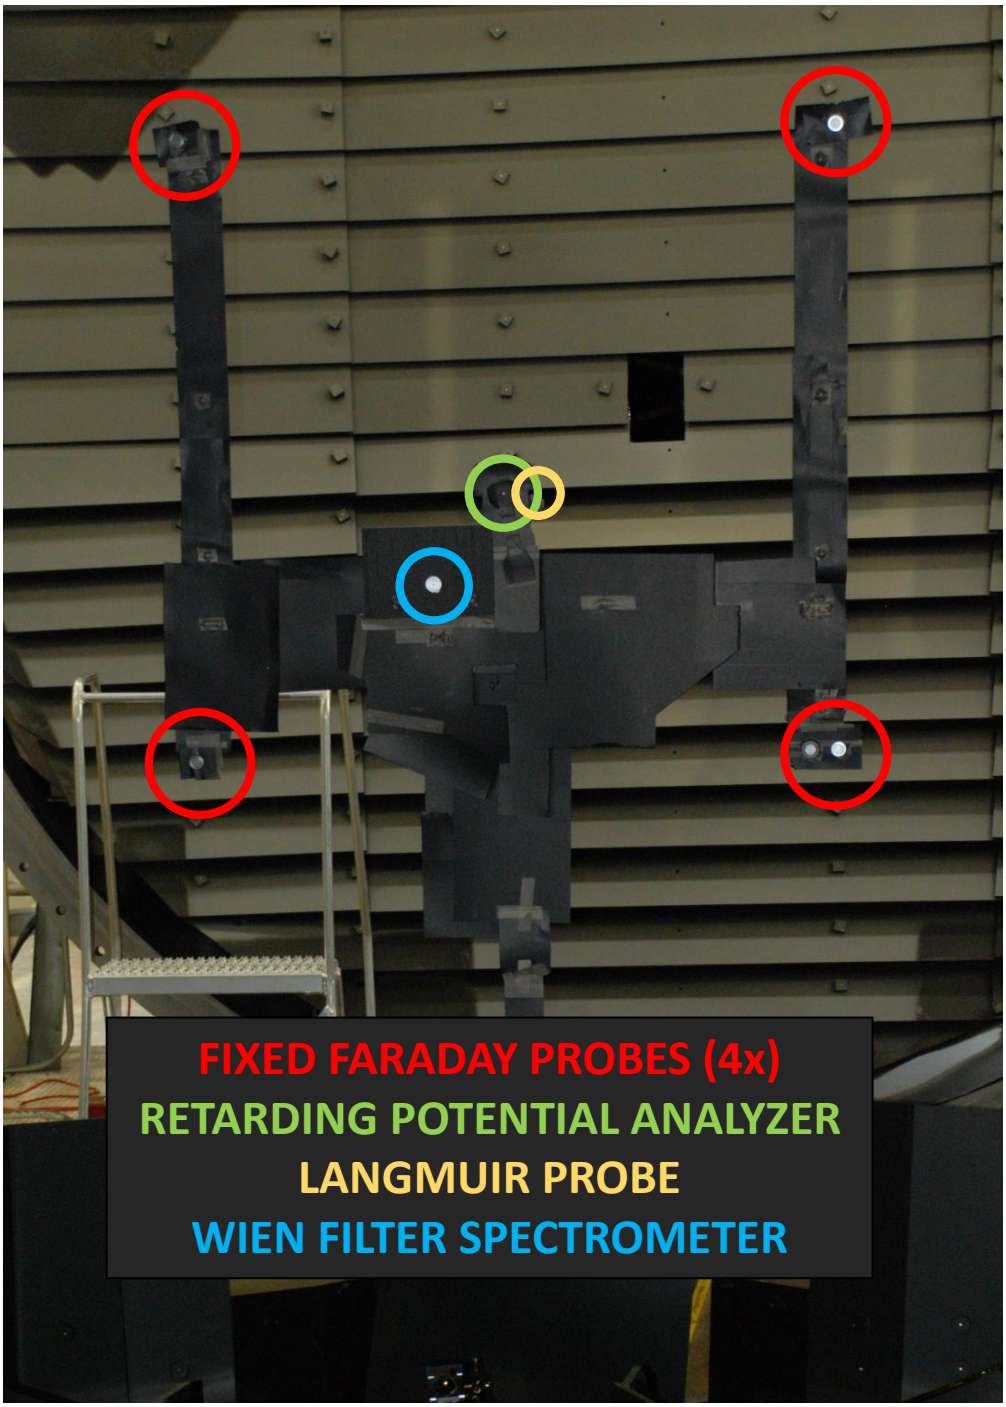

Figure 4.6: A photograph of the far-field plasma diagnostic array mounted in VF5. The location of each diagnostic is indicated.

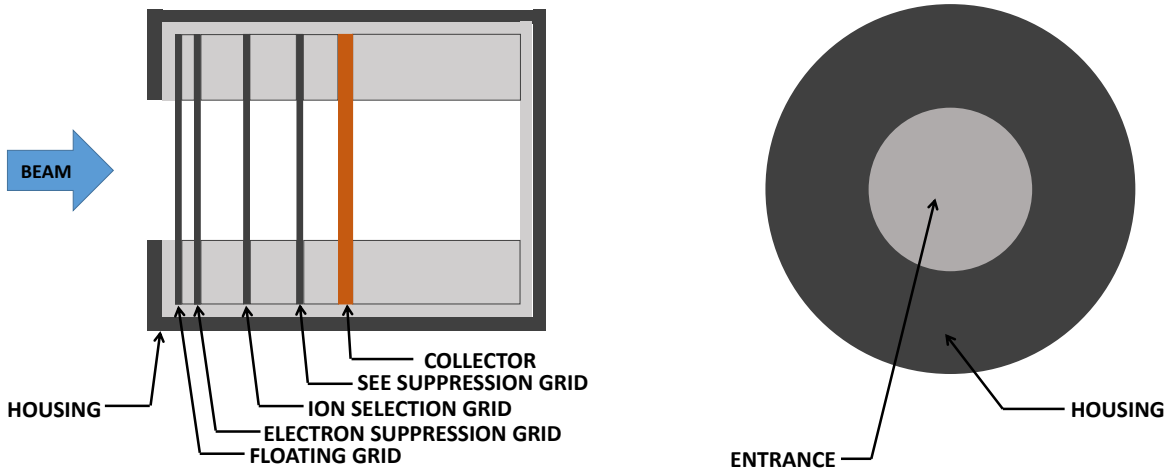

(a) Side view (b) Front view

Figure 4.7: Diagrams of the RPA: (a) the side view and (b) the front view.

RPA was located approximately on thruster centerline. We biased the ion selector grid with a Keithley 2410 sourcemeter and measured the collected current with a Keithley 6485 picoammeter. The RPA had an opening of 6.45 cm^2, but in an effort to minimize the plasma entering the probe at higher power conditions, this opening was reduced to 0.065 cm^2 using a grafoil shield. The acceptance angle of the RPA in this configuration was still large enough that the entire thruster was visible to the probe. Figure 4.7 shows a diagram of the RPA.

We performed data reduction of the RPA traces similar to what was described by Huang [108]. First, we smoothed the raw traces using the Savitzky-Golay method [183], then took a numerical derivative of the collected current I_c with respect to ion selector grid bias V_b. The negative of this derivative ($-dI_c/dV_b$) is proportional to the ion energy distribution function (IEDF) if one can assume a single ion species [182]. For calculation of the thruster voltage utilization efficiency, the RPA is used to calculate the ion energy per charge. Traditionally, most probable voltage is used as a proxy for average ion energy per charge [1, 106]. Most probable voltage V_{mp} is defined as the voltage at which $-dI_c/dV_b$ is maximum (i.e., the peak of the IEDF).

However, the average voltage V_{avg} is a better representation of the average ion energy

per charge. Average voltage and most probable voltage are not expected to be the same for an IEDF that is skewed (i.e., non-Maxwellian), and Hall thruster IEDFs often display varying levels of skew [184]. The reason that V_{mp} is so often used in place of V_{avg} is that the calculation of V_{mp} is significantly less susceptible to noise in the traces (e.g., due to thruster oscillations) or noise-floor issues. Noise is often prevalent in raw traces and magnified by taking the numerical derivative of these data, often to the point of corrupting a calculation of V_{avg}. Huang proposed using a "50% threshold averaging" technique to overcome this impediment [108]. Using this technique, only the portion of $-dI_{coll}/dV_b$ that is greater than 50% of the peak value is used in calculating V_{avg}. This removes the noisiest portions of $-dI_c/dV_b$ from the calculation while retaining a majority of the area under the curve.

Following these recommendations, we calculated both V_{mp} (by locating the peak of $-dI_c/dV_b$) and V_{avg} (by using the 50% threshold averaging technique). Early analysis showed that results (in particular values for V_{mp}) were dependent on the choice of smoothing window size for noisier traces, sometimes varying in magnitude by 5–10 V. To characterize this uncertainty, both V_{mp} and V_{avg} for all traces were calculated across a range of smoothing windows. The mean of these values were taken as the values of V_{mp} and V_{avg} for that condition. This analysis showed that most traces (those without excessive noise) had a V_{mp} standard deviation of 1–3 V and a V_{avg} standard deviation of less than 1 V, demonstrating the 50% threshold averaging technique's avoidance of noise issues. For the noisier conditions, the standard deviation of V_{mp} across the range of smoothing windows could be as high as 9–12 V, whereas the maximum standard deviation of V_{avg} was approximately 5 V. Because of its smaller standard deviation across all conditions and its increased physical relevance, we used V_{avg} for calculating the voltage utilization efficiency.

The bias applied to the ion selection grid in the RPA was with respect to ground. However, the voltage through which ions are accelerated out of the thruster is with respect to the far-field plasma potential. Thus, the V_{avg} value was corrected by:

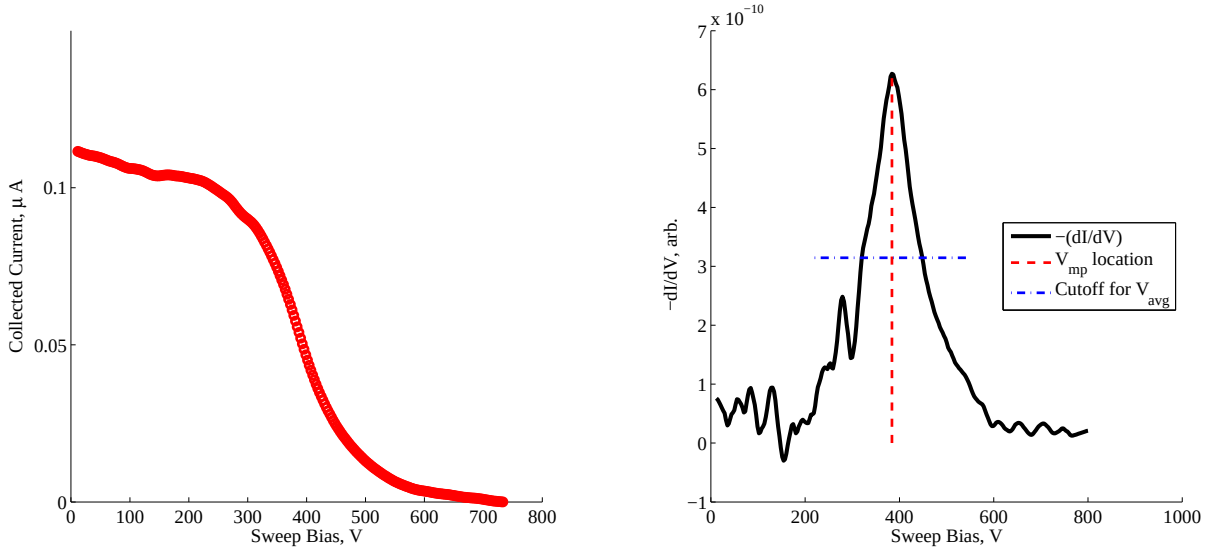

(a) Raw trace

(b) Negative numerical derivative

Figure 4.8: Example RPA data, here for the X3 firing at 80 kW (400 V, 200 A): (a) a raw trace and (b) the resulting negative derivative. The negative numerical derivative plot also shows V_{mp} as identified by the analysis script and the 50% threshold cutoff for calculating V_{avg}.

$$V_a = V_{avg} - V_p, \tag{4.2}$$

where V_a is the acceleration voltage and V_p is the plasma potential, which was measured using the Langmuir probe described below. We then used acceleration voltage V_a to calculate voltage utilization efficiency as defined in Equation 2.14.

Example probe data from the X3 are shown in Figure 4.8. The raw trace (in red on the left) shows the typical shape of RPA traces in Hall thruster plumes [108]. From this, the numerical derivative (in black on the right) shows a strong, clear peak, which the algorithm has identified as V_{mp} as indicated. The blue line indicates the 50% threshold used for calculating V_{avg}.

We found an average uncertainty on V_{avg} of ±1.04 V for our analysis technique. The maximum uncertainty, for a particularly noisy trace, was ±5.2 V, and the minimum was

±0.4 V. An average uncertainty of ±1.04 V corresponds to an average uncertainty in η_v of ±0.003 at 300 V discharge voltage and ±0.002 at 500 V. These uncertainties combine with those from the Langmuir probe analysis to provide the full uncertainty of η_v.

The RPA sweeps took approximately 1 minute to complete, meaning that the results here are time-averaged. The aforementioned thruster oscillations have been shown to move the acceleration region of the thruster and affect the instantaneous ion population at speeds on the order of kilohertz [185–187]. These oscillations are not captured by the RPA swept in this manner, and the results provide the average ion behavior. Time-resolved laser induced fluorescence is one technique that can provide insight into the changes in the ion population over the course of a single discharge oscillation [188] and is a technique that would be of great interest for future experiments on the X3, as discussed in Chapter 9.

4.4.5.2 Langmuir Probe

As described above, the RPA takes its measurements with respect to facility ground, yet the relevant reference voltage for electric propulsion devices is plasma potential. Thus, we used a Langmuir probe (LP) [189] to measure the plasma potential near the RPA and thus correct the RPA measurements. The LP was a circular planar probe featuring a 3.03 cm^2 molybdenum collector area, which was positioned perpendicular to the beam direction. The LP was swept at 10 Hz from -20 V to + 20 V relative to facility ground in a triangle wave using a function generator driving a bipolar power supply. The collected current was measured using a AD210BN isolation amplifier reading the voltage drop across a 100-ohm resistor. The voltage from this isolation amplifier, as well as the voltage applied to the probe, was measured using a National Instruments cDAQ-9178. Figure 4.9 shows a diagram of the Langmuir probe geometry. For simplicity, we used a spare Faraday probe back plate and collector to create the Langmuir used in this test. We used Ceramabond 671 alumina paste around the edges of the collector to limit the exposed electrode area to the circular face.

We performed data reduction using traditional LP theory [46,182,189]. The analysis was

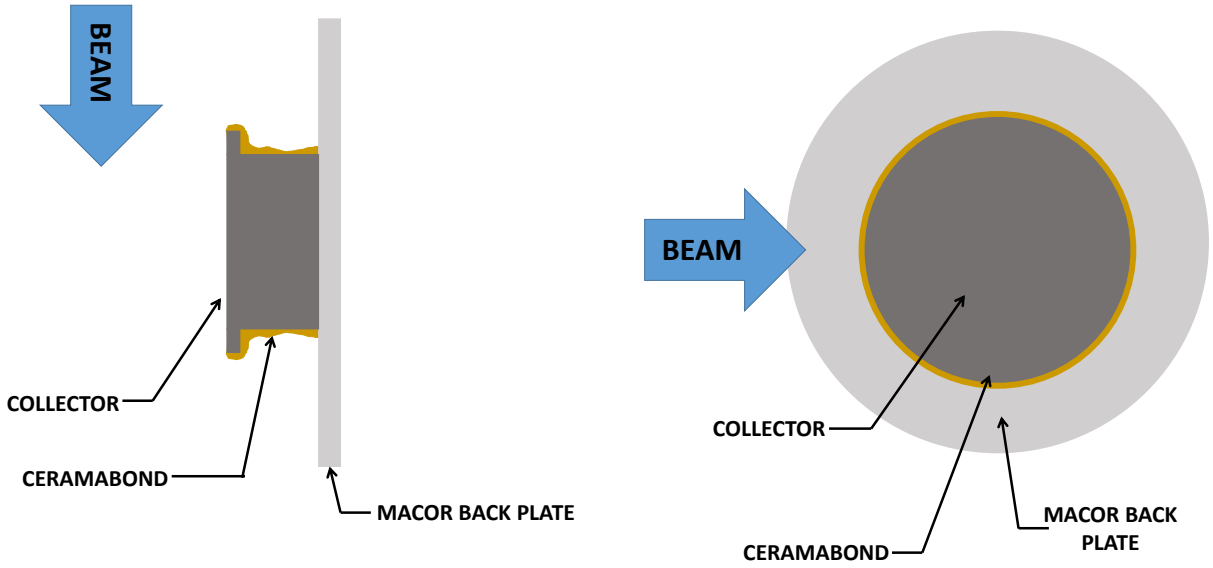

(a) Side view

(b) Front view

Figure 4.9: Diagrams of the Langmuir probe: (a) the side view and (b) the front view.

concerned chiefly with the plasma potential for RPA correction, though floating potential and electron temperature were calculated from the traces as well. Raw traces consisted of ten individual traces (five sweeps up in voltage and five sweeps down over 1 second of data collection at 5 Hz sweep rate). These traces were combined together using a bin-and-average technique with bin sizes of 0.5 V. First, the floating potential was identified. The floating potential is defined as the probe bias at which the electrode is drawing no current:

$$V_f \equiv V_b\,(I_c = 0)\,, \tag{4.3}$$

where V_f is the floating potential, V_b is the probe bias, and I_c is the current collected by the probe. The ion saturation current was next subtracted off the trace. To do so, a least-squares linear fit was applied to the $V_b < V_f$ portion of the trace. This fit is subtracted from the traces to isolate the electron current I_e. We then took the numerical derivative of I_e with respect to probe bias voltage. The location of the maximum of this derivative is defined as the plasma potential:

$$V_p \equiv V_b\left(max\frac{dI_e}{dV_b}\right), \tag{4.4}$$

where V_p is the plasma potential and $\frac{dI_e}{dV_b}$ is the numerical derivative of the electron current with respect to probe bias voltage. To calculate the electron temperature, we determined the least-squares linear slope of the natural logarithm of electron current with respect to probe bias in the linear region between V_f and V_b. The inverse of this slope is the electron temperature:

$$T_e = \left(\frac{d\ln(I_e)}{dV_b}\right)^{-1}. \tag{4.5}$$

The electron temperature does not factor into any thruster efficiency calculations or analysis and was calculated simply for reference.

During LP data reduction we found that the traces displayed varying levels of non-Maxwellian behavior. Figure 4.10 shows an example trace with this behavior. The first derivative of the electron current showed a strong, clear maximum, but this maximum often occurred at a V_b lower than the more obvious knee on the trace. This effect seemed to become more exaggerated at higher power conditions. The LP was oriented normal to the beam as described in the electric propulsion best practices guide [189]. However, this behavior matches what Lobbia describes as the effects of a mis-aligned probe with respect to the beam. To varying degree, a linear region appears in the traces between the maximum of $\frac{dI_e}{dV_b}$ and the visible knee in the I_e vs. V_b trace. The literature shows that this is an effect often seen in streaming plasmas [190, 191]. We attempted a number of alternative V_p identification techniques, including performing two- and three-line least-squares fits on $\ln I_e$ and performing a geometrical technique to identify the visible knee. However, these techniques either continued to identify a V_p below the visible knee (both line fitting techniques) or to overshoot (geometrical technique). Ultimately, we elected to use the V_p identified by the first derivative technique, which was seen as the most reliable technique across the entire

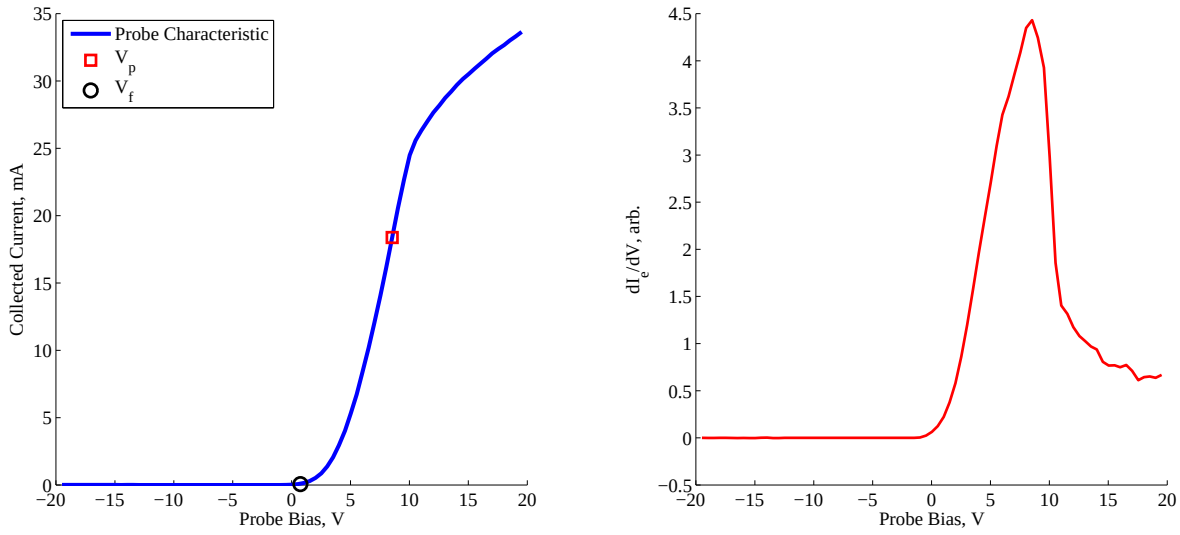

(a) Probe characteristic with V_f and V_p identified

(b) Numerical derivative of electron current

Figure 4.10: Example LP data, showing the non-Maxwellian, streaming-plasma behavior. The plasma potential V_p is identified using the clear maximum in the derivative trace (right), but this point is noticeably below the visible knee on the probe characteristic.

data set and provides a Maxwellian plasma potential value. At most, the variation across the techniques was 4 volts, and was typically closer to 1–2 volts. For the calculation of η_v, this equates to uncertainty on the order of ± 0.01.

Due to the binning process in the LP data analysis method, the uncertainty in V_p was ± 0.5 V at all conditions. This uncertainty contributed an additional ±0.002 at 300 V discharge voltage and ±0.001 at 500 V to the uncertainty of η_v. This brings the total to ±0.015 at 300 V and ±0.013 at 500 V.

4.4.5.3 Fixed Faraday Probes

In an effort to characterize thruster beam symmetry without translating probes through the thruster plume, we mounted four identical Faraday probes of the GRC design [108] at a fixed radius from thruster centerline on the diagnostics arm. These fixed Faraday probes (FFPs) featured a 1.74 cm-diameter molybdenum collector and a molybdenum guard ring with an

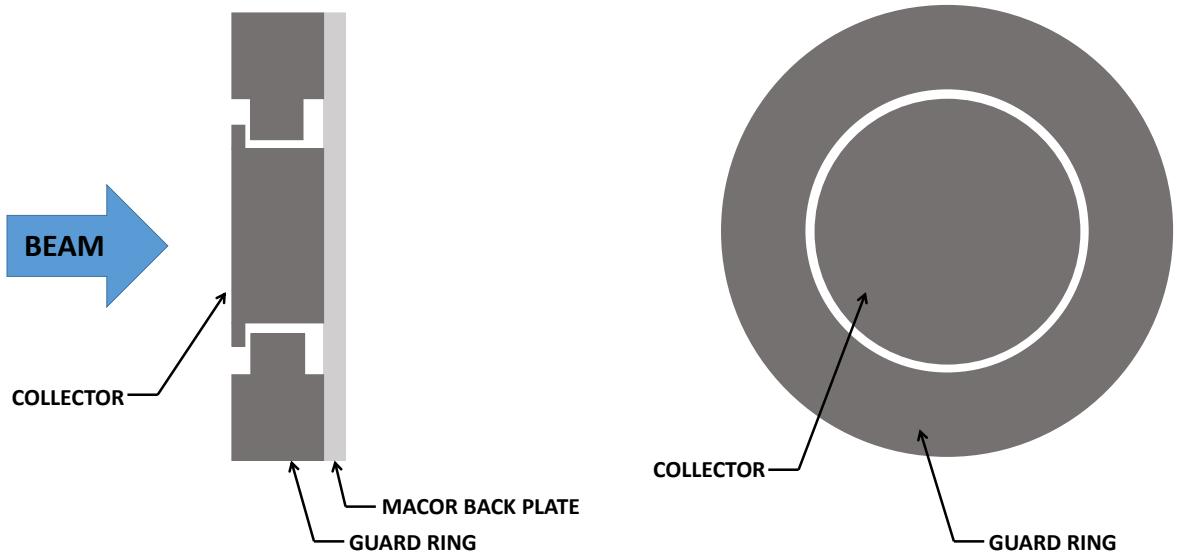

(a) Side view (b) Front view

Figure 4.11: Diagrams of the fixed Faraday probe.

outer diameter of 2.38 cm. The guard ring and collector were separated by a 0.05 cm gap. We biased both the collector and guard ring of each FFP to 30 V below facility ground to ensure the probes were in ion saturation. The collected current was measured using a circuit similar to the LP circuit: it was measured using the voltage drop across a 100-ohm resistor, which was read by an AD210BN isolation amplifier into the National Instruments cDAQ-9178. At each test condition, the current collected by each FFP was measured for 1 second at 2000 Hz. Figure 4.11 shows a diagram of the FFP geometry.

FFP data reduction was straightforward. At each test condition, we took the collected current for a given FFP (I_c) to be the mean value of its trace. These collected current values were converted to current densities:

$$j_{FFP,i} = \frac{I_{c,i}}{A_{probe}}, \tag{4.6}$$

where $j_{FFP,i}$ is the current density at probe i, $I_{c,i}$ is the collected current for probe i, and A_{probe} is the area of the probe collector, taken to be 3.03 cm^2 for all probes. These current density

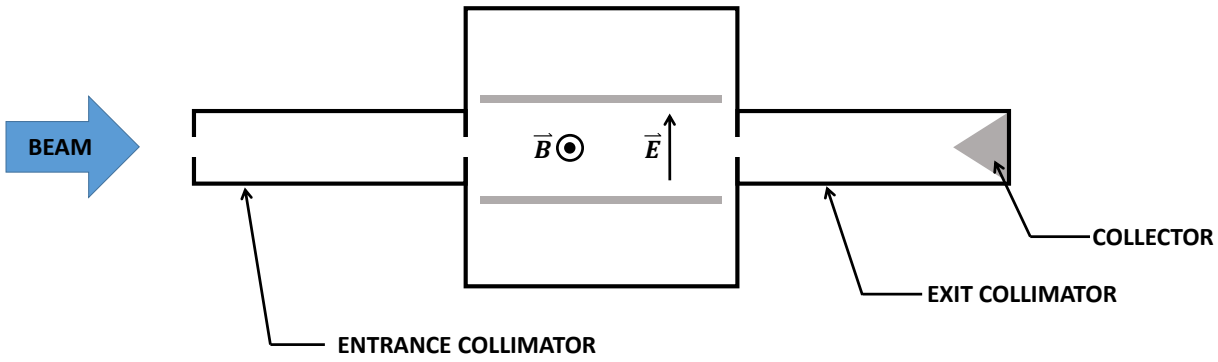

Figure 4.12: A diagram of the Wien Filter Spectrometer.

values were then compared to quantify the variation in beam current density at the beam radius where the probes were located. Uncertainty of the FFP data was taken to be 1% based on the uncertainties of the measurement circuit and the collector area. Other factors such as charge-exchange and secondary-electron effects discussed by Huang in Reference [181] do not affect these measurements because they happen to all probes equally and we are only comparing results between probes at a given condition. These factors only become important when full beam profiles are being measured.

4.4.5.4 Wien Filter Spectrometer

A Wien Filter Spectrometer (WFS) or ExB probe was used to characterize the species fractions of the first four charge states of xenon in the plume. The WFS used was a PEPL probe based on a GRC design and has been used in numerous test campaigns [1, 72, 119]. The WFS features a fixed magnetic field provided by permanent magnets crossed with a variable electric field provided by bias plates. By sweeping the voltage applied to the bias plates, ions can be collected or rejected based on their charge state. We aligned the WFS with the Middle channel of the X3. Figure 4.12 presents a schematic of the WFS design.

WFS data were reduced using a technique similar to that described by Huang [108]. We fit a curve to each of the four detectable peaks (corresponding to Xe^{+}, Xe^{2+}, Xe^{3+}, and

Xe^{4+}) and used those curve fits to calculate current and species fractions for each ion charge state. The X3 WFS traces often (but not always) showed significant peak overlap similar to what Huang found with the NASA-300MS thruster. That study found that a skew-normal distribution provided the best fits and most consistent results for both conditions with and without strong peak overlap, and because of these results we selected it for the X3 data analysis. The skew-normal distribution is a Gaussian fit with an additional fit parameter controlling skewness. For Gaussian curves, this parameter falls to zero and the fit is Gaussian. The skew-normal distribution is of the form:

$$f_i(x) = a \cdot \exp\left[-\frac{(x-c)^2}{2b^2}\right] \cdot \left\{1 + \text{erf}\left[-\frac{d \cdot (x-c)}{\sqrt{2b^2}}\right]\right\}, \tag{4.7}$$

where $f_i(x)$ is the fit result for species i, a, b, c, and d are fit parameters, x is the independent variable (here, x is the probe plate bias V_b), and erf indicates the error function.

The current collected by the WFS for a given charge state i as a function of bias applied to the plates (V_b) is proportional to:

$$I_{c,i}(V_b) \propto V_b^3 \cdot f_i(V_b). \tag{4.8}$$

For the purposes of calculating species fractions, current fractions and charge utilization efficiency, this proportionality is all that is necessary.

We fit a skew-normal distribution to each of the four peaks in order of increasing charge (Xe^+, then Xe^{2+}, etc.). A least-squares fitting technique provided the fit, which we then subtracted from the trace to isolate the next highest peak, a step that was crucial for the traces with significant peak overlap. The curve fit for a given peak was limited to a lower bound of 0.667 of the peak location and an upper limit halfway between the peak location and the location of the next highest peak. For the Xe^{4+} peak, the upper bound was set to be the theoretical location of the Xe^{5+} peak (which was never resolvable in any X3 traces and is not typically seen in Hall thruster plume spectra).

We performed numerical integration of the curve fits to calculate the area under the curve for each fit. These areas were used in calculating current and species fraction for each species. Current fraction was calculated by:

$$\Omega_i = \frac{I_i}{\sum_k I_k} = \frac{n_i Z_i^{3/2}}{\sum_k n_k Z_k^{3/2}}, \tag{4.9}$$

where Ω_i is the current fraction for species i, I is the current for a given species (calculated as the area under the curve for the peak of that species multiplied by a set of physical constants that can be seen to drop out of the calculation [181]), n is the density of a given species, Z is the charge of a given species, and the sums over k are summed over the total number of species present (always equal to 4 here). As noted in Huang's derivation in [181], the physical constants contained within I are equivalent across the peaks of a single propellant, and as such Ω_i can be calculated as the ratio of the area under the curve fit of species i to the summed areas for all species k. Similarly, species fraction calculation followed, using:

$$\zeta_i = \frac{n_i}{\sum_k n_k} = \frac{\Omega_i / Z_i^{3/2}}{\sum_k \Omega_k / Z_k^{3/2}}, \tag{4.10}$$

where ζ_i is the species fraction of species i.

The spectra from the WFS are significantly influenced by charge exchange (CEX) collisions [46, 192] between facility background neutral particles and the beam ions [105]. The beam is attenuated by CEX collisions, which deplete high-energy ions by converting them to fast neutrals via reactions such as:

$$Xe^+_{fast} + Xe_{slow} \rightarrow Xe_{fast} + Xe^+_{slow}. \tag{4.11}$$

The issue is complicated by the fact that CEX collisions attenuate different charge states at different rates due to differences in collisional cross section. Thus, CEX collisions not only attenuate the current collected by the WFS, they actually alter the profile of the spectra through differential attenuation.

Shastry's work showed that these effects increase with increasing background pressure and increasing distance between thruster and probe (both factors increasing the amount of CEX collisions). To minimize facility effects in the plume spectra, Shastry recommended maintaining a "pz" factor less than 2, where p is the facility pressure in units of 10^{-5} Torr and z is the distance from the thruster exit to WFS entrance in units of meters. For chamber pressures on the order of 10-20 μTorr, the recommended probe distance from the thruster is 1-2 meters. High-power operation of the X3 came with high propellant flow rates and thus elevated chamber pressures. Chamber pressure in VF5 during operation ranged from 5–50 μTorr, producing a "recommended" WFS distance from the thruster on the order of 0.5–2 meters. However, this range is in the near-field of the 80-cm-diameter X3, and there were significant probe heating concerns, especially at high-power operation. WFS probes are especially prone to temperature issues due to their permanent magnets, which can lose their magnetic properties at temperatures above their Curie temperature. The geometry of the X3 inside VF5 was such that the probe could not be removed from the plume when not in use. Thus, the WFS was placed with the other far-field diagnostics 8.7 meters from the exit plane of the X3.

Although the CEX attenuation is thus significantly greater than what would be experienced for a recommended pz value, Shastry's correction model is still valid to correct for the attenuation. Shastry developed a "full" CEX correction model and a "simplified" model that took into account only symmetric collisions (i.e., those that that do not produce a new species of ion). Often, Hall thruster analyses use the simplified version [108], but because of the increased significance of CEX collisions on X3 spectra, we employed the full correction model.

The model calculates attenuation factors for each species based on a set of cross sections. These cross sections are provided by:

$$\sigma_1 = 87.3 - 13.6 \log\left(V_d\right), \tag{4.12}$$

$$\sigma_2 = 45.7 - 8.9 \log(2V_d), \tag{4.13}$$

$$\sigma_3 = 2, \tag{4.14}$$

$$\sigma_4 = 87.3 - 13.6 \log(2V_d), \tag{4.15}$$

$$\sigma_5 = 16.9 - 3.0 \log(3V_d), \tag{4.16}$$

where σ is the cross section in units of Å^2 and V_d is the discharge voltage, which Shastry showed was appropriate to use for ion energy for these calculations. These cross sections are used to calculate the following current densities:

$$j_1 = j_{10} \exp(-n_0 \sigma_1 z), \tag{4.17}$$

$$j_2 = j_{20} \exp[-n_0 (\sigma_2 + \sigma_3) z], \tag{4.18}$$

$$j_3 = j_{20} \frac{\sigma_3}{2} \left[\frac{\exp(-n_0 \sigma_4 z) - \exp[-n_0 (\sigma_2 + \sigma_3) z]}{\sigma_2 + \sigma_3 - \sigma_4} \right], \tag{4.19}$$

$$j_4 = j_{30} \exp(-n_0 \sigma_5 z), \tag{4.20}$$

where j is the corrected or actual current density, j_{i0} is the WFS-collected current density for species i, n_0 is the background neutral density in m^{-3} (calculated from chamber pressure measurements) and z is the distance from the thruster exit to probe entrance in m. From these, the attenuation factors of Xe^+, Xe^{2+}, and Xe^{3+} can be calculated as follows (because Xe^{4+} typically constitutes less than 1% of a Hall thruster beam, it is not typically corrected):

$$\left(\frac{j}{j_0}\right)_{Xe+} = \frac{j_1}{j_{10}}, \tag{4.21}$$

$$\left(\frac{j}{j_0}\right)_{Xe2+} = \frac{j_2 + j_3}{j_{20}}, \tag{4.22}$$

$$\left(\frac{j}{j_0}\right)_{Xe3+} = \frac{j_3}{j_{30}}, \tag{4.23}$$

where $\left(\frac{j}{j_0}\right)_i$ is the attenuation factor for species i. These attenuation factors are then used to correct the WFS spectra for the effects of CEX collisions.

Figure 4.13 shows an example WFS trace. Also plotted are the curve fits for the four peaks. This trace demonstrates the blending of peaks that was present in most X3 traces. Here, the Xe^{3+} peak is not visible within the broad tail of the Xe^{2+} peak. By subtracting out the curve fits sequentially during fitting, these hidden peaks can be uncovered, as is shown by the green trace representing the fit for the Xe^{3+} peak. This trace also demonstrates the effect of the CEX correction. The uncorrected current fractions from these data indicated that the population of Xe^{2+} was 11% absolute larger than that of Xe^{+} (0.29 for Xe^{+} as compared to 0.40 for Xe^{2+}), which would be a striking result for a Hall thruster. Yet the CEX correction for the background pressure of $2.5{\times}10^{-5}$ Torr brings these to much more typical values of 0.83 for Xe^{+} and 0.14 for Xe^{2+}.

Uncertainty in the WFS results is mostly due to the uncertainty in the pressure reading used for CEX correction. Huang estimates this to be approximately ±15% for the pressure gauges on VF5 when conductance losses and uncompensated temperature effects are taken into account [181]. A ±15% variation in the facility background pressure measurement results in average uncertainties of the current fractions of ±0.043, ±0.037, and ±0.005 for Xe^{+}, Xe^{2+} and Xe^{3+}, respectively. These, in turn, correspond to an average uncertainty in η_q of 0.004.

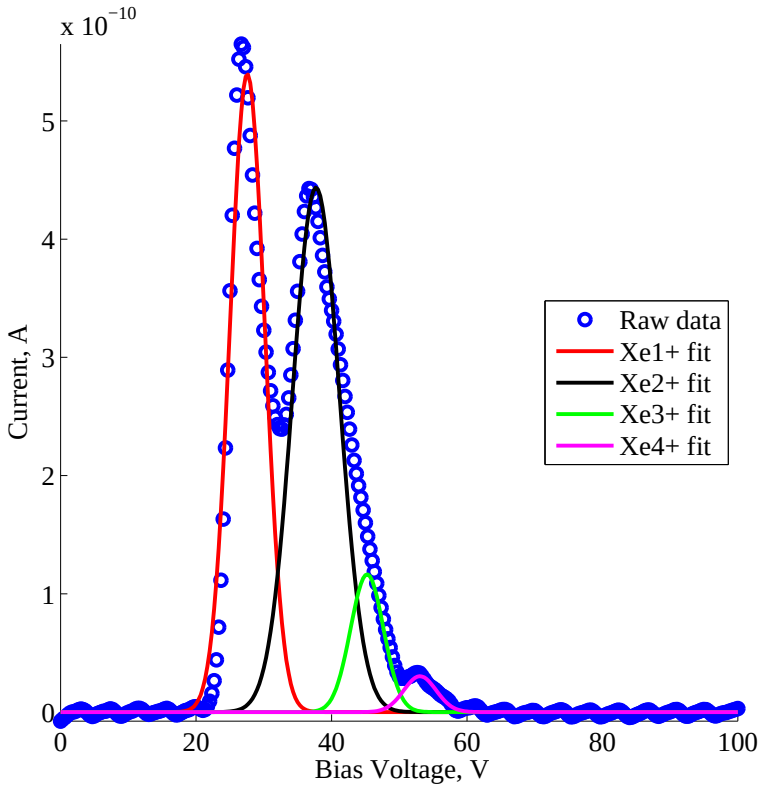

Figure 4.13: Example raw WFS data (in blue) along with the fitted curves for the four Xe charge species peaks.

4.4.6 High Speed Camera

Synchronized with the acquisition of high-speed discharge current measurements with the current guns and oscilloscopes described above we recorded high-speed video of the thruster discharge using a high-speed camera (HSC). A Fastcam SA-Z used in previous experiments with the HERMeS 12.5-kW thrusters [63,67] was positioned at a viewport on the atmosphere side of VF5 and a mirror was placed inside the chamber and aligned such that the camera could see the thruster nearly head-on. Figure 4.4 shows this setup. At each test condition, camera was adjusted such that the entire discharge was captured. The camera was operated at a frame rate of 180 kfps and an image size of 256 pixels by 256 pixels.

For HSC data analysis we followed a procedure that was originally developed by McDonald [61], expanded by Sekerak [60], and used by others [63, 193, 194]. McDonald applied this technique to the two-channel X2 thruster in a previous study [195]. Here, we extend the technique to accommodate up to three simultaneously operating discharge channels. Data analysis took place in Matlab and followed this basic procedure:

1. **Data import:** 2000 frames of the videos were imported. This value was chosen to properly capture relevant phenomena (at 180 kfps, over 100 10-kHz breathing cycles would be captured within the frames) while minimizing computational expense.

2. **Mean image:** The frames were averaged together to create a mean image for the condition. This mean image was then subtracted off of each frame to isolate the AC component of the channel brightness.

3. **Ellipse fit:** A least-squares ellipse fit was performed on the mean image to locate the discharge channel. This ellipse was then scaled to provide inside and outside bounds for each channel. A unique set of inside and outside ellipses was created for each discharge channel in multi-channel operation. Figure 4.14 shows a sequence of inner, middle, and outer channel ellipse sets for an example three-channel operating condition. These ellipses are intended to encapsulate the relevant pixels for subsequent analysis.

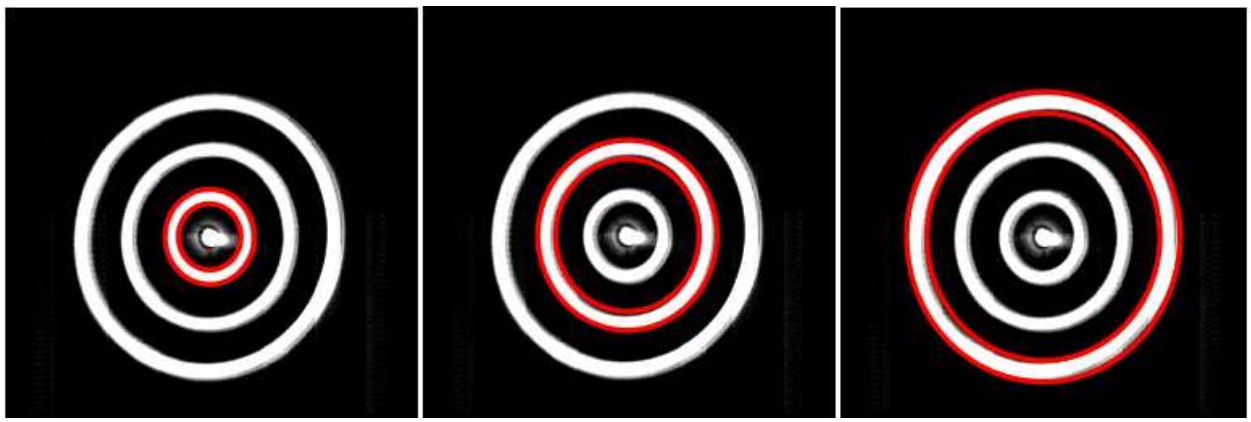

Figure 4.14: Example ellipses fit to identify the inner (left), middle (center), and outer (right) discharge channels from the same mean image. This technique was applied to all multi-channel conditions to isolate each channel.

4. **Bin and average:** The selected pixels were divided into 180 azimuthal bins. For each frame, the brightness of all pixels in a given bin were averaged together.

5. **Viewing angle correction:** Because the thruster was not viewed directly end-on with the HSC, the analysis script made a correction for the viewing angle based on the shape of the fitted ellipse.

6. **Spoke surface:** A two-dimensional heat map was created for the discharge brightness over time. These spoke surfaces (or ϕ-t diagrams) are the equivalent of x-t diagrams for cylindrical coordinates and help to visualize any traveling waves within the discharge [63]. Figure 4.15 shows a representative example spoke surface from the X3. Note that in the example plot, all features are vertical, indicating that the discharge is oscillating as a whole (globally) in time. Spokes, which are localized bright spots that propagate azimuthally within the discharge channel, appear as slanted lines in these plots (c.f. Figure 4a in Sekerak [48]).

7. **Fourier transform:** A two-dimensional Fourier transform was applied to the data to produce a series of power spectral densities (PSDs) corresponding the mode of the oscillation. Following McDonald [61], the mode number m is defined as the number

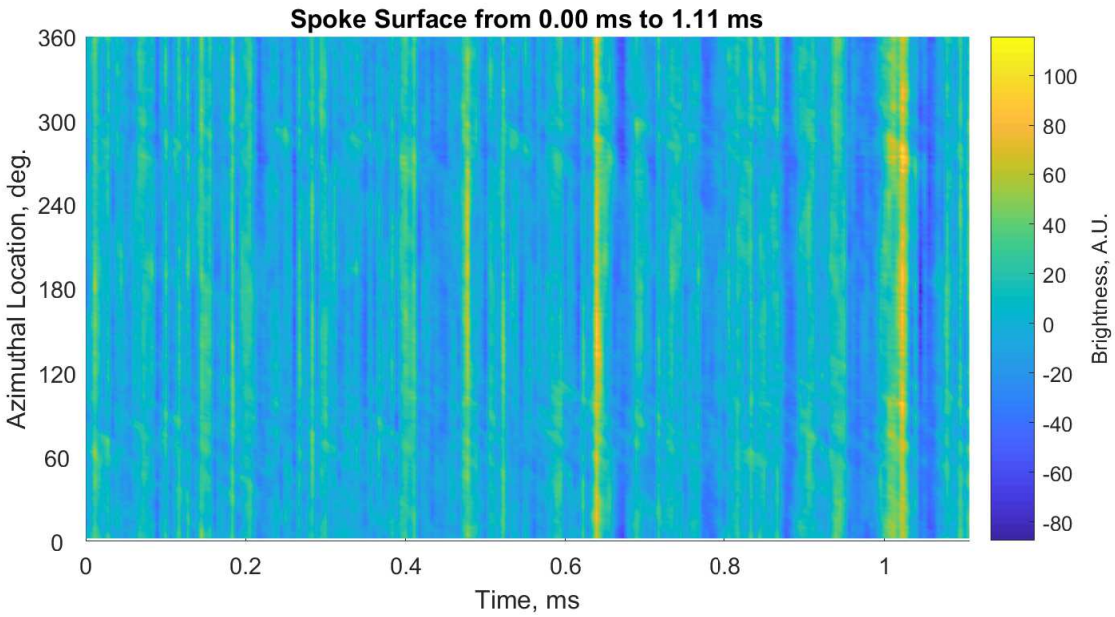

Figure 4.15: An example spoke surface plot from the HSC data analysis. Channel location in degrees is along the ordinate and time is along the abscissa. The relative brightness of the plot corresponds to the relative brightness in the channel. This example shows breathing (global) oscillations with no presence of spokes.

of localized bright spots in the discharge channel. Thus, a m=0 mode corresponds to the entire discharge channel oscillating in time, a m=1 mode corresponds to a single spoke propagating around the channel, etc. Figure 4.16 shows example results from this analysis. In the example, the m=0 mode can be seen to be dominant over any spoke modes ($m \geq 1$) and features two peaks, one at a typical Hall thruster breathing frequency of about 10 kHz and another at a higher frequency (70 kHz).

8. **PSD peak fits:** The two peaks in the m=0 mode of the example PSD were typical to a vast majority of X3 operating conditions. In an effort to characterize changes in the location, width, and relative strength of these two peaks with operating condition, a Lorentzian distribution was fit to each peak. From these fits, values for peak location, width, and strength (height) were calculated. Examples of fitted curves for each of the peaks are shown in Figure 4.17.

9. **Correlation study:** The global thruster oscillation trace (the average of all bins in

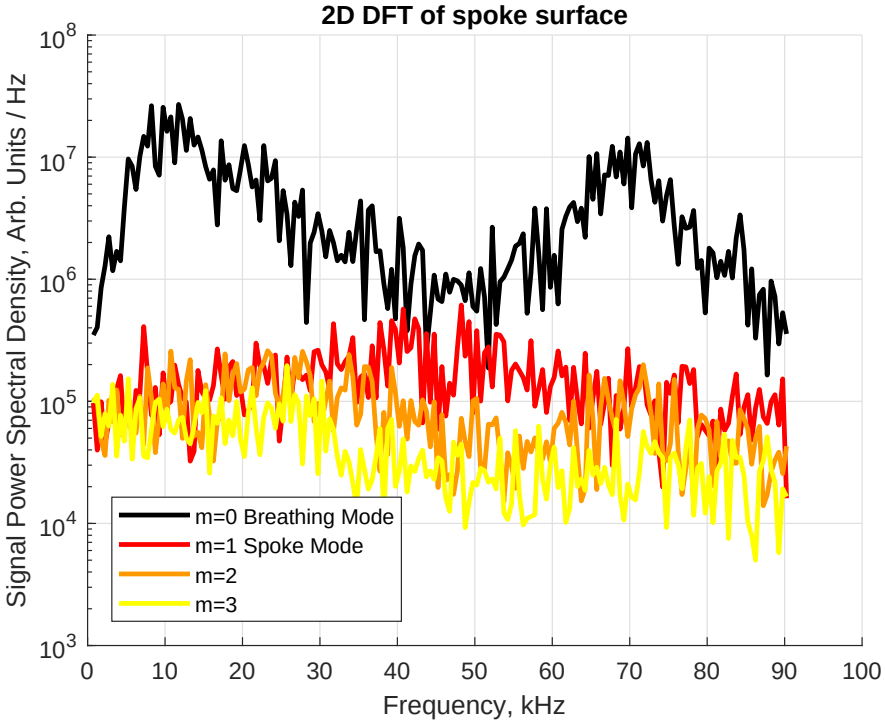

Figure 4.16: Example power spectral densities from the HSC data for m=0,1,2, and 3 modes. Two large peaks, one at 10 kHz and another at 70 kHz, can be seen in the m=0 mode.

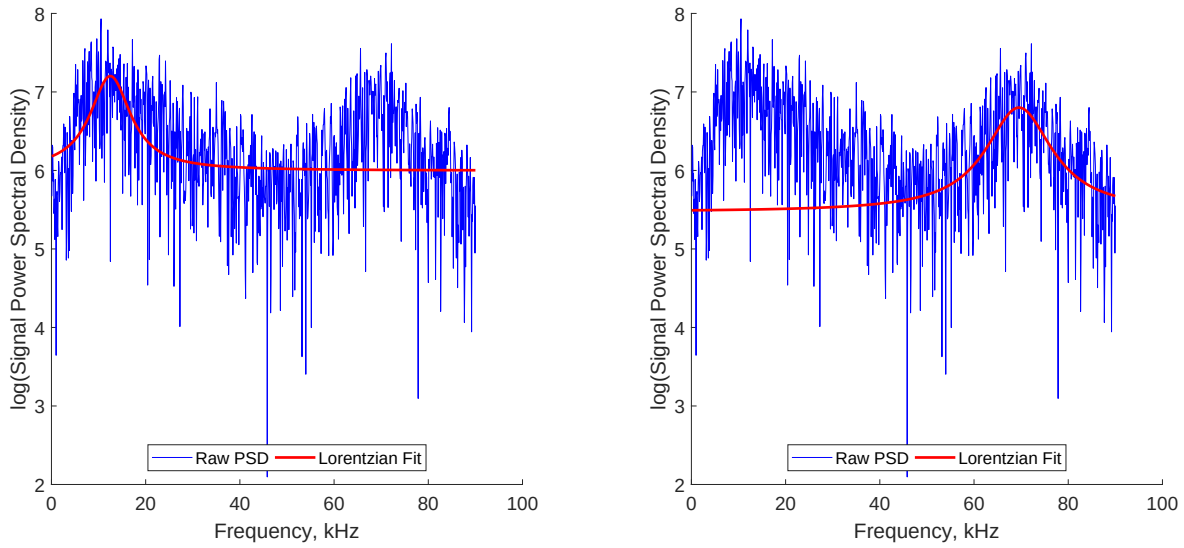

(a) Low-frequency peak

(b) High-frequency peak

Figure 4.17: Example Lorentzian curve fits for the two characteristic peaks in the m=0 PSD: (a) the low-frequency peak and (b) the high-frequency peak.

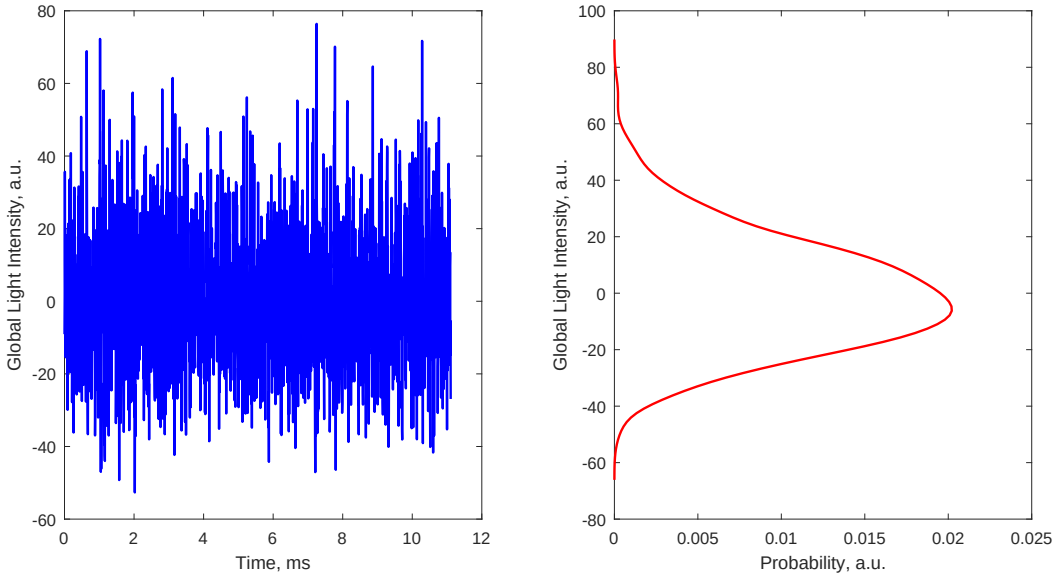

Figure 4.18: Example probability distribution function for the m=0 global oscillation mode on the right, next to the global light intensity signal plotted with time on the left.

each frame) for each channel was cross-correlated to determine whether channels were oscillating in sync or with a phase delay.

10. **Calculate PDF:** Finally, the probability distribution function (PDF) of the global thruster oscillation trace of each channel was calculated, following analysis done by Huang [63]. These PDFs help characterize the nature of the oscillations within the discharge. Purely sinusoidal oscillations have a distinct double-peaked PDF, one that was seen at some conditions on the TDU thrusters in Huang's work. A Gaussian distribution like the one in the example is indicative of seemingly random oscillations that are likely made up of the product of many different sinusoidal oscillations.

4.4.6.1 HSC Cathode Analysis

In addition to the HSC analysis of each of the discharge channels, we attempted to use the same techniques to analyze the cathode oscillations as well. A thorough analysis of cathode oscillations, which often range well above 100 kHz, requires the HSC to be acquiring at a

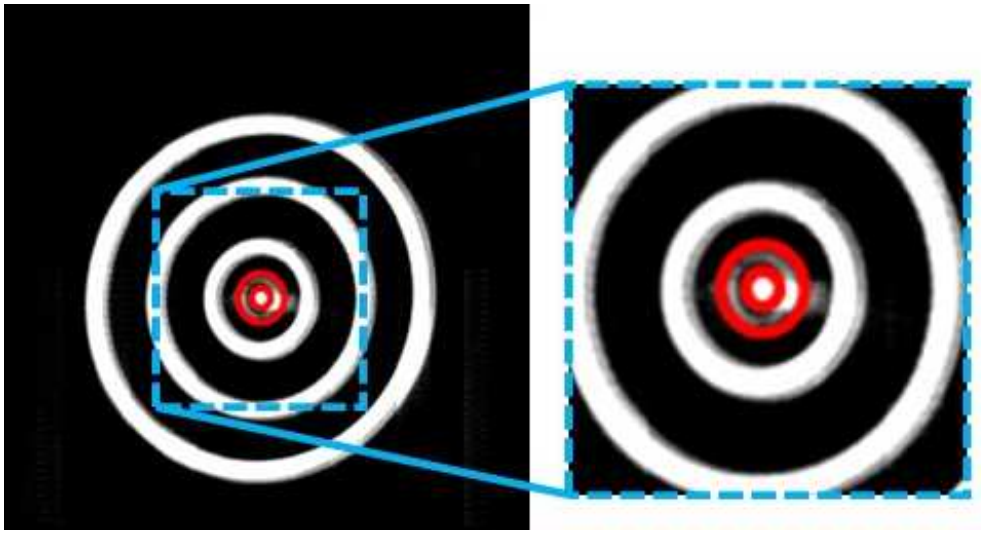

Figure 4.19: Example ellipse to identify the cathode region of the mean HSC image. Due to the viewing angle and the brightness of the cathode spike, no useful information could be extracted from this region.

faster rate than for the discharge oscillations and the camera frame to be readjusted to only capture the cathode region of the thruster [63]. Due to resource limitations during these tests, we instead intended to extract cathode information from the same video frames as the discharge channels. The same basic steps outlined above were used on the cathode region. This region was isolated using a similar ellipse fit as the discharge channels, an example of which is presented in Figure 4.19. Seen in the image, the cathode region of the HSC images has two components: the bright center cathode spike (which extends to the right of the image and occupies the brightest part of the cathode region annulus in the 3–4 o'clock positions) and the dimmer cathode plume (which surrounds the cathode uniformly and makes up the remaining portion of the cathode region). Unfortunately, the combination of viewing angle of the camera, which was such that the large cathode spike intruded significantly into the cathode region, and the intense brightness of the cathode spike, which was so bright during thruster operation that it could not be viewed with the naked eye at higher current conditions, made it such that no useful data could be extracted from this area.

4.5 Thrust Stand

4.5.1 Background

There are a number of different established ways to measure the thrust of a propulsive device. The simplest method is to use a load cell, which measures the thrust force directly. This technique is often used in high-thrust devices such as chemical rockets [196]. Electric thrusters are characterized by their low thrust-to-weight ratios, which are typically much less than unity. This makes using load cells for EP thrust measurement challenging, as the small thrust signal is overwhelmed by the thruster weight or simply not resolvable by the load cell.

A number of thrust stand designs have been developed for EP devices. In general, these designs can be separated into three categories: hanging pendulums, torsional pendulums, and inverted pendulums. The pendulum design is typically used because it allows for high resolution. Of these designs, the inverted pendulum is most often used for Hall thrusters, including the design described here. The reader is directed to Reference [99] for detailed descriptions of each type of pendulum design, including the theory regarding their stability and dynamics.

To facilitate X3 testing and performance characterization, we developed a new inverted-pendulum thrust stand. This stand, called the High-Mass Thruster (HMT) thrust stand, features many characteristics similar to state of the art inverted pendulum thrust stands found in electric test facilities across the industry. However, a number of key components have been modified to accommodate the large mass and expected maximum thrust of the X3. For the initial firing of the X3 by Florenz [4, 120], the inverted pendulum stand inside the LVTF at PEPL was modified to ensure that it was able to support the mass of the thruster. This was intended as a short-term solution, with the long-term vision being two separate thrust stands: the existing stand for low-mass thrusters such as the H6 and the X2 (the "low-mass-thruster" or LMT stand), and a new stand for high-mass thrusters (the

HMT stand). The modifications to the LMT stand were undone to return it to its original condition, and the work described in this section implemented the new HMT stand.

4.5.2 Typical Inverted Pendulum Design

Typical inverted pendulum thrust stands feature two parallel horizontal plates that are connected through a spring and flexure mechanism. The bottom plate is fixed and the top plate is allowed to deflect along the thrust axis. They often contain electromagnetic coils mounted to the bottom plate that actuate the top plate through the application of current. The coils are controlled by a control loop to hold the stand at zero displacement during thruster operation, using the signal from a displacement sensor as the feedback. By recording the applied current necessary to hold the stand at zero displacement (the so-called "null" point), and by calibrating the stand's response through the application of known weights, the thrust being produced by a thruster is known. Inverted-pendulum stands can also operate in "displacement" mode, in which the coils are not used or not installed and the stand is allowed to displace under thrust. The displacement of the stand due to a given thrust force is measured, and using a similar calibration procedure the value of that thrust is known.

For Hall thruster testing, a specific design of inverted pendulum thrust stand called the "GRC-type" is typically seen as the industry standard. This design was originally developed at GRC [197] and has been described [198] and replicated [121, 199] elsewhere across the industry. A schematic of a GRC-type stand is shown in Figure 4.20. The description that follows is specifically of the GRC-style stand at PEPL, referred to here as the LMT stand [121], though the details apply to most GRC-type stands.

The LMT stand features active sensors and electric motors that are controlled from outside the vacuum facility. The stand is enclosed in a copper shroud that is actively cooled by an external chiller/heater unit that pumps coolant through copper lines to the stand. In addition to the shroud, the temperatures of a number of key components of the stand are

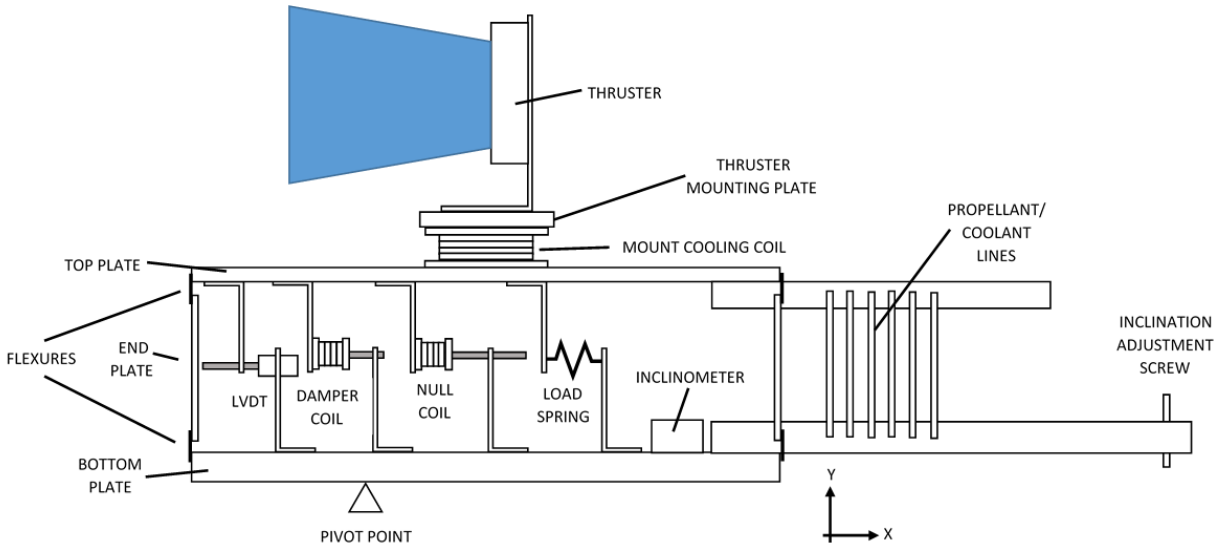

Figure 4.20: A schematic of a GRC-style (or LMT) inverted-pendulum thrust stand for Hall thruster performance characterization. Note that the schematic is notional, items inside the stand are repositioned for clarity, and the figure is not to scale.

maintained with the same chiller. In situ, end-to-end calibration of the stand is accomplished through the application of a set of known weights, which are spooled from an electric motor. Best practices for thrust stand calibrations are detailed in Reference [99] and are not our focus here.

The stand's two horizontal plates are connected using eight thin stainless steel flexures, four of which are connected to each plate. They are linked using a stiff, H-shaped cross-member such that the top plate is restricted to deflecting only in the thrust axis (the x axis in Figure 4.20). As described by Xu [198], the stiffness of these flexures plays an important role in the allowed deflection of the stand and thus its resolution. In addition, they are also major load-bearing columns and must withstand the mass of the thruster without buckling during normal operation or shocks induced during loading. This stand features a "wave"-shaped (or "M" shaped) load spring, the theory of which is described in Reference [198]. This unique shape was designed to minimize the bending moment at each end, where the attachment points are. These springs are typically bent by hand and can be tuned based on the ratio of the lengths of the different components of the spring, the material selection, and

the material thickness. A stiffer spring will reduce the deflection of the stand for a given thrust, but will increase the thruster mass the stand can safely accommodate.

In addition to the flexures and the load spring, the thruster propellant and stand cooling lines play a role in the stand's stiffness. In the LMT stand, these stainless steel lines are bent into spring-like loops that link the bottom and top plates. This allows propellant to be passed to the thruster and coolant to the thruster mount in a way that limits binding or hysteresis. The lines are routed such that the top plate is linked in all locations only to the bottom plate and not to the fixed mounting plate. The inclination of the bottom plate can be actively adjusted while the top plate deflects relative to the bottom plate.

The LMT stand features two electromagnetic coils. Both coils consist of turns of wire around a copper body that is attached to the bottom plate of the stand with a bracket. The coils have actuators consisting of magnets attached to rods that are anchored to the top plate. When current is passed through the coils, a force is applied to the magnet and the top plate is moved. One coil serves as the null coil, which works to counter the thrust force. This coil will be powered throughout thruster operation, and as such, it is actively cooled. The damper coil is used to counteract oscillations inherent to the inverted pendulum's under-damped nature. It only dissipates small amounts of power during thruster operation, and as such is not actively cooled.

Displacement sensing is key to the functionality of the LMT stand. A linear variable differential transformer (LVDT) is typically employed to measure displacement. This signal can be used to calculate thrust directly if the stand is operated in displacement mode. Alternatively, in null mode operation, the displacement signal is used as the feedback signal for a proportional, integral, differential (PID) control loop that holds the stand's displacement at zero by controlling the signal going to the null coil. The measurement from an inverted pendulum thrust stand is susceptible to changes in inclination, since a tilt in the stand will cause a portion of the thrust vector to no longer be positioned along the control axis, replaced by a portion of the thruster weight vector. The LMT stand features a system to precisely control

the stand inclination throughout testing. A precision electrolytic inclinometer is mounted on the bottom plate of the stand to measure the inclination throughout testing. This sensor is typically kept isothermal using the stand's cooling loop.

The bottom plate is placed on two pivot bolts that are located forward of the thruster mount location to provide increased sensitivity for the leveling system. The arm on the bottom plate that the propellant loops are attached to extends out beyond those lines, and at the end a fine-threaded bushing (typically 1/4"–80 thread) is inserted. A fine-threaded rod is inserted into the bushing (as indicated in Figure 4.20) and connected to an electric motor controlled from outside the vacuum facility. Small movement of the motor provides very precise inclination adjustments.

As has been mentioned, the null coil and inclinometer must be held at a constant temperature for proper stand operation. Additionally, the entire bottom plate is held isothermal using the same cooling loop to help prevent against thermal warping, which can introduce error to thrust measurements. The thruster is the dominant source of heat for the stand, so the thruster mount features a cooling loop. This acts as a thermal isolator, removing any heat conducted through the thruster's mount before it can enter the stand. Additionally, the entire stand as depicted in Figure 4.20 is enclosed in a copper shroud that is cooled using a separate loop of the same cooling system. The shroud acts as a Faraday cage for any electromagnetic interference that may come from the thruster during operation, and because it is cooled it isolates the stand from any radiative heat from the thruster or plasma plume.

4.5.3 The High-Mass-Thruster Thrust Stand

4.5.3.1 Overview

The HMT thrust stand is an inverted-pendulum stand based on the GRC-style design. It operates on the same principles: the top plate is allowed to deflect in the thrust axis; a null coil is used in a PID control loop to counteract that displacement due to the thrust force, which holds the stand displacement at zero; the inclination is measured and controlled; calibration

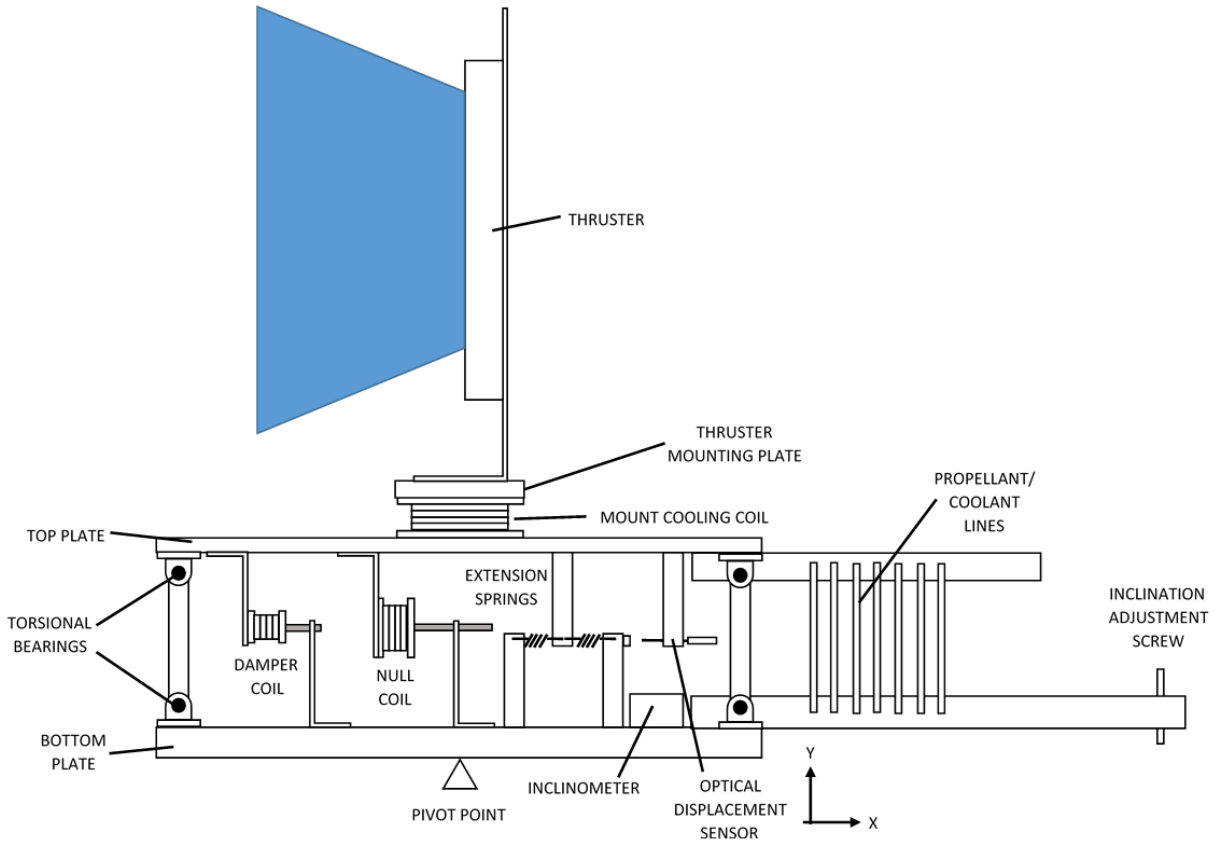

Figure 4.21: A schematic of the University of Michigan's High-Mass Thruster thrust stand. The stand is similar to the GRC-style thrust stand, but differs in a number of key ways. Note that the schematic is notional, items inside the stand are repositioned for clarity, and the figure is not to scale.

is done through the application of known weights; the temperature of key components of the stand is maintained using coolant loop with an external chiller/heater unit. However, a number of components have been changed in the HMT stand, a schematic of which is shown in Figure 4.21.

Differences between the LMT stand and the HMT stand include a flexure design that uses torsional bearings in place of thin stainless steel beams; safety posts to prevent large deflections in the case of stand failure (not shown in Figure 4.21 for clarity); off-the-shelf extension springs in place of the wave-spring; larger coils and magnets to provide increased thrust force; a change in the pivot location; the use of an optical displacement sensor in place of the LVDT; and an aluminum shroud with multiple quick-release panels in place of

the typical copper. The specifics of these differences are detailed below.

4.5.3.2 Upgraded Structural Design

The most significant design driver for the HMT stand was the large mass of the X3. The new design needed to be capable of not only supporting the static weight of the X3 against flexure buckling but being able to withstand the increased shocks that can occur during thruster loading and unloading. The stainless steel flexures were determined to be a weak point in the GRC-style design for application to high-mass thrusters. Though a flexure-based design was used successfully with the NASA-457M testing [10], the further increase in mass associated with the X3 warranted investigation of a new approach. As described by Xu [198], a flexure-based design has been developed for thruster masses up to 250 kg. However, the new approach detailed below provides an extra factor of safety for high mass thrusters such as the X3.

In place of shim-stock flexures, the HMT stand uses a flexure design that was originally provided by the Air Force Research Laboratory (AFRL) for preliminary firing of the X3 [4] which uses double-ended torsional bearings in each corner of the stand. These bearings are connected through rods. These rods are connected perpendicular to the thrust axis (into the page of Figure 4.21) in the front and rear of the stand. These connecting arms serve the role of the H-shaped end plates of the LMT stand. The selected torsional bearings are rated to 280 kg of load each, which provides an adequate factor of safety against collapse for the X3. The bearings also provide a greater restoring force, which is helpful for maintaining the thruster in the null position. As described above, the stiffness of the various stand components (flexures, load spring, and propellant tubing) affects how much thruster mass a given stand design can support. Despite the addition of a fifth propellant line, the propellant routing system in the HMT stand is similar to the GRC-style design and thus not significantly stiffer. The flexures and the extension springs are therefore stiffened substantially in the HMT to provide the increase in load capability.

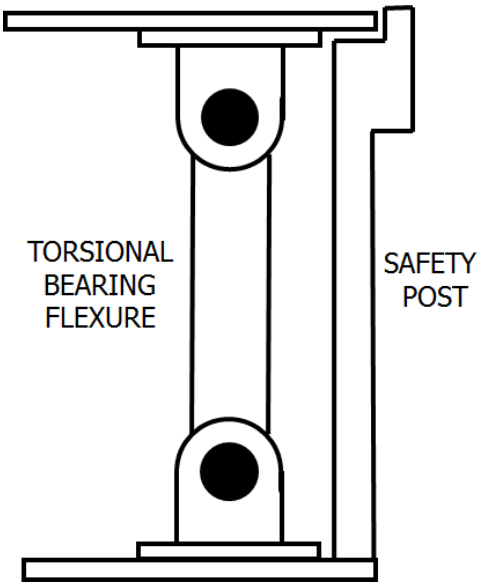

Figure 4.22: An schematic illustrating the safety post and torsional bearing flexure design of the HMT stand.

In addition to the change in flexure design, we added safety posts to the stand as shown in Figure 4.22. These posts are positioned in the front and rear of the plate and provide hard physical stops against both buckling and excessive displacement from the null position. They are machined to provide about 1.5 mm of clearance in the y axis and about 2 mm of displacement in either direction along the thrust (x) axis.

4.5.3.3 Extension Springs

Another change in the design of the HMT stand is that of the load spring. This design was also part of the design provided by AFRL and has been modified and expanded in the HMT stand. The HMT stand uses commercially-available coil extension springs in place of the GRC-type wave spring [198]. Whereas the GRC-type single spring is adjusted to provide a certain amount of pre-load, the extension springs are installed in opposing pairs such that there is a spring (or multiples) pulling in tension in either direction to provide a restoring force to center the stand. The benefits of this design include the repeatability, availability, and wide variety of off-the-shelf parts. However, a hand-made GRC-style spring can be

tuned precisely to the desired spring constant. To increase adjustability of the load spring setup in the HMT stand, two features were added.

The first feature was multiple rows of multiple spring mount holes. This provides the ability to both add multiple springs in both directions and to coarsely adjust the pre-load of the springs. It was found during initial testing of this stand that the pre-load of the springs was a key factor in the balance of the stand. The holes are spaced 6.4 mm apart, and there are twelve different mounting options in each direction. The second feature is a "trim" spring mount, which adds the ability to more precisely adjust the pre-load in each direction. The trim spring mount is on a threaded screw which allows much finer control of the pre-load. The trim spring is intended to be a spring with lower spring constant that is adjusted to precisely balance the stand once the coarse tuning is complete. Figure 4.23 shows a schematic of the spring mounting scheme.

4.5.3.4 Inclination Control

With a thruster the weight of the X3, a misalignment with true level of 0.1° will provide a force along the thrust axis of nearly 4 N, which is approximately 50% of the expected peak thrust of the X3 of 8 N at 200 kW, 800 V. Additionally, the X3 is also capable of operating stably at powers as low as 2 kW and producing thrust values on the order of 200 mN in these conditions. Inclination misalignment is especially detrimental for these low thrust conditions.

Thus, precise and reliable inclination control is key. The HMT stand uses a system very similar to that used in the LMT stand, in which a 1/4"–80 fine-threaded rod is attached to an electric motor. This provides an inclination adjustment step size of approximately 0.3 arcsecond. The LMT stand featured a DC electric motor for inclination control instead of a stepper, which limited the precision of inclination adjustability. In its place, the HMT stand uses a stepper motor capable of 400 steps per revolution. The inclinometer (which is the same model in both the LMT and HMT stands) provides a resolution of approximately 3

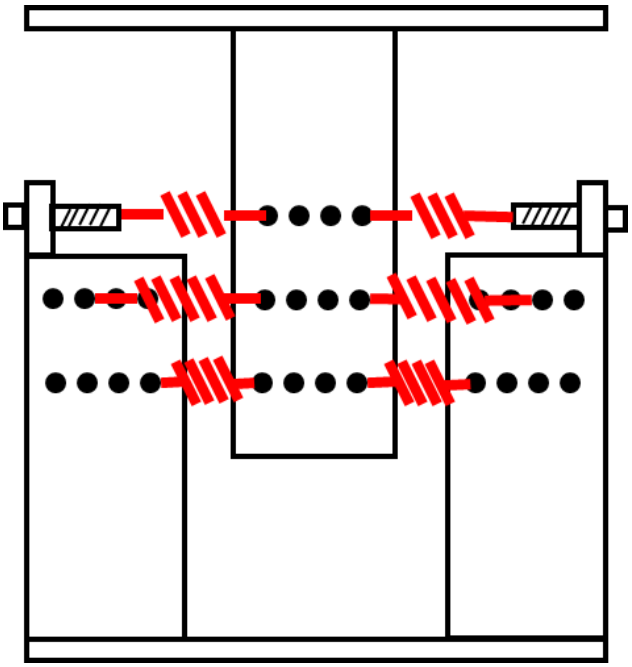

Figure 4.23: A schematic of the spring mounting scheme in the HMT stand. The middle member is attached to the top plate of the stand, and the left and right members are attached to the bottom plate. Six spring positions are shown, in which the top-most two springs are trim springs (attached to threaded rods as shown) and the remaining four locations are for higher-spring-constant load springs. On the other side of these members, six additional identical spring mounting positions are available.

mV per arcsecond of tilt.

In addition to the change in inclination adjustment motor, the pivot point for the stand was moved to be directly underneath the center of gravity of the thruster (i.e., directly under the thruster mount). This provided a direct load path through the pivots and into the table on which the stand is mounted. The LMT stand has the pivot forward of the thruster mount, which distributes a small amount of thruster mass to the inclination screw but increases the inclination adjustment resolution by extending the adjustment arm length. Though the HMT stand loses a small amount of inclination adjustment resolution by shortening the arm length, the added security of the direct load path was seen as worth the sacrifice.

Previous PEPL stands have relied on manual inclination adjustment in which the motor is actuated by the operator at intervals necessary to keep the inclination within the desired bounds. The new HMT stand instead uses a simple control loop executed by LabVIEW, which provides constant monitoring and adjustment. Whenever the inclination signal steps out of the bounds set by the operator, the stepper motor is actuated to move the signal back to zero. Depending on the value of these bounds, the stand inclination usually adjusts every few minutes, as well as in response to any step-changes in thruster operation condition.

4.5.3.5 Null Coil Design

Stands of the GRC-style design have been specified as being capable of measuring up to 5 N of thrust [197,198], so designing a stand capable of measuring up to the 8 N of thrust expected of the X3 did not require a significant overhaul of the actuation mechanism or design. The HMT stand uses the same basic principles described by Xu, featuring an actively-cooled wire-wound null coil setup. However, instead of a single null coil, the HMT stand features two separate coils in series. This setup provides two benefits: it allows for symmetric application of the restoring force (as there was not space for a single coil on stand centerline) and allows for double the force at a given current (thus reducing the current necessary at large thrust values).

The stand was designed to allow for easy coil swapping. This way, multiple sets of coils can be wound to provide higher resolution at low thrusts (fewer turns of wire) as well as low enough current at high thrust to avoid thermal issues (more turns of wire). Based on a given expected test matrix, the proper set of coils can be installed in the stand. By scaling the coils in this manner, the thrust stand resolution (quantified in mA of coil current per mN of thrust force) can be kept as high as possible while avoiding excess thermal load or wire current-rating problems. The permanent magnet chosen for the coils also plays a role in how much restoring force is provided per mA of applied current. For the HMT stand, 25.4-mm-diameter samarium cobalt magnets were chosen for each null coil.

It was found during the design phase of this stand that minimizing the gap between the copper coil body and the permanent magnet actuator increased the force provided by the coil at a given current by increasing the coil/magnet coupling. Thus, coil diameters have been chosen to provide a gap of approximately 3 mm between the permanent magnet and the inner diameter of the coil body.

4.5.3.6 Displacement Sensing

Whether the stand is run in null or displacement mode, an accurate, reliable displacement measurement is crucial for stand operation. In place of the traditional LVDT detailed above, the HMT stand features an optical displacement sensor. This sensor, a Philtec D63, features a fiber optic sensor head that mounts to the top plate of the stand. It reflects off a mirror mounted to the bottom plate of the stand. The sensor head affixes to a fiber optic cable that runs outside of the vacuum chamber to a control box, which outputs a 0–5 V analog signal. The sensor provides a sensitivity of 2.7 mV per μm. Because the system is fiber optic all the way from the stand to the atmosphere-side facility control room, it has been found to be much less susceptible to electrical noise. This is a very significant advantage for the optical system. The LVDT system at PEPL used 15 meters of electrical cabling, whereas the optical system, with the control box situated in the same rack as the data acquisition

system, uses less than a meter.

During initial thrust stand assembly and verification, the sensor tip was positioned inside the stand and the standoff distance from the mirror was set. The gain on the control box was adjusted as per the instructions provided by manual, and a series of tests were conducted to confirm proper response. Once these checkouts were complete, the sensor did not need to be serviced, adjusted, or re-aligned through five months of consistent testing. It provided very repeatable and reliable results across many facility cycles and different thruster operating conditions.

4.5.3.7 Stand Tuning

The dynamics of an inverted pendulum thrust stand will change when the stand is loaded versus when it is not. For smaller thrusters, this effect may be negligible. But with thrusters the mass of the X3, the stand dynamics change significantly when the stand is loaded. To accommodate detailed stand tuning while not placing the thruster at risk, an equivalent dummy weight system was created to allow a single operator to safely add and remove 250 kg of weight from the stand.

The stand's natural frequency and response changed drastically when fully loaded, and the springs required to maintain proper stand balance changed as well. With an X3-equivalent mass on the stand, load springs providing a total effective spring constant of 10.8 N/mm and a pre-load force of 137 N in each direction were required to keep the stand off the stops and moving freely under applied calibration loads. These spring settings were found iteratively by adding spring constant and pre-load in increments until the stand was able to return to the null position even after large displacements.

4.5.3.8 Operation

Versions of the HMT stand were developed for use at PEPL and at GRC for X3 testing. The operation of the PEPL version is described briefly in Reference [200]. It was tuned to

accommodate thrusts of up to 1.5 N and demonstrated linear calibrations with R^2 values typically around 0.99996. The zero drift of the stand was on the order of 1% of the thrust signal across many hours of thruster operation. The uncertainty of the PEPL HMT stand in the configuration tested was approximately ± 11 mN.

The GRC HMT stand featured a number of small modifications and improvements over the PEPL version. Because of the larger expected thrusts for the GRC test, we fabricated custom brass calibration weights with axial through-holes for mounting on a 20-pound test monofilament fishing line. We re-routed the internal cooling and propellant lines of the stand to adopt GRC best practices, which de-coupled the bottom top plate from the mounting structure. We also modified the cooling of the internal plate further such that the inclinometer came before the coils in the cooling path (thus ensuring that the heat from the coils did not affect the inclinometer). Based on data collected throughout the test, the thrust stand was found to have a statistical uncertainty of approximately 2%, plus an additional 14 mN uncertainty due to the resolution of the inclination reading. We performed in situ calibrations of the stand at the beginning and end of each test day, and additionally took zeros periodically throughout the day. Over the course of this test campaign, we found that the thrust stand calibration slope (in mN/V) varied around the mean with a standard deviation of about 2% day to day. Thermal drift of the stand was typically around 1-2% of full scale across a day of operation. In an effort to assess whether electrical noise from thruster operation was affecting thrust measurements, a calibration weight was dropped during operation of the thruster in a particularly oscillatory condition. The thrust value for the weight matched the "thruster off" calibration value to within the thrust measurement uncertainty.

4.6 Magnetic Field Mapping Test Setup

Magnetic field mapping was done on a new dedicated rig constructed at PEPL based on plans provided by JPL. The rig consisted of three 36-inch screw-driven Velmex linear motion

stages used to maneuver a Lake Shore Cryotronics MMZ-2508-UH three-axis Hall probe. The Hall probe was positioned in a custom holder that consisted of a Teflon sleeve pinned by set screws. This allowed for rotation of the probe during alignment. The Hall probe was attached to a Lake Shore Cryotronics 460 gaussmeter. A LabView virtual instrument on a dedicated computer controlled the motion stages via serial connection and read the data from the gaussmeter via GPIB. The motion stages were mounted to a 2-ft by 3-ft Thor Labs optical table. The table was in turn mounted to the bottom of the X3 using a custom-made adapter plate.

The coordinate system used throughout testing was that of the motion stages: X is the axial direction, Y is the horizontal radial direction, and Z is the vertical radial direction, all with respect to the thruster (see Figure 4.24). The Lakeshore probe had a different coordinate system, but the LabView code automatically compensated for this. As such, any references to coordinate directions here will be with respect to the coordinate system of the motion stages.

4.7 Propellant Uniformity Mapping Test Setup

The setup for the propellant uniformity tests consisted of a Stabil Ion vacuum pressure gauge, a Pitot tube attachment for the gauge, a rotational motion stage, a frame to hold the anode, a rotational motion stage to position the anode, Viton propellant lines, and a 3000 sccm mass flow controller. Following previous work in the literature [201, 202], a pitot tube attachment was fashioned to take pressure measurements. This Pitot tube was made out of 0.25-inch outer-diameter stainless steel tubing, attached to a 2.75-inch Conflat flange with a Swagelok-plumbed feedthrough and fashioned with a 90-degree bend to allow for flow settling. The Pitot tube was 167 mm by 87 mm. The Pitot tube features a 45-degree tip similar to that described by Huang [202]. The 45-degree tip serves to deflect any impinging atoms to the side; a square-cut 90-degree tip would reflect gas particles straight

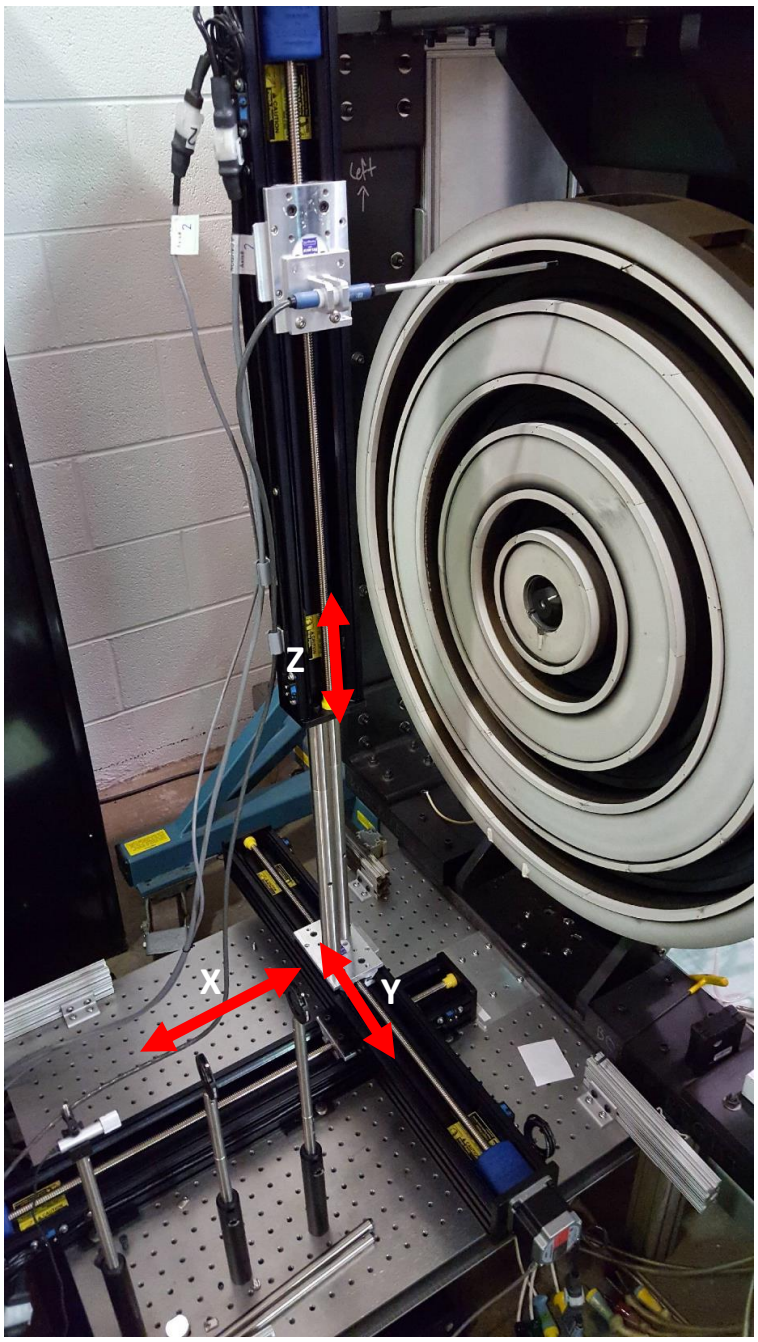

Figure 4.24: A photograph of the experimental setup of the magnetic field mapper. Visible are the X3, the Lake Shore 3-axis Hall probe, the three Velmex linear motion stages, the laser and irises used for alignment (in a "stowed" position for mapping), and the optical table on which the entire rig and thruster are mounted. The coordinate axes are labeled X, Y, and Z on the drawing.

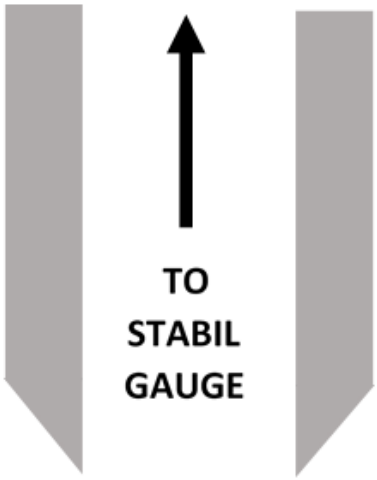

Figure 4.25: Illustration of the probe tip configuration used for the Pitot tube. The tip was ground to a 45° angle such that any gas atoms that impinged on the tip were reflected to the side as opposed to straight down.

back and potentially cause erroneously-high pressure measurements. Though no comparisons to simulation were intended for this experiment, it was determined that matching the test setup of similar experiments would minimize potential sources of error. A schematic of the probe tip is shown in Figure 4.25. The Pitot tube was then attached to the end of the Stabil gauge, which is compatible with 2.75-in Conflat, with a copper gasket to seal the joint. The pressure gauge/Pitot tube assembly was then attached to an aluminum extrusion boom arm. The pressure gauge boom arm was fixed to a 21-in by 32-in by 0.5-in aluminum plate.

The anode to be tested was mounted onto an aluminum extrusion frame affixed to the rotational stage. Previous flow testing, which was done with all three anodes at once, tested the anodes outside of the channel cups to minimize weight and used thin aluminum sheets to provide approximations of channel walls. However, because this test only probed a single anode at a time, weight was not seen as an issue. We thus decided to leave the anodes inside their channel cups (this assembly is referred throughout the remainder of this work as the "anode assembly" for simplicity). This way, the true channel wall contour was provided from the anode exit plane to about halfway down the channel length. The boron nitride segments, which provide the channel wall contour for the remainder of the channel length, were left

Figure 4.26: A photograph of the test setup, installed in the LVTF. Here, the outer channel is aligned and ready for testing. The Pitot tube and Stabil gauge can be seen in the background; the anode and rotational stage in the foreground; and the rigid aluminum plate connecting the two underneath.

out due to their fragility. The rotational stage was mounted to the same aluminum plate as the Pitot probe boom arm, such that the pressure gauge and anode assembly portions of the test setup were rigidly connected. This was done to aid in maintaining alignment through vacuum facility pump down and testing. A photograph of the setup, here with the Outer channel, is shown in Figure 4.26.

Xenon propellant was routed to the anode using the LVTFs existing propellant feed system, which consisted of 0.64-cm-diameter stainless steel lines. A single line was used for these tests which was metered by an Alicat-brand 3000 sccm mass flow controller. At the bulkhead connections for the thrust stand, a Swagelok tee was installed onto the line that spilt the flow into two tubes. Equal lengths of approximately 3.7 m of 0.64-cm inner diameter Viton tubing were affixed to each branch of the tee and routed to the two propellant inlet tube studs on the anode. The length of Viton tubing was chosen to provide enough slack such that the anode assembly could be freely spun a full 360^{o} without the lines becoming taut, crimped, or tangled. The Viton lines were joined to the tee and to the anode propellant inlets using a technique provided by NASA GRC.

This testing occurred during the off-hours of another program that was performing high-voltage thruster operation, during which significant carbon backsputter was apparent within the LVTF. Because the anode and test setup would be inside the facility during many hours of thruster operation, a Grafoil shield was installed to protect the test setup (most importantly the small exit holes of the anode) from carbon backsputter. Visual inspection of the setup and anode after each test indicated that the shielding was working as intended, as no backsputter was detected.

4.8 Summary

This chapter provided an overview of the experimental apparatus used in the characterization of the X3. The X3 was tested in two different vacuum facilities in two different states over three different test campaigns. In both facilities we used a specially-designed inverted-pendulum thrust stand to characterize the performance of the thruster. During high-power testing at GRC we used a suite of plasma diagnostics to characterize the plasma beam of the thruster which included a retarding potential analyzer, a Langmuir probe, a Wien Filter Spectrometer, and a set of four fixed Faraday or ion saturation probes. We used results from these diagnostics to characterize the various phenomenological efficiencies of the X3 across its operating envelope, details of which we provide below. Additionally, we used a high-speed camera to study the oscillation behavior of the thruster. Before this high-power testing, we performed detailed mappings of the magnetic field and propellant distribution of the thruster on dedicated apparatus at the University of Michigan described at the end of this chapter.

CHAPTER 5

Low-Power Performance

"If you [cut corners], it's going to be an illusion that it's going to make it faster... You have to do it step by step, but you want to do it ferociously."

– Jeff Bezos [203]

5.1 Introduction

Prior to recent facility pumping speed upgrades [1], the LVTF at UM was capable of supporting X3 firing at flow rates up to approximately 100 mg/s and powers at or below 30 kW. These are the bounds within which previous firing of the X3 [4,92,120], as well as testing reported in this chapter, were restrained. As such, only the low-power range of the thruster's performance envelope could be tested at UM. In this chapter we present the results of this performance evaluation, which consisted of two experiments. The first, detailed in Section 5.2, was intended to repeat the results from the first firing of the X3 with a higher-fidelity thrust stand. This test evaluated thruster performance at 300 V discharge voltage up to 30 kW discharge power. The second experiment, detailed in Section 5.3, employed a number of lessons learned from the first test and explored improved magnetic field settings for the X3. That test also characterized the thermal profile of the thruster in four different operating configurations with an aim to improve thermal modeling efforts. Both of these tests were

[1] Ongoing as of this writing.

also intended to assess the readiness of the X3 for high-power operation at GRC.

5.2 Low-Power Thruster Characterization 1

Prior to proceeding to high-power operation of the X3, there was interest in repeating the measurements performed by Florenz described in Chapter 2 with an improved thrust stand. This was to more precisely determine the low-power performance of the thruster and identify any areas for improvement or re-optimization before proceeding to high-power testing.

5.2.1 Test Description

During the initial characterization work by Florenz described above, the X3 was operated at two current densities. At the higher current density (for which performance data were not presented by Florenz), chamber pressure was approximately 67 μTorr-Xe. The backpressure limit for Hall thruster performance testing has been suggested to be 30 μTorr-Xe [176]. There were also issues with the cryogenic pumps "crashing" at the higher current density, wherein the pumps would suddenly begin to warm and shed their gas loads, causing a sudden increase in chamber pressure [92]. In an attempt to limit the backpressure of these experiments to as close to this established limit as possible within the LVTF, only the lower current density was tested here.

At these settings, each channel was operated at constant power. That is, discharge current for each channel was held constant, controlled by anode mass flow rate. This constant-power operation was chosen over constant-mass-flow-rate operation because it better accounted for the variation in backpressure across different thruster operational modes. By maintaining power, thrust should theoretically be the same regardless of backpressure. Work by Hofer [111] showed that this was not strictly the case for the H6MS (a single-channel, magnetically-shielded thruster), but that a centrally-mounted cathode reduced the backpressure sensitivity to negligible levels as compared to an externally mounted one.

Table 5.1: The throttling table for low-power thruster characterization 1. All operation was performed at 300 V anode potential.

Condition	$I_{d,I}$	$I_{d,M}$	$I_{d,O}$	$I_{d,T}$	$P_{d,T}$
I	13.5 A	0	0	13.5 A	4.1 kW
M	0	31.5 A	0	31.5 A	9.5 kW
O	0	0	54.8 A	54.8 A	16.4 kW
IM	13.5 A	31.5 A	0	45 A	13.6 kW
IO	13.5 A	0	54.8 A	68.3 A	20.5 kW
MO	0	31.5 A	54.8 A	86.3 A	25.9 kW
IMO	13.5 A	31.5 A	54.8 A	100 A	30 kW

A summary of the test conditions for this experiment are presented in Table 5.1. The thruster was operated at "optimized" magnetic field settings at each operating point based on previous optimization work done by Florenz during the first firing of the thruster. This was to reproduce the testing completed above with improved thrust stand resolution. A common magnetic field across all conditions would help illuminate any performance boost caused by multiple-channel interaction, as in work by Georgin [187] and Cusson [204], but that was not the goal of this experiment.

One major difference between these test points and those presented in Chapter 2 above is that here, the thruster was operated with all magnets on for all conditions. This was due to thruster stability issues discovered during this campaign. This should be noted while comparing these data to those in Chapter 2. Additionally, it should be noted that low-power performance characterization 2 presented below in Section 5.3, this issue was overcome and the thruster was operated with only the electromagnets energized for the channels that were firing in any given condition. During operation, each channel of the thruster went through a bakeout of 1.5 hours before any experimental data were obtained. Due to the large thermal mass of the thruster, no attempt was made to reach full thermal equilibrium before performance measurements were taken (a similar concession was made with the NASA-457Mv1 [10]). Instead, the thruster was typically operated for 30 minutes before data collection, at which point the thruster discharge current was steady and had

stopped drifting.

A version of the JPL-designed LaB_6 cathode described in Chapter 4 was used for these tests. The particular version used during these tests was a 275-A cathode with external gas injectors. For this test, the cathode was operated at a constant 10% TCFF of the total flow rate through all of the anodes. Up to 2 mg/s of propellant was flowed through the cathode center. Any remainder of the 10% TCFF was flowed through the external injectors.

5.2.2 Results

Thruster performance was characterized by thrust, anode efficiency, anode specific impulse, and thrust/power ratio (T/P), all of which are described in Chapter 2. Here, the anode quantities were studied because no effort was made to optimize cathode performance or minimize electromagnet power. In addition to the plots below, full performance data is presented in Appendix B.

Thrust versus discharge power is presented in Figure 5.1. It can be seen that the thrust increases as discharge power increases, which is to be expected. It appears that the M, O, and MO conditions fall onto a thrust/power line with a shallower slope than the other configurations. It is unclear from this plot whether this is evidence of the M, O, and MO configurations under-performing what is expected or the other conditions over-performing. Comparisons are made to other high-power Hall thrusters below which help to illuminate the trends.

Figure 5.2 shows anode specific impulse versus discharge power. This plot indicates that for all multi-channel operating conditions, the anode specific impulse is at least equal to that of the highest relevant single-channel case. The IM condition saw an increase in specific impulse over the Inner channel alone, though for that condition the uncertainty is large enough that the difference may not be statistically significant. The uncertainty in both anode specific impulse and anode efficiency predominantly come from the uncertainty in the thrust measurement.

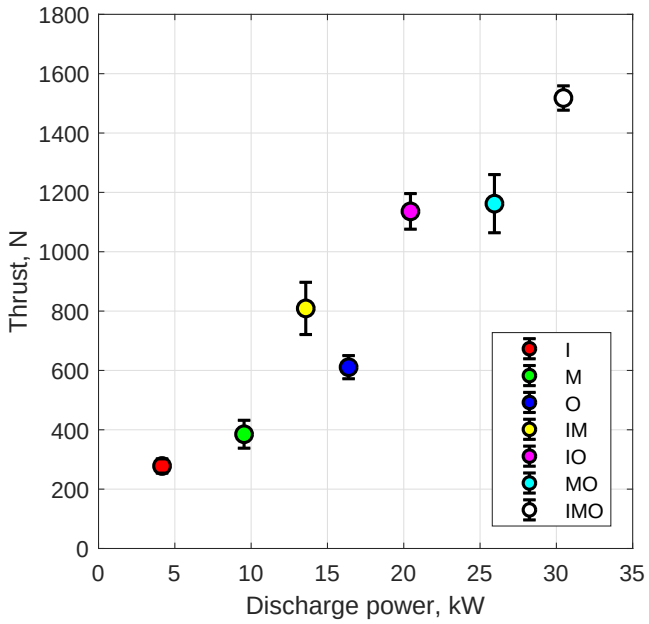

Figure 5.1: Thrust versus power across all operating conditions for low-power performance characterization 1.

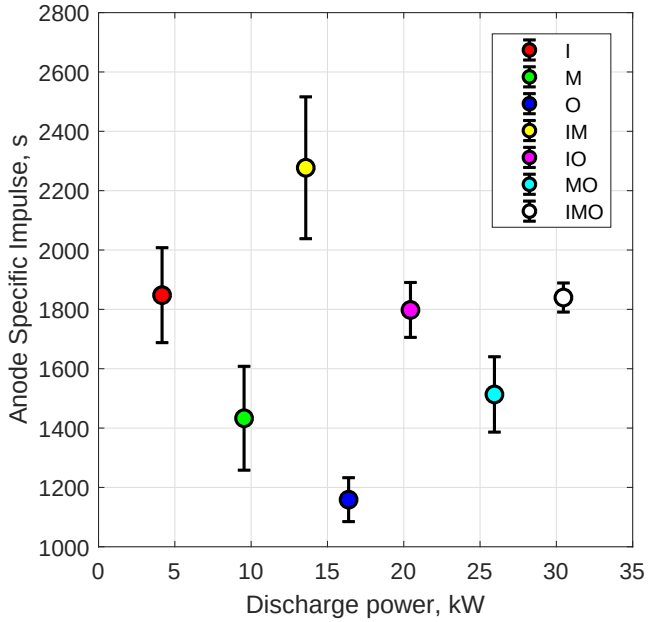

Figure 5.2: Anode specific impulse versus discharge power across all operating conditions for low-power performance characterization 1.

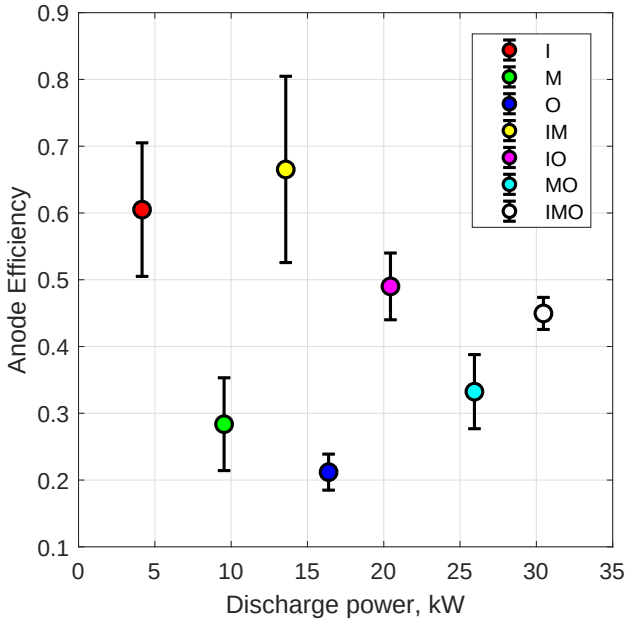

Figure 5.3: Anode efficiency versus discharge power across all operating conditions for low-power performance characterization 1.

Anode efficiency versus discharge power is shown in Figure 5.3. Here, the trend is less clear. The IM condition provides an efficiency higher than that of the inner or middle alone (and in fact the highest efficiency of the entire set); MO also provides a slightly higher efficiency than either channel separately; the remaining multi-channel cases (IO and IMO) have efficiencies that fall between the values of the channels operating alone. The reason for this is unclear. Again, the large uncertainty on two of the conditions obscures any trends.

5.2.3 Comparison to Other Thrusters

To better frame the relevance of these performance data, despite the significantly off-nominal current density, comparison was made to other state-of-the-art Hall thrusters at similar current densities. The thrusters chosen for this comparison were the NASA-300M [75], NASA-400M [76], NASA-457Mv1 [10], NASA-457Mv2 [74], and H6 [106].

The first of these comparisons is presented in Figure 5.4. This indicates that the X3's operation falls along the same general thrust/discharge power line. This is to be expected

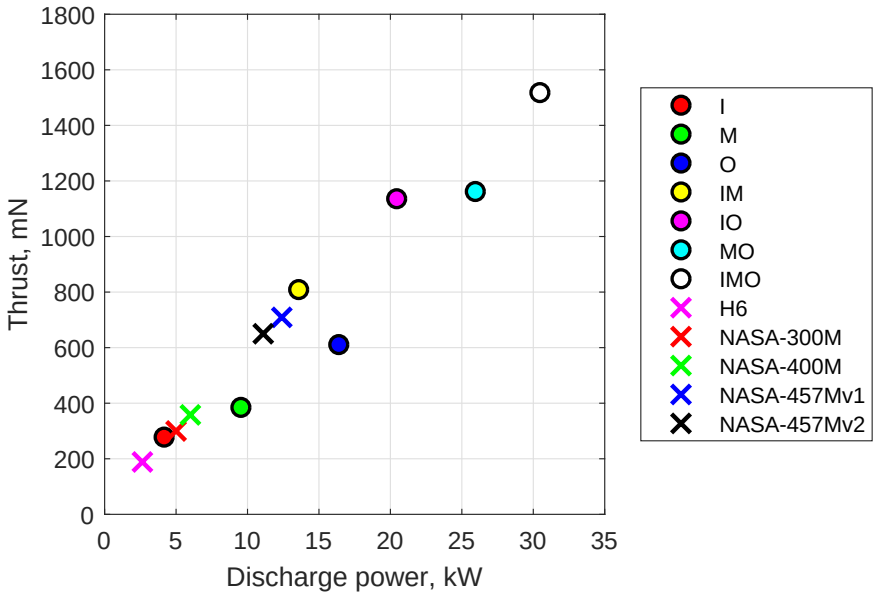

Figure 5.4: A comparison of measured thrust to discharge power between the X3 in low-power performance characterization 1 and previous data from other state of the art Hall thrusters at similarly low current densities.

because the X3's design is based heavily on these thrusters. It can be seen that the IM condition, seen as high compared to the M and O conditions, actually falls closer to these other thrusters than the M or O conditions, suggesting that the M and O conditions are exhibiting poor performance, not that the IM condition is exhibiting a boost.

Additionally, thrust-to-power ratio is plotted against anode specific impulse in Figure 5.5. It can be seen here that other Hall thrusters at these comparable current densities and 300 V discharge voltage operate at specific impulses around 1800 seconds and thrust-to-power ratios ranging from approximately 55 mN/kW to just above 70 mN/kW. This emphasizes that the outlying conditions, specifically the M, O and MO conditions, are operating at both lower T/P and lower anode specific impulse, which in turn indicates that there is likely a major loss mechanism at work in these conditions. Work detailed both in the following section and in Chapter 6 below will show that this low performance is likely due to a combination of poorly optimized magnetic field settings and propellant leaks in the anodes.

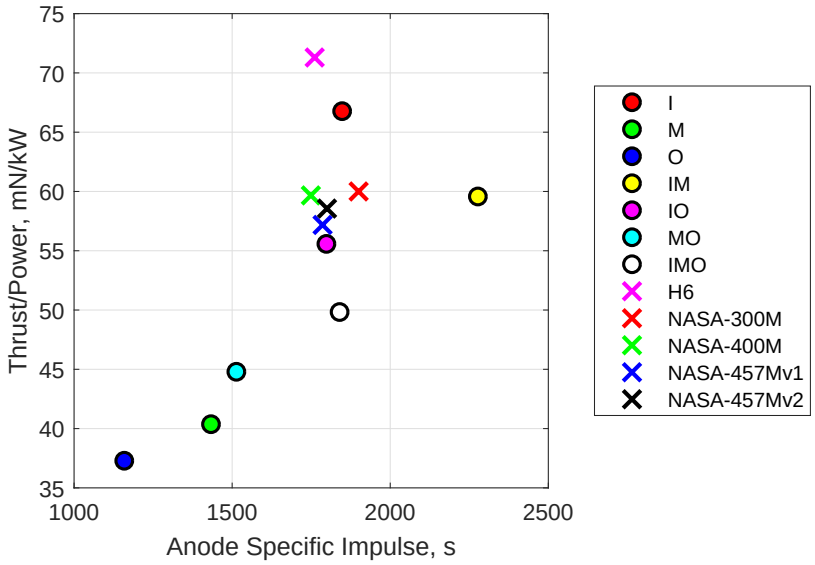

Figure 5.5: A comparison of T/P versus anode specific impulse for the X3 in low-power performance characterization 1 and previous data from other state of the art Hall thrusters at similarly low current densities.

5.2.4 Conclusions

The efficiencies, specific impulses, and T/P of the M, O, and MO conditions were unexpectedly low in performance characterization 1. Reasons for these anomalies were not immediately clear following this test. At the end of this test, it was concluded that significant additional work was necessary with the thruster's magnetic field. During this test, all six electromagnets were operated for all conditions, as it was found that the inner-channel magnets helped stabilize the cathode during middle- and outer-channel operation. However, the thruster was designed to operate each channel with only its two magnets. It was identified that further work with magnetic field settings might allow for stable operation with only the magnets for the channels that are firing [205]. The magnets of non-operating fields had an observable effect on plume structure during this testing, indicating that they likely affected performance as well. A similar effect was seen with the X2 NHT [1].

5.3 Low-Power Thruster Characterization 2

A second low-power test of the X3 was designed as further preparation for high-power testing. First, we desired a thorough evaluation of optimum magnetic field settings. Second, as part of the NextSTEP program testing, we performed a thermal characterization of the thruster to aid in the high-power thermal simulations occurring as part of that project.

5.3.1 Test Description

A main goal of low-power characterization 2 was to re-evaluate the optimization of the magnetic field of the thruster. In analysis after low-power characterization 1, it was determined that the "optimized" coil settings used by Florenz strayed from those that provide an optimum magnetic lens topology. Because of that, it was desired to re-optimize the settings. First, we conducted a series of firings using the thruster's baseline coil settings (which were the coil settings from which Florenz started during his optimization). These test conditions were then repeated with improved coil settings, which were determined by a series of simulations using Infolytica MagNet software described by Cusson [205]. Performance between the two conditions was then compared. In addition to this performance evaluation, we performed a high-voltage checkout of the thruster. This test involved carefully taking each channel of the X3 to higher discharge voltages at a low current density. Finally, we performed a series of firings to thermal steady-state. These runs were primarily for the NextSTEP program. The thermocouple data were used by the thermal modeling team at JPL to fine-tune and validate their simulations of the X3 [93].

During the course of the test campaign, a ceramic electrical isolator on the outer channel was identified as having cracked. Because of its position inside the thruster, it was difficult to determine the extent of the damage without pulling the thruster from the chamber and disassembling it. This was deemed too disruptive to the tight test schedule, so to minimize potential damage to the outer channel hardware due to arcing at high voltage, that channel

was not pushed above 300 V until the end of testing.

This test took place in the LVTF at UM in Configuration 2 described in Chapter 2. Any time after the vacuum facility was vented to atmosphere the thruster was put through a bakeout procedure, which consisted of operation for 1.5 hours at 300 V and the test discharge current on each channel. The high ambient humidity during this test campaign made this bakeout procedure vital. At times during the bakeout, discharge current artificially rose more than 30% over nominal conditions due to moisture outgassing from the thruster ceramics. For any performance data collection, as was done in the first low-power test, the thruster was operated for a minimum of 30 minutes prior to any performance data being collected. This period allowed for the discharge current of the thruster to settle. For the thermal steady-state runs, there was interest in collecting the transient thermal data throughout the run, but the first-priority data were the final temperatures of key components at the end of thruster operation. The criterion for achieving thermal steady-state for these tests was chosen to be a rate of temperature rise less than 10 °C/hour on the thermocouples mounted to the back pole of any firing channel(s). The back pole is the largest continuous mass of material in the whole thruster and was thus expected to be the slowest to reach thermal equilibrium. Thus, it was expected that the other components of the thruster would have temperature rates of change smaller than the back pole at the time of shut off, provided there were no last-minute changes to thruster operating condition (changes in mass flow rate or electromagnet settings). Later analysis by JPL showed that this criterion was not sufficient for the modeling, and the traces we present below indicate that significant temperature changes were still occurring. However, these data still provide estimations of relevant temperatures of the X3.

Two versions of the JPL LaB_6 cathode were used for these tests. For the first portion of this test, we used the same cathode that was used for low-power characterization 1. A heater failure in that cathode caused the remainder of the tests to be run with an updated version, which included an updated heater and orifice plate design that provided an improved current rating of 330 A. Throughout this experiment, the cathode was operated with a 7%

TCFF with a similar scheme for splitting the flow between the main cathode tube and the external injectors. Typically, modern Hall thrusters, especially those on which the X3 design was based, use a TCFF around 7% [100]. The 10% TCFF value had been used in the first low-power test for historical reasons: Liang operated the X2 almost exclusively with a 10% TCFF for stability reasons [1], and Florenz followed that precedent during the first firing of the X3 [120]. However, it was suspected that the X3 would be able to operate stably with a 7% TCFF. That was determined to be true during the first handful of firings during this experiment and it was decided that for the remainder of the test the TCFF would be set at 7% of the total combined anode flow rate.

5.3.2 Results: Thruster Performance

Thruster performance results are split into two studies. The first looks at optimization of the thruster's magnetic field, and the second look at thruster stability at key operating conditions.

5.3.2.1 Magnetic Field Optimization Study

The study of magnetic field optimization itself was split into two sub-components. The first looked at thruster performance using the baseline coil current ratios established during the design phase of the X3 and compared that performance to performance gathered during low-power performance characterization 1. Following those tests, in-depth magnetic field modeling yielded updated coil ratios. This effort, detailed by Cusson [205], worked under the assumption that each configuration (single channel vs. two-channel vs. three-channel) would require a different coil ratio to provide an optimized plasma lens topology. The baseline coil ratios had been established under the premise that the coil ratios were essentially isolated and that the optimized coil ratio for, e.g., the Inner channel in the I condition would be the same as for any multi-channel condition. The more recent modeling showed that to not be the case. As such, new coil ratios for each channel in various configurations were established

Table 5.2: Performance results from a series of magnetic field improvement efforts (the various "2-coil" conditions), compared to those seen above in the low-power characterization 1 (the "6-coil" conditions). A notable improvement in performance is seen for the M and O conditions, whose performance was seen as sub-par during low-power characterization 1.

CONFIG.	COND.	V_d	I_d	T	$I_{sp,a}$	η_a
I	6-coil	300 V	13.1 A	0.278 N	1848 s	0.605
I	2-coil, baseline	400 V	9.8 A	0.212 N	1697 s	0.449
I	2-coil, improved	400 V	9.7 A	0.256 N	1830 s	0.595
M	6-coil	300 V	31.7 A	0.385 N	1433 s	0.283
M	2-coil, baseline	400 V	24.0 A	0.533 N	1848 s	0.504
O	6-coil	300 V	54.7 A	0.611 N	1159 s	0.212
O	2-coil	300 V	51.3 A	0.962 N	1826 s	0.560

and the performance of the thruster with these new ratios was measured. These ratios were developed solely using the computer model, and during this experiment no further ratio optimization was performed while the thruster was operating.

The thruster was operated (and optimized in the modeling effort) with only the coils energized for the channels that were firing. This differed from how the thruster seemed to operate best during low-power characterization 1 but is more in line with what was established during the design of the thruster and during the first firing [4]. Selected results from the series of magnetic field studies are presented in Table 5.2, and are compared to selected relevant data from low-power characterization 1 above.

The test matrix for this test was constructed around the needs of the NextSTEP program, so the test points do not always offer a direct comparison. However, even with that caveat, the improvement in performance with the "two-coil" magnet conditions (that is, only having the electromagnets powered for the channel that was firing) is notable, though not uniform. The Middle and Outer channels both display striking improvements with only their two coils energized. At a higher-voltage, lower-current condition, the Middle jumped from 0.28 anode efficiency to 0.50. The improvement for the Outer channel is larger, and the conditions are more similar than those of the Middle channel. At very similar operation conditions (300 V, 54.7 A versus 300 V, 51.3 A, a discharge current difference of less than 10%), the anode

efficiency of the Outer channel improved from 0.21 to 0.56. The thrust produced increased by approximately 67%. We suspect that the improvement for the Middle channel was likely higher than reported. A substantial leak was detected in the Middle channel propellant line inside the LVTF after this test was complete. Any propellant leaking from the line represents an artificial drop in efficiency because it represents propellant that is not actually being provided to the thruster.

The Inner channel shows a slightly different situation: for 6-coil operation, it displayed a reasonable anode efficiency value, and for the baseline 2-coil case the performance actually dropped. However, the improvements to the magnetic lens topology in the "improved" 2-coil condition bring the efficiency up to a comparable value as the 6-coil condition.

These results suggest that efficient operation of the X3 is attainable through proper tuning of the electromagnet ratios. As discussed in the next section, optimization of the magnetic field strength at a given coil ratio is also key to operating the thruster in a stable, efficient mode.

5.3.2.2 Thruster Stability

Throughout these tests, the magnetic field strength was optimized for each operating point at whichever coil ratio was being used. This optimization also served as a preliminary exploration of thruster stability. At a given operating condition, once the thruster discharge current had leveled off and the discharge current oscillations appeared steady, a sweep of the magnetic field strength was performed. During these maps, the anode and cathode flow rates and discharge voltage were kept constant, as was the coil ratio as defined by Cusson [205]. The magnetic field strength was then adjusted using the current to the electromagnets. The current for both coils was swept simultaneously to maintain proper field topology throughout the map. During these sweeps, thruster telemetry was monitored, including discharge current, RMS of discharge current oscillations, cathode-to-ground voltage, and thrust.

Sweeps of this nature are typically used to optimize thruster operation for a given pa-

rameter. Operation can be tuned via the magnetic field to, for instance, maximize thrust or efficiency, or minimize discharge current oscillations. During these sweeps, it was typically seen that thrust did not change, at least outside the uncertainty of the thrust stand (which as noted previously was approximately ± 11 mN as deployed in LVTF Configuration 2). It should also be noted that cathode-to-ground voltage is not typically a parameter that is optimized using this type of mapping, but is instead tuned by adjusting the TCFF. However, cathode-to-ground voltage is often used as a quantification of thruster-cathode coupling and represents a portion of the cathode-coupling voltage loss (the other contribution is from the plasma potential) [206]. Cathode coupling voltage is typically viewed as an indication of the ease with which electrons are extracted from the cathode and transported into the thruster beam and to its anode [100]. Cathode-to-ground voltage was changing during these maps, indicating a change in the thruster-cathode coupling, and as such, it is reported with the results here.

The discharge current, normalized discharge current RMS oscillations, and the cathode-to-ground voltage are plotted for the Inner channel at 400 V, 9.8 A in Figure 5.6. The RMS discharge current, as measured using the high-speed current diagnostics described in Chapter 4, is normalized by the DC discharge current value, which was measured by the precision current shunt/isolation amplifier setup described in Chapter 4. All metrics are plotted against the magnetic field strength normalized by the nominal magnetic field strength. It can be seen that the discharge current follows a curve similar to that seen in the H6-US unshielded single-channel thruster [48] as well as in the H6-MS and NASA-300MS-2 magnetically shielded single-channel thrusters [12]: as the magnetic field strength decreases, discharge current decreases. What is not seen in this plot is the turn-around point, the point at which continuing to decrease the magnetic field strength causes the discharge current to rapidly begin rising. It is not clear from these data whether that transition does not occur on the X3 or if simply low-enough magnetic field strengths were not attained during this map. The discharge current oscillations show a region of increasing, then decreasing, magnitude

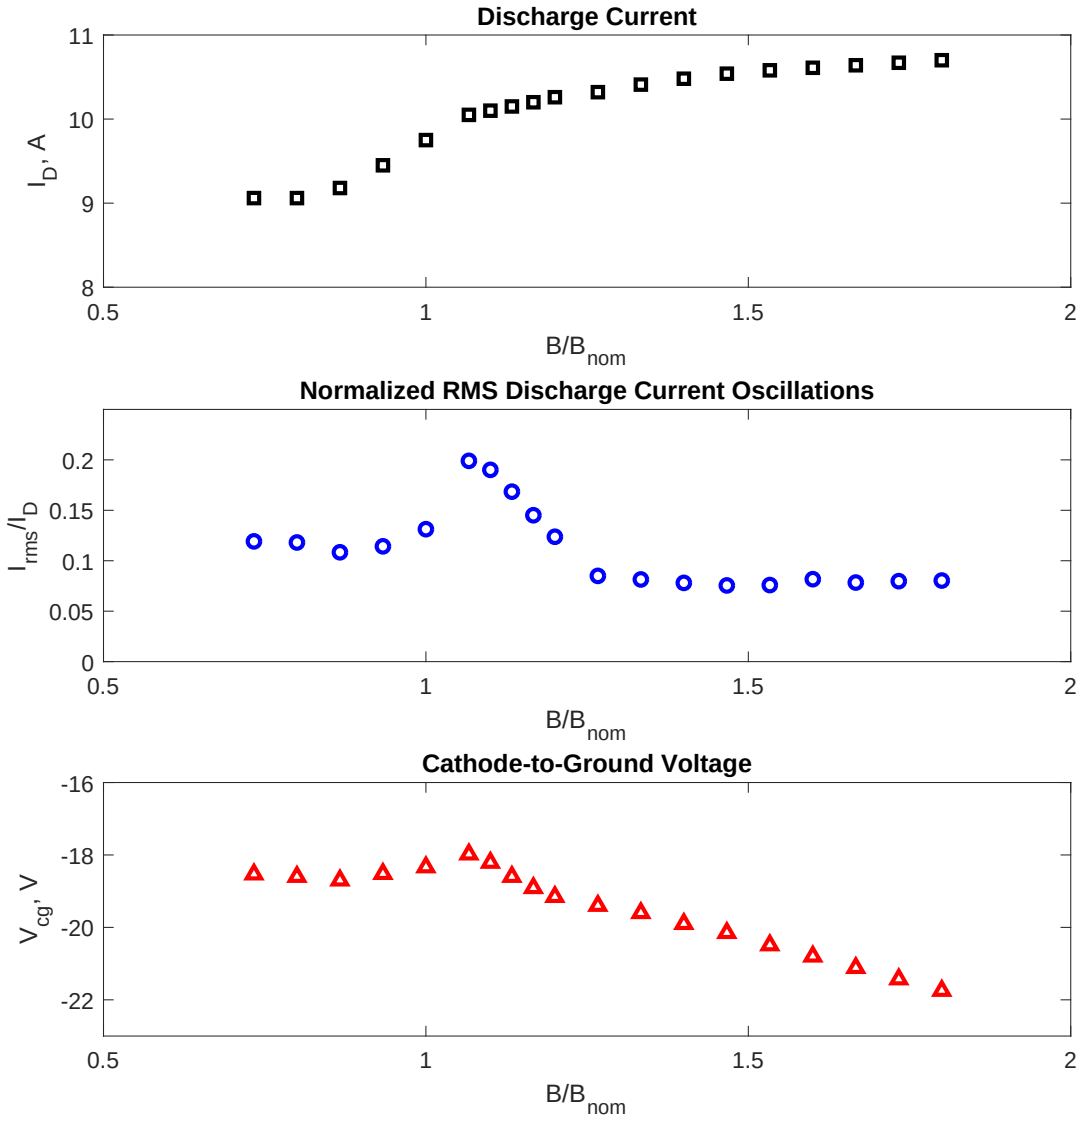

Figure 5.6: The results of a magnetic field optimization map during low-power performance characterization 2 for the Inner channel at 400 V, 9.8 A. Discharge current, normalized RMS discharge current oscillations, and cathode-to-ground voltage are all plotted against the magnetic field strength (which has been normalized by thruster nominal magnetic field strength).

as magnetic field is reduced, but again there is not a point in which the thruster transitions to a largely oscillatory condition. Finally, cathode-to-ground voltage is seen to decrease as magnetic field strength is decreased, followed by a local minimum (seen as a peak on the plot due to the negative nature of the voltage) that coincides with the local maximum seen in the discharge current oscillation trace, and finally with a plateau at low magnetic field strength that coincides with a plateau in the discharge current oscillations.

A matching set of plots for the Middle channel at 400 V, 24 A is presented in Figure 5.7. Similar trends can be seen here. It appears that the magnetic field strength was not lowered far enough to see a similar plateau to the one in the parameters from the Inner channel in Figure 5.6. Otherwise, though, the features appear the same. There is a relatively steady decrease of discharge current with decreasing magnetic field strength; discharge oscillations are low at high magnetic field strength, see an increase at values of B/B_{nom} slightly larger than unity, and then begin to approach similarly low values as magnetic field strength continues to decrease; cathode-to-ground voltage decreases with decreasing magnetic field strength before leveling off at the magnetic field strength corresponding to the peak in discharge current oscillations.

Finally, a set of plots for the Outer channel at 300 V, 40 A are presented in Figure 5.8. The Outer channel map does not cover as wide a range of magnetic field strength as those for the Inner and Middle channels, and some key features seen in those results are not present here. Due to the limited nature of these data, as well as the fact that the Inner and Middle maps were performed at 400 V discharge voltage whereas the Outer map was performed at 300 V, no major conclusions about the Outer channel as compared to the Inner or Middle should be drawn. The discharge current oscillation trace lacks as strong of an increase at lower B/B_{nom}, and the cathode-to-ground voltage never levels off like for the other two channels. However, this may be due to the more limited extent of this sweep. The behavior otherwise is not surprising: discharge current decreases with decreasing magnetic field strength, and cathode-to-ground voltage gets closer to ground (i.e., less negative) with

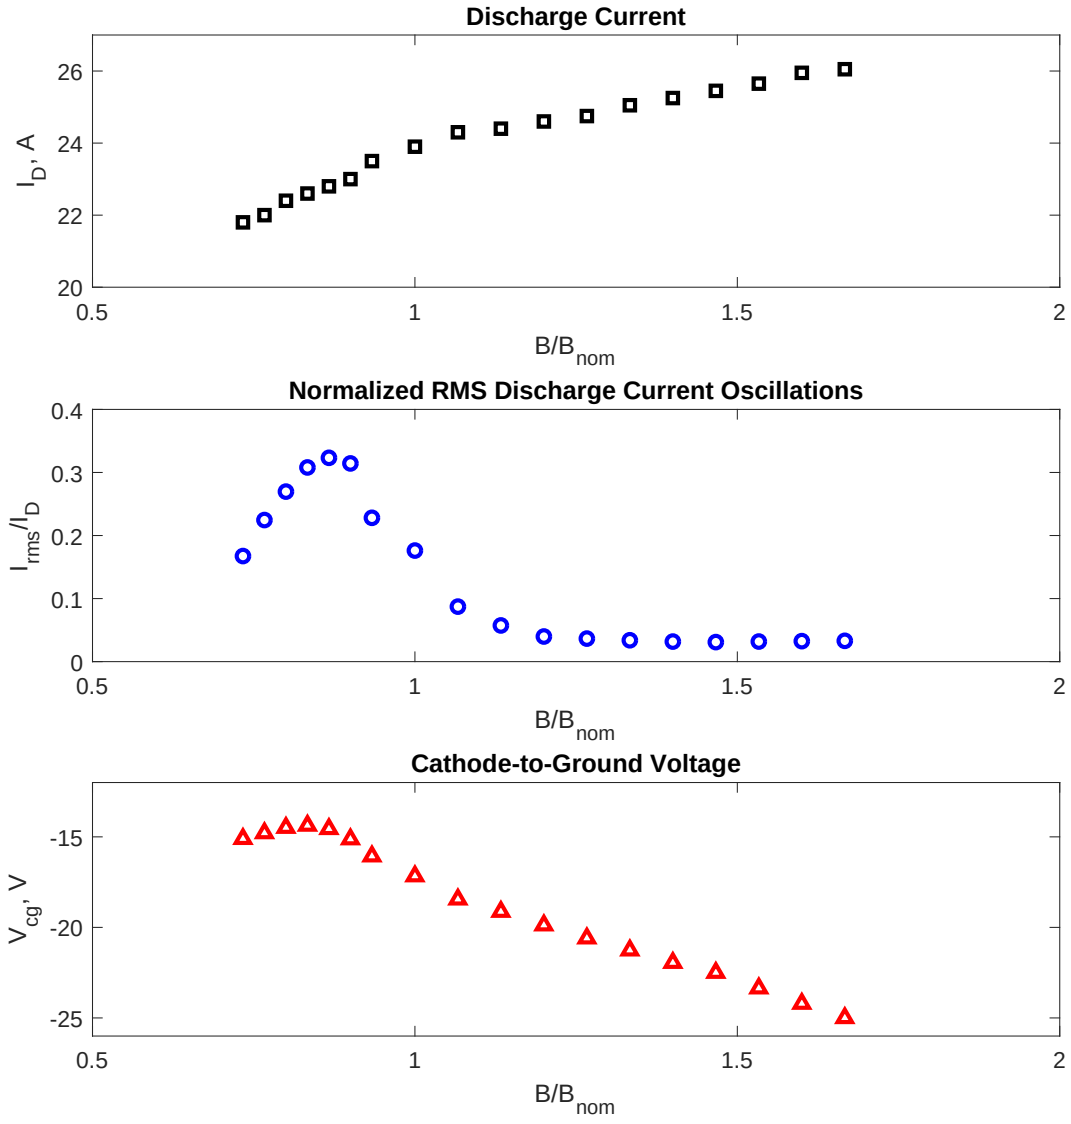

Figure 5.7: The results of a magnetic field optimization map during low-power performance characterization 2 for the Middle channel at 400 V, 24 A. Discharge current, normalized RMS discharge current oscillations, and cathode-to-ground voltage are all plotted against the magnetic field strength (which has been normalized by thruster nominal magnetic field strength).

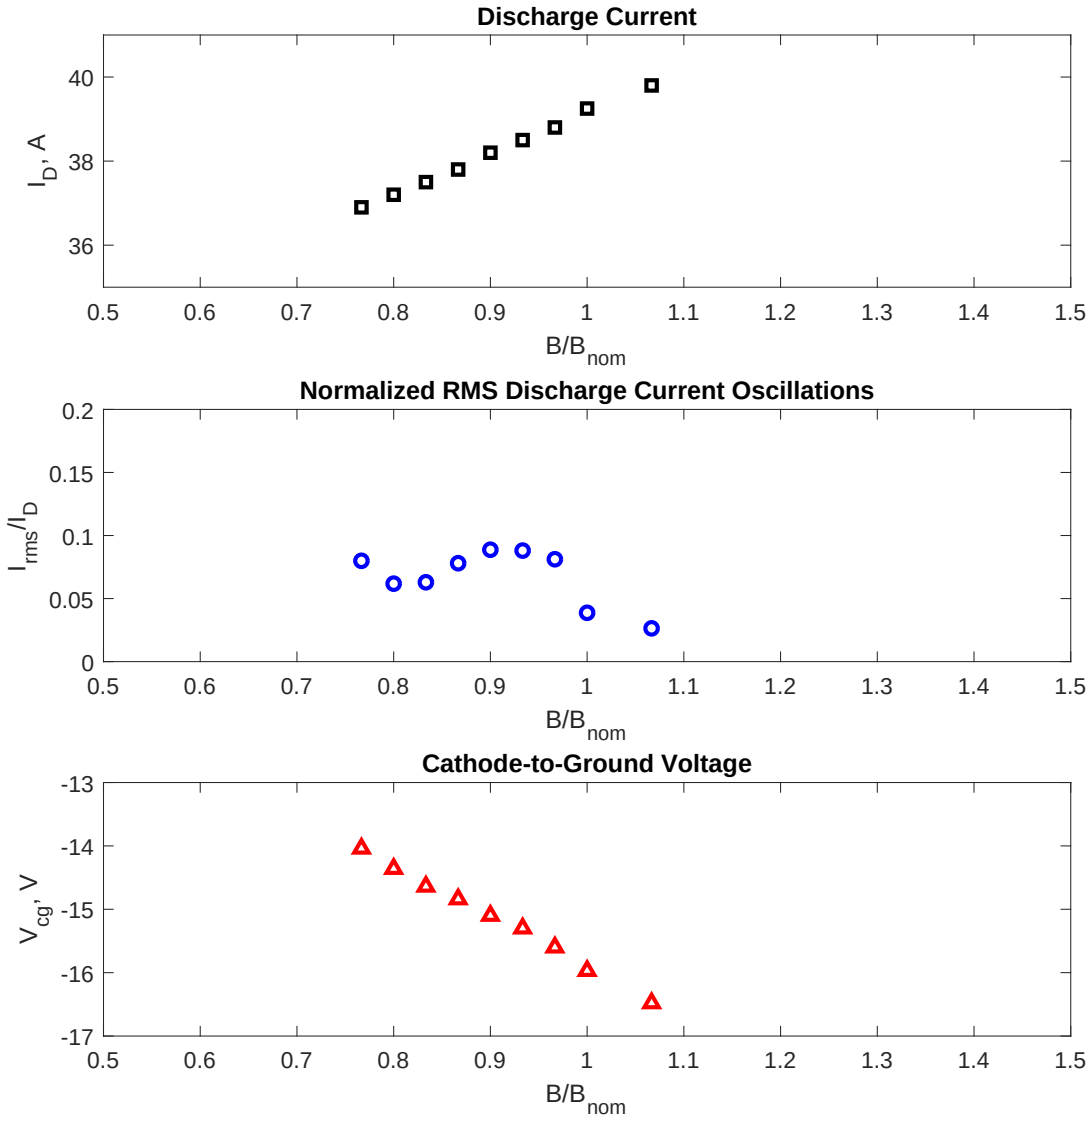

Figure 5.8: The results of a magnetic field optimization map during low-power performance characterization 2 for the Outer channel at 300 V, 40 A. Discharge current, normalized RMS discharge current oscillations, and cathode-to-ground voltage are all plotted against the magnetic field strength (which has been normalized by thruster nominal magnetic field strength).

decreasing magnetic field strength. Just as with the Inner and Middle data, there is no rapid, large increase in discharge current or oscillations as the magnetic field strength decreases. In a general sense, the three channels display similar characteristics.

In addition to the single-channel maps, a small map was undertaken with the thruster firing in the IM condition. This map consists of two components. The first, presented in Figure 5.9, swept the Inner channel's magnetic field while that of the Middle channel was held constant at B/B_{nom}=1.07. The second, shown in Figure 5.10, swept the Middle channel's magnetic field while that of the Inner channel was held constant at B/B_{nom}=1.07.

The maps are only over a small range of magnetic field strengths, and any conclusions drawn cannot be taken for the thruster as a whole or any operating conditions other than this particular one. However, these plots demonstrate that, at least for the IM condition at 300 V and low current density, and over a relatively small range of magnetic field strengths (and magnetic field strength differences), the behavior of the channels is isolated. That is, changes to a given channel's field strength affects its discharge current behavior but not those of the other channel. Here, increasing the Inner channel's magnetic field strength while holding the Middle channel's constant causes the Inner channel's discharge current to rise approximately 0.6 A over the field strengths mapped and causes its oscillation magnitude to drop from approximately 75% of the mean discharge current to below 50%. Similar trends can be seen with the Middle channel when its field strength is swept: its discharge properties are affected, while those of the Inner channel stay roughly constant.

It is also worth noting is that both maps demonstrate an effect on cathode-to-ground voltage, yet that trend is opposite between the two channels. Lowering the Inner channel's field strength causes a decrease in cathode-to-ground voltage, as expected and as demonstrated by all three single-channel maps. Yet decreasing the Middle channel's field strength actually causes cathode-to-ground to become more negative, suggesting worsened cathode coupling (though without plasma diagnostics, this cannot be confirmed). Reasons for this are unclear from these limited data, and further exploration is recommended.

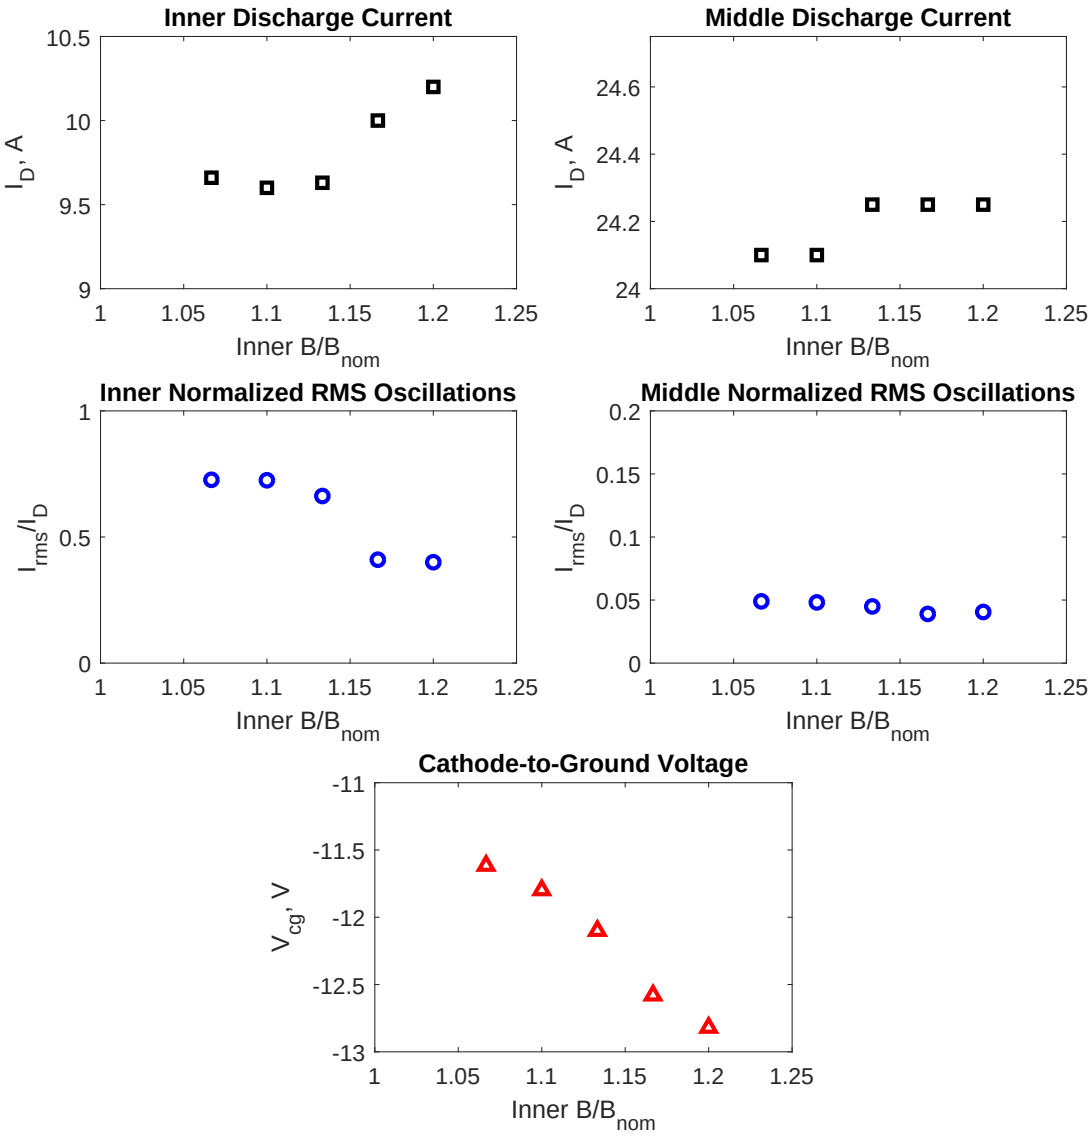

Figure 5.9: The results of a magnetic field optimization map during low-power performance characterization 2 for the thruster firing in the IM configuration at 300 V, 10.4 kW (9.8 A Inner discharge current, 24 A Middle). Here, the Inner channel's magnetic field strength was swept while the Middle's was held at B/B$_{nom}$=1.07. Discharge current and normalized RMS discharge current oscillations for each thruster, along with cathode-to-ground voltage, are all plotted against the Inner-channel magnetic field strength (which has been normalized by thruster nominal magnetic field strength).

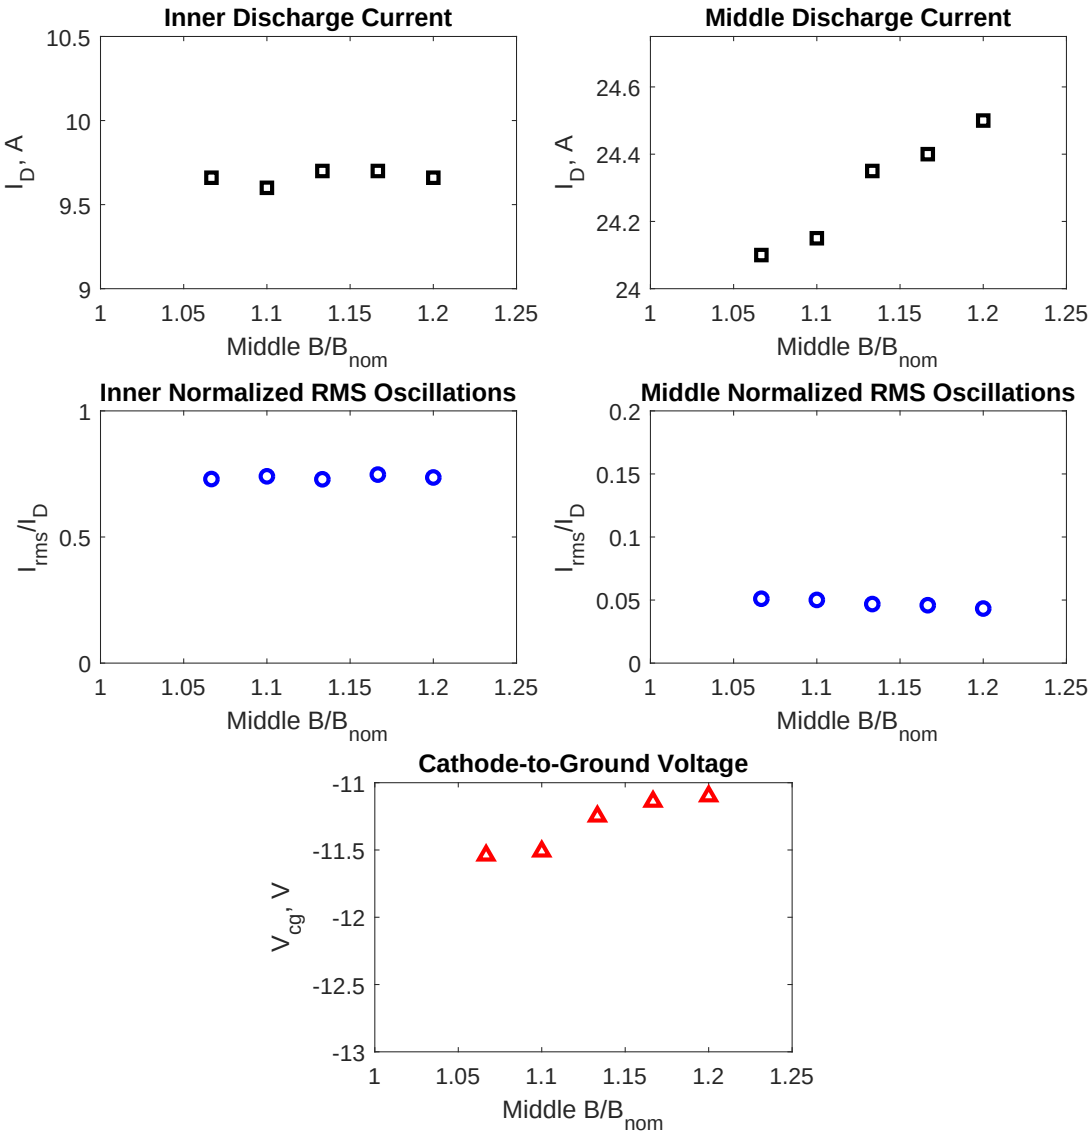

Figure 5.10: The results of a magnetic field optimization map during low-power performance characterization 2 for the thruster firing in the IM configuration at 300 V, 10.4 kW (9.8 A Inner discharge current, 24 A Middle). Here, the Middle channel's magnetic field strength was swept while the Inner's was held at B/B$_{nom}$=1.07. Discharge current and normalized RMS discharge current oscillations for each thruster, along with cathode-to-ground voltage, are all plotted against the Middle-channel magnetic field strength (which has been normalized by thruster nominal magnetic field strength).

Modeling work done by Cusson after this test demonstrated that the optimized coil current ratios for multi-channel operation are only valid for equal field strengths between channels [205]. Operating channels at different field strengths will skew one or both field shapes due to a change in the magnetic flux being shared between the channels. For the limited range of field strengths tested here we suspect that the skewing was minimized, but maps covering a larger range of field strength difference between the channels will need to compensate for this skewing. This becomes a very complex problem with all six electromagnets energized.

5.3.3 Results: Thermal Steady-State

The thruster was operated to limited thermal steady-state (as defined above) in four configurations: I, M, O and IM. Results from these runs are presented as temporal thermocouple traces. For each channel, the temperatures of the relevant magnets, channel cup(s), and back pole location(s) are plotted, except in the cases where the relevant thermocouple was providing bad data. Those cases were the thermocouple on Magnet 1 (which was shown to have a lower-than-acceptable isolation against the windings of the electromagnet and was thus disconnected for this test) and the thermocouple located on the Middle channel cup, which was providing temperatures near ambient throughout all thruster runs and was confirmed to have come loose from its mounting location in the subsequent thruster disassembly, discussed below in Chapter 6.

Thermocouple traces for the Inner channel thermal steady state run are shown in Figure 5.11. It can be seen that the channel cup reaches thermal equilibrium well before the back pole does, indicating that the method of using the back pole temperature change to indicate thermal steady-state of the whole thruster was valid. Though the plateau is not as striking as that seen in the channel cup trace, the back pole does flatten out to be within the 10 ^{o}C/hour criterion set for the test before thruster shut-down.

Thermocouple traces for the Middle channel thermal steady state run are presented in

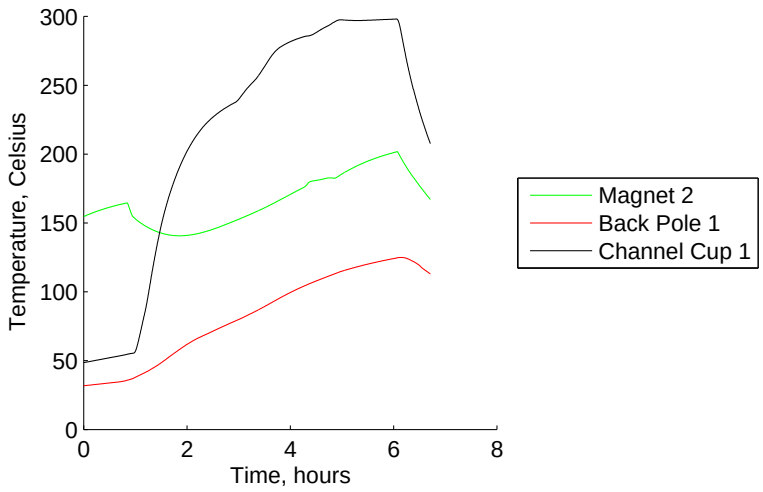

Figure 5.11: Relevant temperatures from the Inner-channel thermal steady state run. The sudden drop in temperature seen in all traces just past t=6 hours coincides with thruster shut-off.

Figure 5.12. Here, as noted above, the channel cup thermocouple was providing inaccurate readings and is thus not plotted, but it can be seen that both magnets arrive at an obvious plateau well before thruster shut-down, and once again the back pole temperature arrives within its criterion before shut-down.

We see that Magnet 3 operated approximately 100 °C hotter than Magnet 4. This is due to the coil ratio used for the Middle channel, which puts significantly more power into Magnet 3 than Magnet 4. In general, Magnets 1-6 are all operated at comparable coil currents to provide equivalent field strengths (with minor variation due to variation in exact magnetic circuit between the channels). This explains the fact that Magnets 3 and 4 run hotter than Magnet 2: they are operating at similar currents, but the larger magnets have higher resistances and are thus dissipating more power.

Similar traces for the Outer channel are presented in Figure 5.13. The channel cup once again arrives into thermal steady-state almost two hours prior to shut down. The large

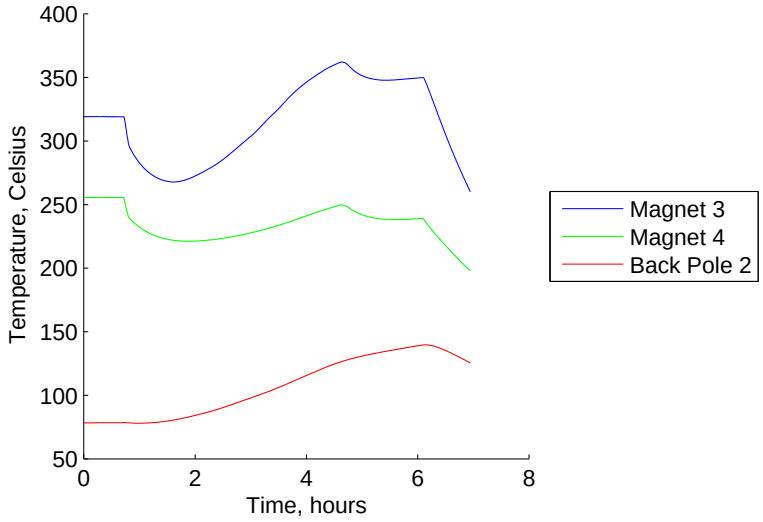

Figure 5.12: Relevant temperatures from the Middle-channel thermal steady state run. The sudden drop in temperature seen in all traces just past t=6 hours coincides with thruster shut-off.

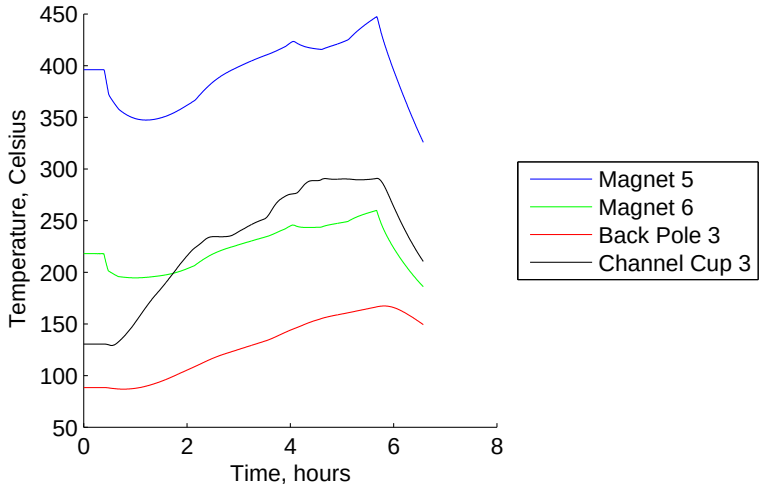

Figure 5.13: Relevant temperatures from the Outer-channel thermal steady state run. The sudden drop in temperature seen in all traces just prior to t=6 hours coincides with thruster shut-off.

spike in the magnet temperatures, especially visible in the trace for Magnet 5, is due to an increase in magnetic field strength near the very end of the test in an attempt to settle thruster instabilities. It can be seen that these changes in the magnet temperatures do not affect the channel cup or back pole traces in a detectable way.

It is interesting to note that Magnet 5 operates at a significantly higher temperature (nearly double) than Magnet 6. We observed this behavior throughout testing. This is somewhat surprising considering that Magnet 6 is physically larger and has a higher resistance, meaning it consumes more power at a given coil current. However, Magnet 6 is the outer-most coil of the thruster and features ample surface area that radiates directly to the vacuum chamber. Thus, it is able to reject much more of its heat than Magnet 5.

Finally, the set of thermocouple traces for the two-channel IM condition are shown in Figure 5.14. This run was performed as such: first, the Middle channel was lit and allowed to operate for approximately 3 hours until it reached steady operation and had done initial thermal soaking. This was done out of concern regarding whether the LVTF cyropumps would be able to operate steadily with the increased gas load of the IM condition for long enough to reach thermal steady-state; by operating with only the Middle channel, the thruster was allowed to heat without the extra propellant flow of the Inner channel. Once the Inner channel was lit around t=4 hours, large changes in the slopes of Channel Cup 1 and Back Pole 1 are noticeable, as expected, along with a smaller but detectable increase in the slope of Back Pole 2. Once the thermal steady-state criterion was met for both back pole thermocouples, the Inner channel was shut off first, which accounts for the drop in Channel Cup 1 seen at approximately t=7.5 hours. Then, the Middle channel was shut off, providing the large drop off just before t=8 hours similar to those seen in other traces.

5.3.4 Asymmetric Anode Heating Issue

During the second low-power test, asymmetric heating of the Outer channel anode was detected. This presented itself as a “hot spot” or “glow region” on the anode, wherein a

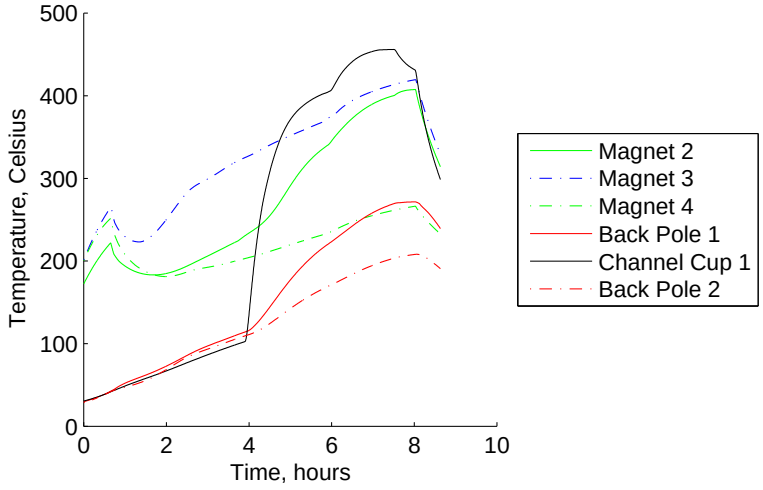

Figure 5.14: Relevant temperatures from the IM thermal steady state run. Temperatures for the Inner channel are plotted as solid lines, and counterparts from the Middle channel are plotted as dot-dashed lines in the same color. The sudden large change in slope of Channel Cup 1 at approximately t=4 hours coincides with the lighting of the Inner channel. The sudden drop in temperature seen in all traces just prior to t=8 hours coincides with Middle channel shut-off, and the drop seen in the Channel Cup 1 trace seen prior to that at approximately t=7.5 hours coincides with Inner channel shut off. Note that the plot begins approximately 1 hour before Middle channel light.

portion of the anode was visibly glowing orange after thruster shut-down while the rest of the anode was not. This issue appeared during thermal steady-state testing. It was first detected in the fourth hour of Outer channel firing at 16 kW (400 V, 41 A). By viewing the thruster through the endcap viewport of the LVTF, an orange glow was detected emanating from the outer channel in the 10-11 o'clock region.

Reviewing information from the first firing of the thruster showed that a similar hot spot appeared in the same location during the krypton burn-in portion of the thruster firing. The spot was noted in the test log at 500 V, 55 A, but testing quickly proceeded to low-voltage xenon performance characterization and the phenomenon was not investigated further. We undertook a much more substantial investigation of the phenomenon here. We identified a number of potential causes, including asymmetric propellant distribution, asymmetric magnetic field topology, uneven thermal contact between the anode and channel cup or channel cup and thruster back pole, and asymmetry in anode height (thus pushing the anode further into the plasma in some regions). As described in the next chapter, the first portion of the investigation occurred during the thruster test window and included both thruster inspection at atmosphere and operational checkouts. When no obvious cause was identified by the end of the test window, we undertook a second phase of investigation to more thoroughly characterize the magnetic field and propellant distribution in the thruster. Finally, a series of mechanical improvements were made to the thruster at GRC following a detailed inspection of the thruster. These included repairing welds on the anodes and adding gasket material to provide a better seal between the anode and channel cup.

5.3.5 Conclusions

A substantial amount of information was gathered during low-power thruster characterization 2. All of this information was important for preparation for the subsequent high-power performance mapping. Stable operation was achieved across a wide range of thruster operating conditions. We demonstrated that the magnetic field topology (as set by the coil ratios)

and strength (as set by the magnet currents) have a substantial effect on the performance of the X3. By only operating the coils for the channels firing, and by updating the coil ratios for various configurations to better provide proper magnetic lens topology, significant performance gains were obtained, including most notably an efficiency increase of 0.33 and a thrust increase of 67% at the same discharge conditions for the Outer channel. Finally, the thruster was operated to limited thermal steady-state conditions, demonstrating reasonable thermal behavior and temperatures. One concern that was identified was that the anode of the Outer channel appeared to be heating asymmetrically during operation, resulting in a portion of the anode visibly glowing after thruster shut down.

5.4 Summary

We conducted an important set of tests at the low power levels attainable in the LVTF in Configurations 1 and 2. These tests explored the performance and stability of the X3 and helped prepare for subsequent high-power performance mapping. Low-power characterization 1 identified a number of concerns with X3 operation, including performance well below that predicted by theory and results from other state of the art Hall thrusters. However, that test provided no insight into the mechanisms causing these results. Subsequent modeling identified improvements to be made to the thruster magnetic field topology by modifying the coil ratio. Using these improvements, low-power thruster characterization 2 demonstrated improved performance at similar operating conditions. Specifically, the Outer channel demonstrated an improvement in anode efficiency of +0.33 and an increase of 67% in the thrust produced at similar operating conditions with the new magnetic field settings. This test also characterized magnetic field strength optimization of the thruster as well as its thermal profile at powers up to 16 kW in different channel configurations. All these data provided important insight into thruster operation prior to high-power testing at GRC.

CHAPTER 6

Thruster Checkout

"The development of high-power Hall thrusters as well as of any large-scale device represents certain design and technical challenges... [A] linear increase in Hall thruster dimensions involves certain difficulties with providing an azimuthally homogenous magnetic field and gas distribution in the thruster channel required for normal drift motion of electrons."

– Robert Jankovsky, Sergey Tverdokhlebov, and David Manzella [207]

6.1 Introduction

Before high-power testing at GRC we thoroughly inspected the X3 for any evidence of issues that would prevent high-power operation. As we describe in Chapter 5, the Outer channel anode appeared to be asymmetrically heating during low-power testing. This was concerning as it indicated some kind of defect in the thruster. We undertook a full investigation, starting with a series of tests during thruster operation as described in Section 6.2. When these did not identify a cause for the heating, we conducted thorough magnetic field and propellant flow uniformity maps as described in Section 6.3. We found no major defects in either test, which was a positive result for high-power testing but still did not provide a cause for the anode heating issue. Finally, we shipped the thruster to GRC in preparation for testing and performed a thorough inspection of the thruster. This inspection, described in Section 6.4,

yielded a number of important improvements to be made to the thruster, including welds requiring repair on the anodes. Ultimately, no clear cause of the anode heating issue was identified, but a number of critical thruster improvements were made that incorporate GRC best practices.

6.2 Investigation During Thruster Testing

Immediately after identifying the hot spot, we vented the facility to atmosphere and thoroughly inspected the thruster while mounted on the thrust stand. No obvious causes were identified. A limited map of the magnetic field was undertaken in situ, which demonstrated that the thruster magnetic field on channel centerline at the channel exit plane did not vary by more than 2.2%. An attempt was made at performing a propellant distribution cold-flow study in situ, but this was technically challenging and did not yield useful results. Without a cause identified, we returned to thruster testing in an attempt to identify how the hot spot changed with changing thruster operating conditions. These tests were all qualitative, as no thermal camera was available to us during the test window and placing thermocouples on the high-voltage anode was deemed risky and challenging. Photographs were taken of the thruster immediately after shut down through the end cap of the LVTF at similar camera settings. By comparing the location and apparent brightness of the anode glow, certain trends could be identified. The results indicated minimal dependence on magnetic field strength but that the glow region appeared to diminish with higher anode flow rates. Most interestingly, the region moved location (mirrored over the horizontal thruster centerline) when the magnetic field direction of the channel was reversed. Despite identifying these trends, we were unable to identify a clear cause for the glow region.

6.3 Investigation Post-Testing

With no obvious cause of the glow region identified at the close of the test window, we undertook a more thorough and rigorous investigation of the thruster hardware. This consisted of mapping the magnetic field of the thruster in four azimuthal locations, as well as studying saturation; disassembling the thruster to inspect it in detail and to ensure there were no damaged or broken pieces; and finally a more thorough and rigorous mapping of the propellant distribution for each channel.

6.3.1 Magnetic Field Mapping

6.3.1.1 Introduction and Scope

Magnetic field azimuthal uniformity is critical for efficient Hall thruster operation. We performed magnetic field mapping to verify uniformity and thruster saturation. For the Outer channel of the thruster, this was also used as part of the glow region investigation.

We made the following measurements:

1. Single-point saturation data for each channel at one azimuthal location
2. Two-dimensional channel maps of single-channel conditions (I, M, O) at four azimuthal locations
3. Two-dimensional channel maps of two-channel (IM) and three-channel (IMO) conditions at one azimuthal location

The two-dimensional maps covered the channel cross section available with the BN installed, as well as a small amount of area downstream of channel exit that included area in front of the thruster pole pieces. A notional grid is shown in Figure 6.1, where the anode face is located at X/L = 0 and the channel extends into the positive X/L direction. The two-dimensional maps were taken at 12, 3, 6, and 9 o'clock on the thruster face such that

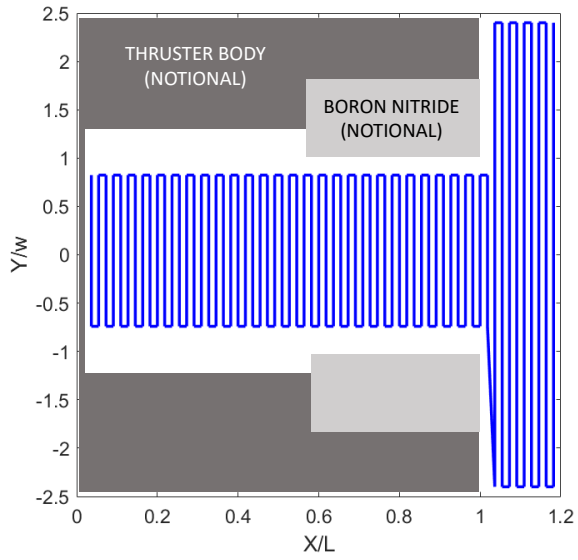

Figure 6.1: A notional magnetic field sweep grid. Here, the X coordinate has been normalized by the channel length L, and the Y coordinate by the channel width w. Notional thruster channel geometry is shown. Similar grids were created for each of the X3's channels based on the given channel's width.

only two motion stages were necessary to accomplish the measurements. A single grid was created for each channel and used for all maps of that channel.

6.3.1.2 Test Procedures

The first step of testing involved the construction of the rig and alignment of the motion stages themselves. Although the stages were designed to interface with an optical table and with each other in a way that allowed for easy alignment, there was still enough play in their interfaces to require alignment verification. The X stage was aligned using a tape measure to ensure it was square to the optical table. The Y stage was then aligned physically using optical posts placed at either end of the stage: the X stage was advanced (with the Y stage mounted on top, but loose and able to pivot) until each end of the Y stage was touching its optical post. The Z stage was then aligned on top of the Y using a measuring tape and a digital protractor.

With the stages aligned to the optical table, the next step of alignment involved aligning the thruster to the stages. The following procedure was used, which was replicated from that used at JPL. It utilizes a small Thor Labs laser and two irises, all mounted to optical posts, as well as a 1-inch diameter mirror.

1. We placed the laser and both irises along the same row of holes on the optical table. For the X3, we chose a row such that the laser beam would be hitting the thruster on the BN of one of the channels.

2. Using a reference length or set of calipers, we set the irises to be at the same height. For the X3, we chose a height as close to thruster centerline as the optical equipment allowed.

3. We turned the laser on and adjusted its height and angle until the beam passed through both irises and reached the thruster face.

4. We moved one of the irises to a different location along the same optical table row, and ensured that the laser beam still passed through both irises and reached the thruster face. This step ensured that the laser is truly aligned with the optical table—because the irises have a finite area, it is possible to have the laser aligned to pass through both irises in a given location without actually having the laser be square to the table.

5. With the laser aligned and passing through both irises, we held the 1-inch mirror against the thruster face to reflect the beam back toward the irises and laser.

6. We then adjusted the thruster alignment (in two axes: side-to-side and forward-back tilt) until the reflected beam passed through both irises and entered the laser face.

As a final alignment check once this was complete, the Hall probe was brought forward until it touched the BN at each of the channel locations across both axes of the thruster (the 12-6 o'clock axis and the 3-9 o'clock axis). With the X3, we found that these values varied by a few tenths of a millimeter, but the variation was random (i.e., there was no obvious

slope across a particular axis). This was attributed to the fact that the BN in the X3 is many separate parts held in place by different pole pieces and springs.

With the entire system aligned, data collection began. Data collection consisted of two phases: locating of (X,Y) = (0,0) for the given channel and the actual mapping. Location of (0,0) used the following technique:

1. Location of X=0 was done by approaching the probe tip in 0.1 mm increments until it was touching the BN of the given channel. A piece of paper was used to protect the BN, and when the probe was touching the BN firmly enough to pin the paper, we declared it to be at X=0.

2. Location of Y=0 then proceeded, which involved advancing the probe 8 mm into the channel and then approaching the inner and outer BN edges in turn. Here, because of the flexibility of the probe in the long axis, the paper technique was not used, and instead a flashlight beam was shined behind the probe tip. When the light disappeared, the probe was declared to be touching the wall. Both walls were located in this manner, and the point halfway between each wall's location was taken to be Y=0.

The LabView virtual instrument collected data by stepping through a grid of points supplied by the operator. The Lakeshore 3-axis Hall probes feature Y and Z sensors that are offset 2.08 mm from the centerline of the probe. To accommodate for this, the code automatically moved the probe 2.08 mm in the correct direction to gather the Y data. The gridding process thus was the following:

1. Move to new position

2. Record X data

3. Move +2.08 mm Y

4. Record Y data

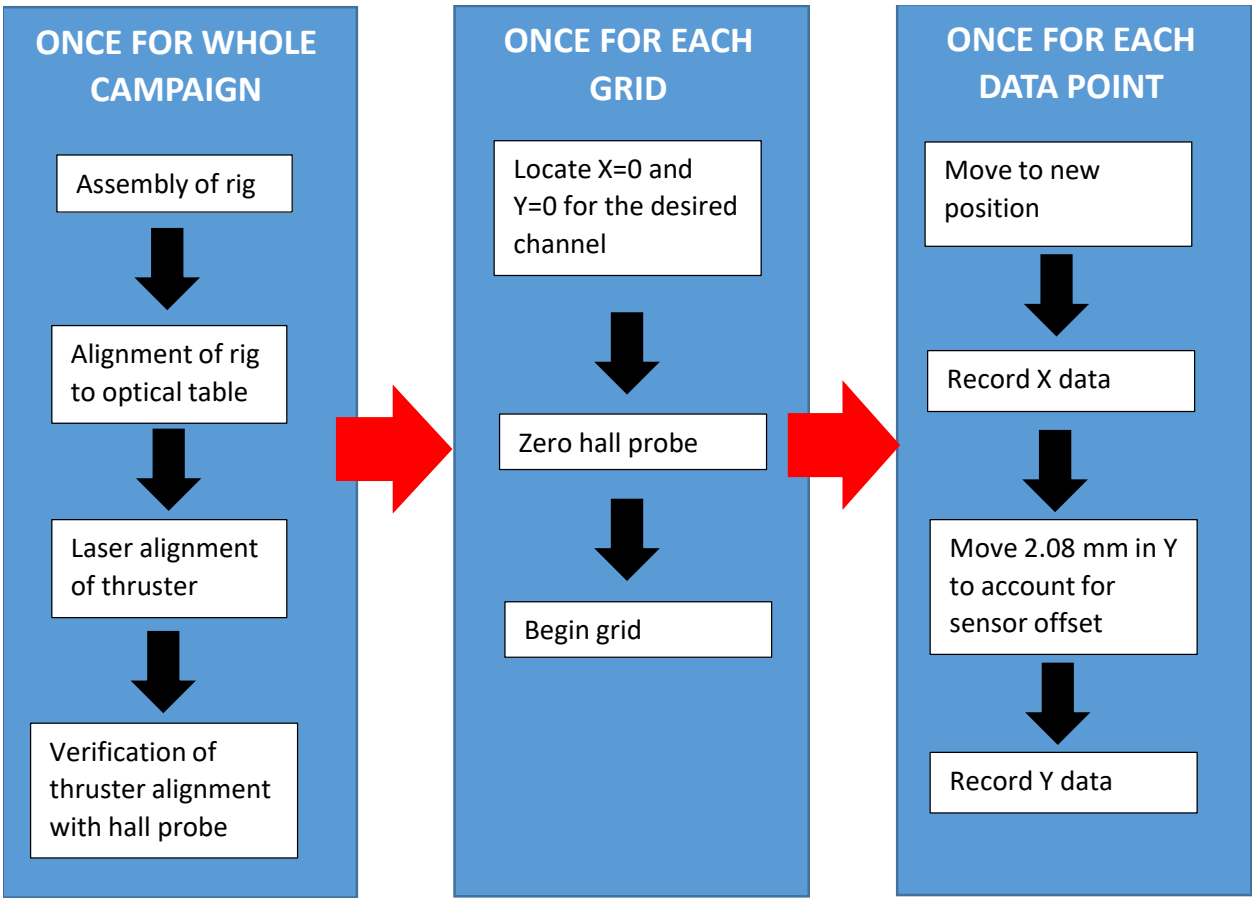

Figure 6.2: A flow chart detailing the magnetic field mapping procedure.

5. Move to new position

Grids thus had to be designed around this adjustment to ensure that the probe would not touch the BN walls. Figure 6.2 presents a flow chart that details the basic process followed for the magnetic field mapping of the X3 presented here.

6.3.1.3 Test Setup Validation

Before proceeding with maps of the X3, a validation campaign was undertaken with the H9 thruster [208, 209], which had recently been mapped at JPL using a similar rig, Hall probe, and gaussmeter. The intention of this test was to validate both the UM test setup and the alignment procedures. Results from this test are omitted here due to the proprietary nature of the H9's magnetic field design. Ultimately, the results were able to match relatively well

in shape, but the UM setup was not able to replicate exactly the results taken at JPL. Centerline field strengths matched between the two facilities, but the streamtraces of the data collected on the UM test apparatus had a slight tilt as compared to the JPL results that was never able to be removed. After many re-alignment attempts, the Hall probe itself was settled on as a likely culprit for the mismatch. Due to time constraints, X3 mapping proceeded with the old probe. Thus, this H9 validation can be seen as a measure of the error associated with the measurements presented below and provide an estimate for how closely the measurements made with the test apparatus in this configuration can be expected to match modeling results.

6.3.1.4 Saturation Results

Magnetic circuit saturation occurs when the current–field strength characteristic for the channel is no longer linear. Eventually, the characteristic flattens such that increased current provides no increase in magnetic field strength. Knowing the point at which the characteristic stops being linear is important for thruster operation, as it provides an upper limit for magnetic field strength. Because the curve flattening is gradual, thrusters can be operated in the non-linear region if the curve is known and the field topology is maintained. The X3 magnetic circuit was designed to provide saturation for each channel at a particular magnetic field strength, referred to here simply as $B_{sat,design}$. All plots below are normalized by this value to provide qualitative information about the magnetic circuit saturation behavior without providing specific proprietary information.

As a comparison tool, a linear least-squares fit was made to the linear region of the Hall probe data for each channel (here, the first nine or ten data points to ensure linearity) and then extended through the full range of coil currents investigated. These linear fits provide an un-saturated reference to which the Hall probe data can be compared. When the circuit begins to saturate, the Hall probe data points are expected to fall beneath the linear fit line. Hall probe data was collected for each channel at the point of maximum field strength on

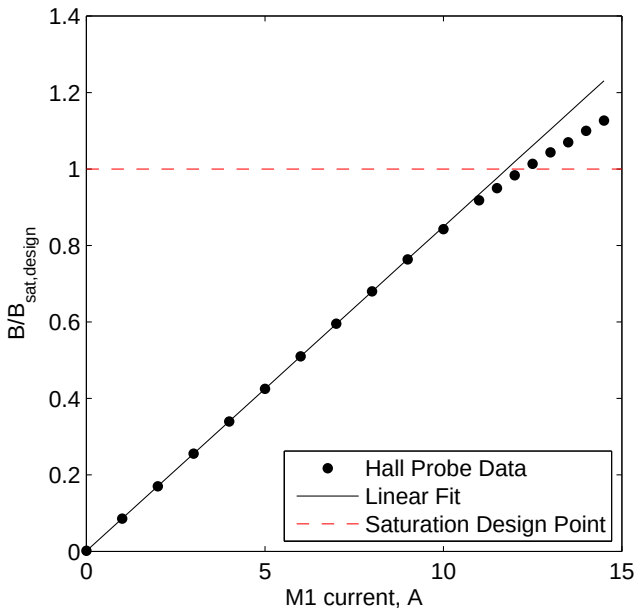

Figure 6.3: Inner channel saturation results. The Hall probe data is plotted against a linear least-squares fit of the first ten data points of the Hall probe data to provide an "un-saturated" reference line. The magnetic field strength is normalized by the designed saturation point. The data are plotted against M1 current on the abscissa for simplicity, but M1 and M2 were both powered for these data at a constant ratio. Data were collected at the point of maximum field strength on channel centerline at the 12 o'clock position.

channel centerline. This location varied slightly for each channel but was generally at the channel exit plane or slightly downstream of that. These data were collected after a grid was performed for the given channel, and as such, the probe alignment (as performed by the techniques described above) carried over from the grid.

Inner channel saturation data is presented in Figure 6.3. It can be seen that the Hall probe data begins to deviate from the linear fit at a $B/B_{sat,design}$ value of approximately 0.9, with the Hall probe data clearly deviating at $B/B_{sat,design}$=1. However, even up to measured magnetic field strengths of $B/B_{sat,design}$=1.13 (the maximum attainable field strength with the available power supplies at UM), the curve has not fully flattened. Middle channel saturation data is presented in Figure 6.4. Here it can be seen that the Hall probe data begins to deviate just slightly from the linear fit at values of $B/B_{sat,design}$ very close to 1. For the Middle channel, values of $B/B_{sat,design}$=1.10 were attainable with the power supplies

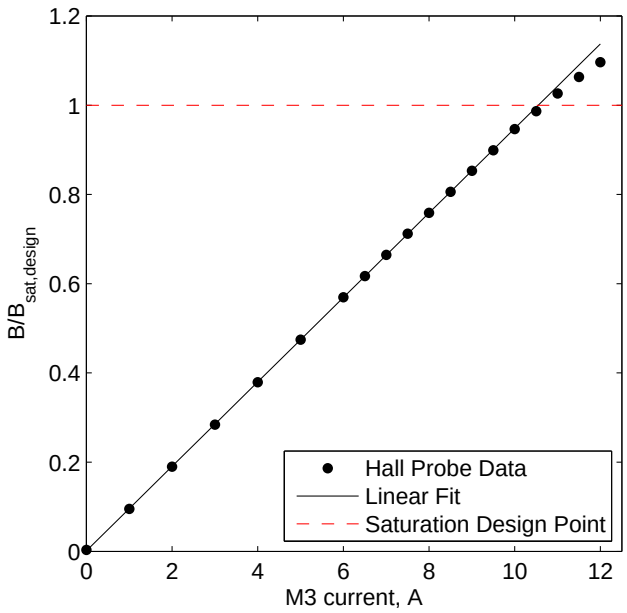

Figure 6.4: Middle channel saturation results. The Hall probe data is plotted against a linear least-squares fit of the first nine data points of the Hall probe data to provide an "un-saturated" reference line. The magnetic field strength is normalized by the designed saturation point. The data are plotted against M3 current on the abscissa for simplicity, but M3 and M4 were both powered for these data at a constant ratio. Data were collected at the point of maximum field strength on channel centerline at the 12 o'clock position.

available, and although deviation is clear at that field strength, the curve has not leveled off. Outer channel saturation data is presented in Figure 6.5. Here it can be seen that the curve has just started to deviate from linear at the maximum attainable field strength of $B/B_{sat,design}$=1. Overall, the data for all three channels shows that they are performing at or near the design limit. The curves display very linear trends until that point, indicating proper magnetic circuit response.

In addition to the single-channel saturation studies, a brief multi-channel verification was performed. Theory suggests that a given channel's magnetic circuit should saturate at a particular field strength regardless of whether any other channel's coils are powered. This is due to the fact that saturation is an effect of the magnetic flux and the geometry of the magnetic circuit components. The coil current at which this occurs will likely vary with a neighboring channel's coils powered, as magnetic flux from the extra coils will be present in

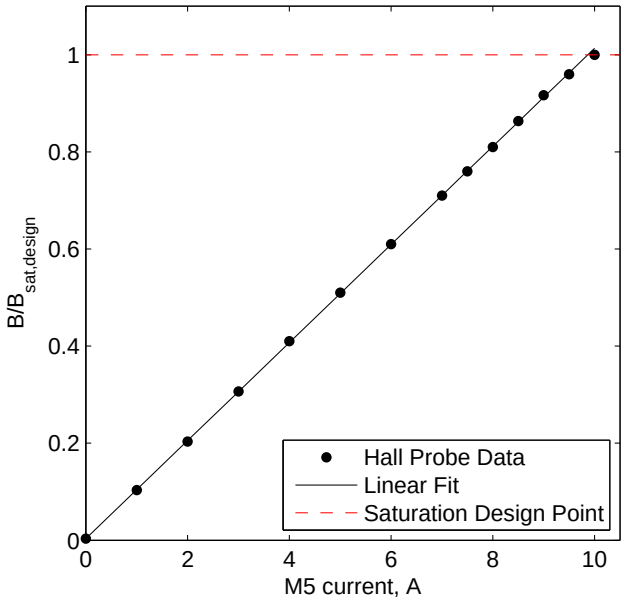

Figure 6.5: Outer channel saturation results. The Hall probe data is plotted against a linear least-squares fit of the first ten data points of the Hall probe data to provide an "un-saturated" reference line. The magnetic field strength is normalized by the designed saturation point. The data are plotted against M5 current on the abscissa for simplicity, but M5 and M6 were both powered for these data at a constant ratio. Data were collected at the point of maximum field strength on channel centerline at the 12 o'clock position.

the circuit as well.

To verify that the X3 was behaving in this manner, additional saturation studies were performed on the inner and middle channels in the IM coil configuration. As detailed by Cusson [205], the coil ratios for the channels varied between single-channel (I, M, O) conditions and the multi-channel (IM, IMO, etc.) conditions. For each channel, only that channel's coil currents were swept (holding to the prescribed ratio), while the other channel was held at a constant, moderate coil current. For each channel, the Hall probe was placed in the same location as for the single-channel studies.

The results of the Inner channel sweep in the IM condition are plotted in Figure 6.6. A linear least-squares fit is plotted alongside that Hall probe data. For reference, the Hall probe data and linear fit for the single-channel case (the same data as in Figure 6.3) are reproduced here. It can be seen that the saturation behavior is essentially identical: around $B/B_{sat,design}$=0.9, slight deviation from linear is detected, and at $B/B_{sat,design}$=1, there is clear deviation. However, even at the peak $B/B_{sat,design}$ of 1.11, the Hall probe data has not leveled off.

It can be seen that these field strengths are achieved at much lower coil currents than the single-channel condition, which suggests substantial sharing of magnetic flux from the Middle channel's coils. This is further supported by the fact that the coil ratio for the Inner channel was significantly different between the I and IM cases. However, this change does not ultimately affect the saturation field strength for the Inner channel, just as theory suggested would be the case.

Similar data for the Middle channel in the IM condition are plotted in Figure 6.7. Unlike the Inner channel, the two sets of data fall very close to each other. This suggests dis-similar sharing between the two channels, a conclusion further supported by the very slight change in optimum coil ratio for the middle channel between the M and IM conditions and the results from Cusson's study in Reference [205]. Power supply limitations prevented testing above $B/B_{sat,design}$=1, but up to the maximum field strength of $B/B_{sat,design}$=0.926 the data

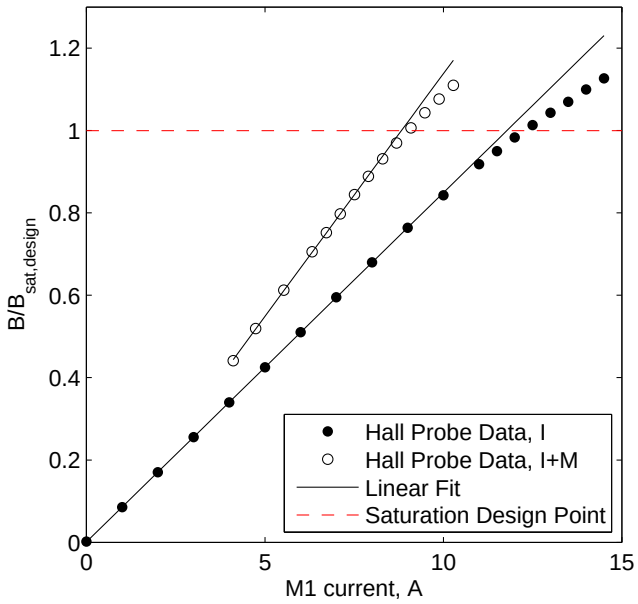

Figure 6.6: Inner channel saturation results for the IM condition. The Hall probe data is plotted against a linear least-squares fit to provide an "un-saturated" reference line. The magnetic field strength is normalized by the designed saturation point. For reference, the Inner channel saturation results for the I condition are reproduced from Figure 6.3. As above, the data are plotted against M1 on the abscissa for simplicity. For this study, M1 and M2 (the Inner channel magnets) were swept at a constant ratio while M3 and M4 (the Middle channel magnets) were held at fixed values. Data were collected at the same point as the single-channel data.

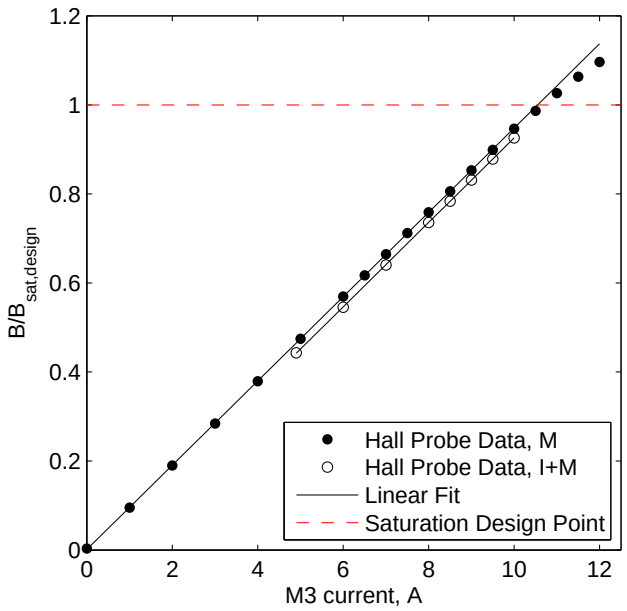

Figure 6.7: Middle channel saturation results for the IM condition. The Hall probe data is plotted against a linear least-squares fit to provide an "un-saturated" reference line. The magnetic field strength is normalized by the designed saturation point. For reference, the Middle channel saturation results for the M condition are reproduced from Figure 6.4. As above, the data are plotted against M3 on the abscissa for simplicity. For this study, M3 and M4 (the Middle channel magnets) were swept at a constant ratio while M1 and M2 (the Inner channel magnets) were held constant. Data were collected at the same point as the single-channel data.

are very linear and track closely with the single-channel M data.

6.3.1.5 Azimuthal Uniformity Results

Two-dimensional in-channel maps were taken at four locations around the thruster: 12, 3, 6, and 9 o'clock on the thruster face. These maps were conducted for each channel in its single-channel configuration. The maps at the four locations were intended to compare field shape at each location to confirm uniformity.

Inner data at each of the four locations is presented in Figure 6.8. Each color represents a different azimuthal clock location. It can be seen in general that the data line up very closely, especially in the region of highest magnetic field near the exit plane of the channel.

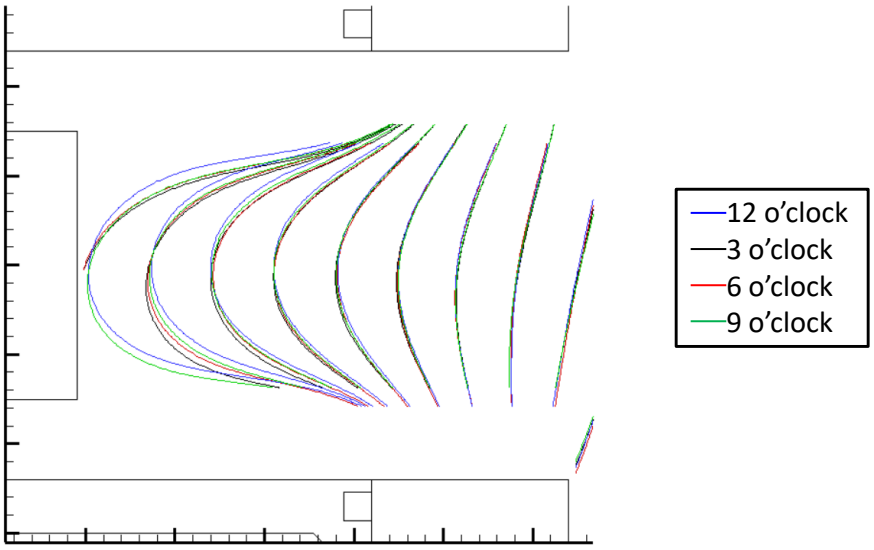

Figure 6.8: Inner channel azimuthal uniformity results for the nominal single-channel condition. Each color represents the Hall probe data from a different clock position on the thruster as noted. Notional channel geometry is included for reference: the channel walls run parallel to the abscissa, and a representation of the anode is visible near the ordinate. Units are deliberately omitted.

Deep in the channel near the anode, the thruster's magnetic field is on the order of Earth's magnetic field, and the slight differences are not surprising. Data in the area of low field strengths were close to the noise floor of the instrument. Middle channel data are presented in Figure 6.9. Here, only three clock locations are presented, as the 9 o'clock map was not collected at the proper magnetic field settings. The same trend is apparent here as in the Inner channel data, that there is very good agreement in the downstream region of the channel and slight variations in field lines deep in the channel near the anode. Outer channel azimuthal uniformity maps are presented in Figure 6.10. The trend of very uniform field lines within the downstream section of the channel is repeated here as well. This is of increased importance for the Outer channel, as a non-uniform magnetic field was identified as a possible culprit for the asymmetric anode heating observed during thruster testing. There is no sign in these maps of any defects that would cause this issue.

We present the variation in centerline field strength, which provides additional insight into azimuthal uniformity, in Table 6.1. The location of these measurements within the channel

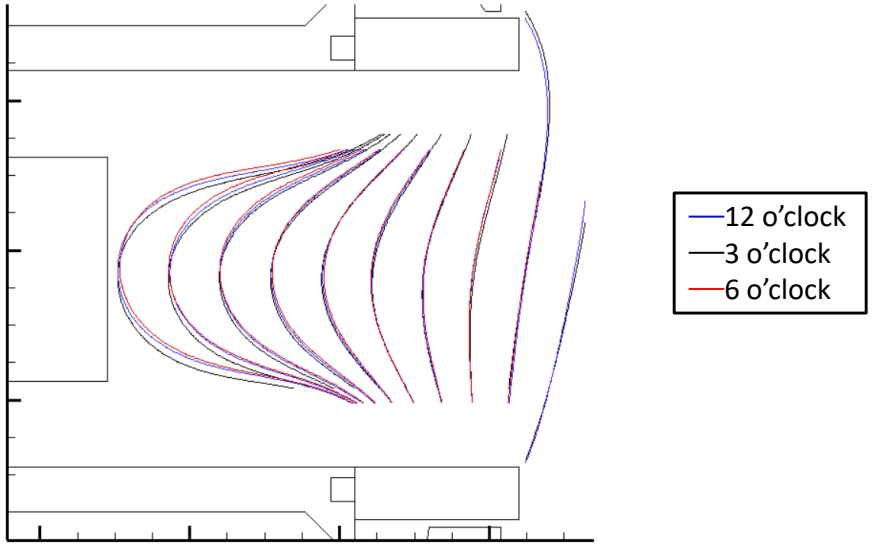

Figure 6.9: Middle channel azimuthal uniformity results for the nominal single-channel condition. Each color represents the Hall probe data from a different clock position on the thruster as noted. Notional channel geometry is included for reference: the channel walls run parallel to the abscissa, and a representation of the anode is visible near the ordinate. Units are deliberately omitted.

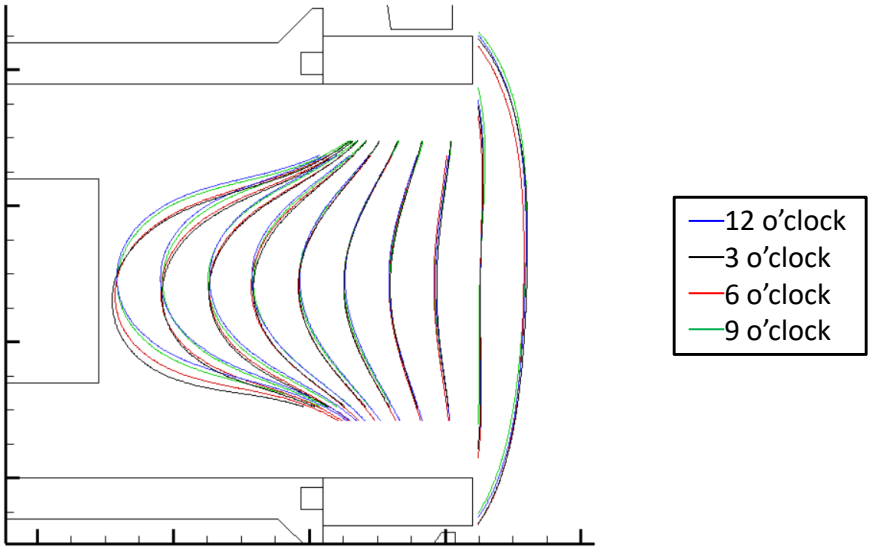

Figure 6.10: Outer channel azimuthal uniformity results for the nominal single-channel condition. Each color represents the Hall probe data from a different clock position on the thruster as noted. Notional channel geometry is included for reference: the channel walls run parallel to the abscissa, and a representation of the anode is visible near the ordinate. Units are deliberately omitted.

Table 6.1: Percent variation between peak channel-centerline field strengths at each of the four locations mapped during the azimuthal uniformity study.

Channel	I	M	O
Variation	0.46%	0.94%	2.3%

cross-section of each channel match the location of the saturation study, as indicated above. The percent variation was calculated by taking the difference between the maximum and minimum field strengths and dividing that by the mean field strength. In general, there are no large variations in field strength across any given channel. Uniformity decreases as channel size increases, which is not surprising, as maintaining concentricity and tolerances is more difficult with increased channel size.

6.3.1.6 Conclusions

The data presented here suggest two significant conclusions. First, the thruster was producing an azimuthally uniform magnetic field on all three channels. No major defects or disagreements between the four azimuthal locations mapped on the thruster were discovered, and peak centerline field strength on each channel varied less than 1% for the Inner and Middle channels and less than 2.5% for the Outer channel. Second, no indication was found that the magnetic field is the cause of the asymmetric heating issue on the Outer channel.

6.3.2 Propellant Uniformity Mapping

6.3.2.1 Introduction and Scope

Propellant uniformity within the discharge channel is critical for efficient Hall thruster operation [72, 202, 207]. Propellant can be distributed using a number of different schemes, most of which couple the propellant distribution function with the discharge circuit into a single component referred to as the anode. As noted in Chapter 2, three anode design types were

considered for the X3. An extensive modeling and simulation campaign was used to investigate all three options thoroughly [4]. Ultimately, a drop-in style design was selected. The modeling work suggested that uniform flow would be achieved in all cases for each anode, but verification of this was necessary.

A limited flow uniformity verification test was conducted prior to first firing of the X3. This test consisted of radial and azimuthal pressure maps at a single flow rate for each channel: a high-flow condition for the Inner and Outer channels and a moderate-flow condition for the Middle channel. Additional low-flow conditions had been planned, but the pressures achieved were below the resolution of the pressure gauge selected. In the interest of time, the low-flow conditions were skipped. Unfortunately, the lowest flow condition is the most likely to produce non-uniform propellant distribution. Because propellant is routed into the anodes of the X3 via two inlets positioned 180^{o} apart, uniform propellant distribution relies on choked flow occurring within the baffles of the anode. At low enough flow rates, the choked condition can be lost, which may cause a significant decrease in propellant uniformity. The Outer channel asymmetric heating issue occurred at flow rates near the lower end of the range of that channel: the region was first discovered at a flow rate of approximately 400 sccm, whereas the pre-assembly flow test only tested that anode at 1500 sccm. Non-uniform propellant distribution was identified as a possible cause of the heating issue.

To both investigate the asymmetric heating and to better verify the propellant distribution of all three anodes across their range of flow rates, a new test was conducted before high-power testing of the X3 using the apparatus described in Chapter 4. This test was intended mostly to verify azimuthal uniformity at the low-flow condition for each anode. As a verification of previous results, high-flow data was taken as well. These tests consisted of placing a single anode of the X3 at a time into the test setup, which for expediency was housed inside the LVTF at UM off to one side of the facility, allowing for these tests to take place alongside ongoing thruster testing for a different project. To limit setup complexity and more easily allow the test to occur alongside the on-going testing, a single radial/axial

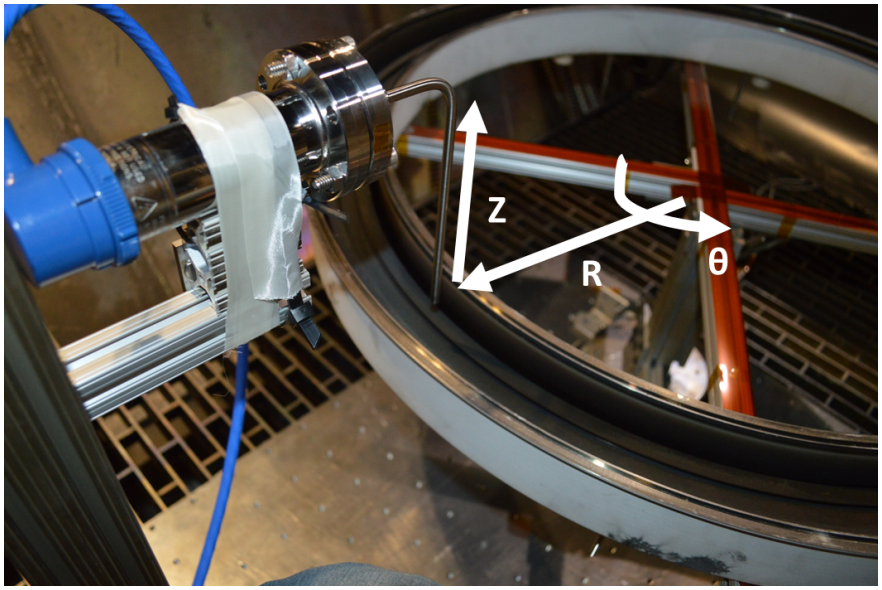

Figure 6.11: The coordinate system used throughout the test campaign. The Pitot probe was held at a constant (R,Z) location and the anode was spun beneath it to sample at different θ locations.

point in the channel was measured around the channel azimuthally. That is, the pressure probe was aligned to a single point in the channel cross-section (R,Z), and the anode was then rotated beneath the probe tip to measure that same point at various azimuthal locations (θ) around the channel. A photograph of the setup, including the coordinate axes, is shown in Figure 6.11.

Simulations suggested that maximum non-uniformity would occur approximately halfway down the channel length [4]. The (R,Z) location of the probe tip was set to be on channel centerline at a Z close to the middle of the channel length to ensure that measurements were taken near this region. Due to slight differences in alignment, the Outer channel's Z location was approximately 3 mm farther downstream than that of the Inner and Middle channels, but this variation is approximately 5% of the channel length and still puts the measurement location within the region of interest inside the channel. Typically, azimuthal variation of somewhere between ±5% to ±12% is deemed acceptable for Hall thrusters, with higher uniformity preferred [202, 209]. For the X3, a criterion of ±12% was chosen due to the laboratory and proof-of-concept nature of the thruster.

Table 6.2: The test matrix for the azimuthal propellant uniformity test.

CHANNEL	$\Delta\theta_{exithole}$	$\Delta\theta_{datapoint}$	$\dot{m}_{min,design}$	$\dot{m}_{max,design}$	$\dot{m}_{min,test}$	$\dot{m}_{max,test}$
Inner	18°	10°	60 sccm	600 sccm	70 sccm	600 sccm
Middle	10°	8°	150 sccm	1200 sccm	150 sccm	1000 sccm
Outer	7.5°	10°	300 sccm	1500 sccm	300 sccm	1000 sccm

6.3.2.2 Test Procedures

Alignment of the anode to the Pitot tube was done before each test run. The anode was rotated repeatedly in 90-degree increments, and the radial and axial location of the Pitot tube was measured using a ruler with 0.5 mm indications. This was repeated as the anode was shifted around until alignment was consistent in both directions at each of the four 90-degree locations to ± 1 mm. The starting (or zero-degree) point on each anode was adjusted before each pumpdown to coincide with the top or "12-o'clock" propellant inlet. Alignment occurred with the propellant lines attached to the anode to minimize bumping or jostling once alignment was complete.

Each anode of the X3 features a similar design, but the exact spacing of the downstream holes varies channel to channel. To ensure that a representative pressure distribution was being measured for each anode, the spacing of each pressure measurement was adjusted accordingly. The low-flow and high-flow flow rates were chosen based on the expected range of flow rates for each anode. However, it was found during testing that the Stabil gauge would reach the upper limit of its range at flow rates above 1000 sccm. Thus, the high-flow conditions for the Middle and Outer channels were set to be 1000 sccm instead of their true maximum flow rate. Based on results from other thrusters, as well as the results discussed below, it was expected that the flow would become more uniform at these higher flow rates, and the data support this conclusion. Details of the test matrix are provided in Table 6.2.

At the beginning of each test, we powered the pressure gauge and allowed it to warm up for approximately one hour. Once the pressure measurements were steady, we began flowing gas to the anode. To allow the vacuum facility, pressure gauge, and mass flow controller

to reach steady-state operation, we began measurements approximately 2-3 minutes after starting the flow. Typically, the pressure measurements were steady after less than one minute.

Testing proceeded as follows:

1. The anode was rotated to a new measurement location
2. The system was allowed to settle for 10 seconds
3. The pressure was recorded
4. Repeat

We recorded the data from the readout on the front of the Stabil gauge controller and marked settling time with a stopwatch with 1-second resolution. For each flow rate, we spun the anode counter-clockwise from 0°–360°, then clockwise from 360°–0°. Because there was a small amount (approximately 2-3°) of backlash in the rotational stage, there were two benefits to mapping both directions at each flow rate: because of the offset, the forward and backward data sets provided a more thorough mapping of the pressure distribution across the anode at a given flow rate; additionally, the data are more directly comparable between the low- and high-flow conditions (as opposed to mapping 0°–360° at one flow rate and 360°–0° at the other). We tested the low-flow condition first for each anode, followed by the high-flow condition.

6.3.2.3 Results

We present the results as normalized pressure, wherein the pressure value at position θ are normalized by the mean value of the data set at a given flow rate:

$$\|p\|(\theta) = \frac{p(\theta)}{\overline{p}} \tag{6.1}$$

where $\|p\|(\theta)$ is the normalized pressure at position θ, $p(\theta)$ is the measured pressure at position θ, and $\overline{p}$ is the mean pressure of the set of data. The deviation values are then simply calculated as such:

$$\sigma_{max} = \|p\|_{max} - 1 \tag{6.2}$$

$$\sigma_{min} = 1 - \|p\|_{min} \tag{6.3}$$

where σ is the deviation and the subscripts *max* and *min* represent the upper and lower values, respectively.

The Outer channel was the first anode tested, and the most critical. During the design phase, the Outer anode was identified as the most likely to have uniformity issues due to its size [4]. Also, it has the largest range of flow rates, which meant that the single-point uniformity test done before thruster assembly was especially lacking. Finally, as mentioned above, there was potential that a non-uniformity in flow was causing the asymmetric anode heating issue. Results from the low- and high-flow cases are shown in Figure 6.12. For the 300 sccm case, the normalized pressure ranged from 0.887 to 1.117 (a spread of +11.7%/−11.3%). For the 1000 sccm case, the normalized pressure ranged from 0.908 to 1.092 (a spread of ±9.2%).

We can draw three major conclusions from these data: First, the pressure distribution falls within the uniformity criterion set for the thruster, even at the low-flow condition. Second, the flow increases in uniformity as the flow rate increases, as expected. Third, there are no indications in these data of any low or high region of flow that might cause the asymmetric anode heating.

The results from the Middle channel are presented in Figure 6.13. The Middle channel results are presented on the same scale as the Outer channel's results to facilitate more meaningful comparisons. At 150 sccm, the normalized pressure ranged from 0.932 to 1.079

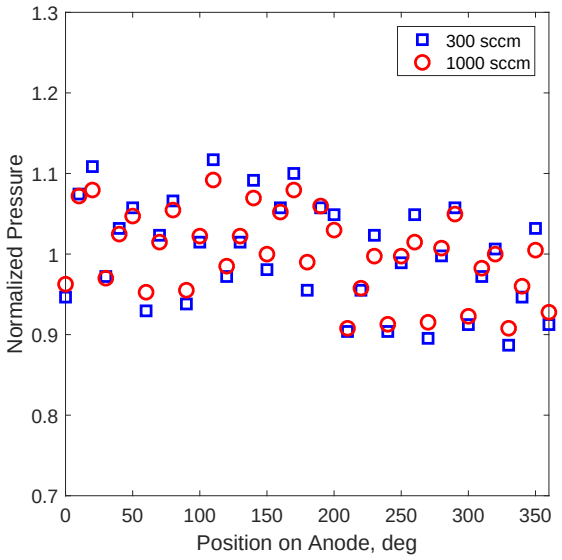

Figure 6.12: Propellent azimuthal uniformity results from the Outer channel at 300 sccm and 1000 sccm.

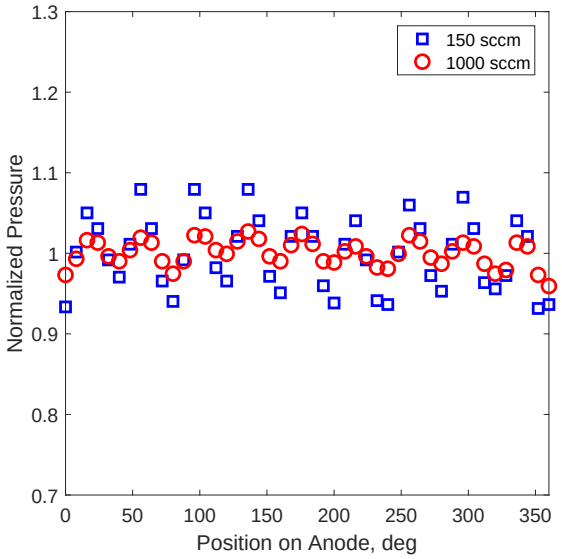

Figure 6.13: Propellant azimuthal uniformity results from the Middle channel at 150 sccm and 1000 sccm.

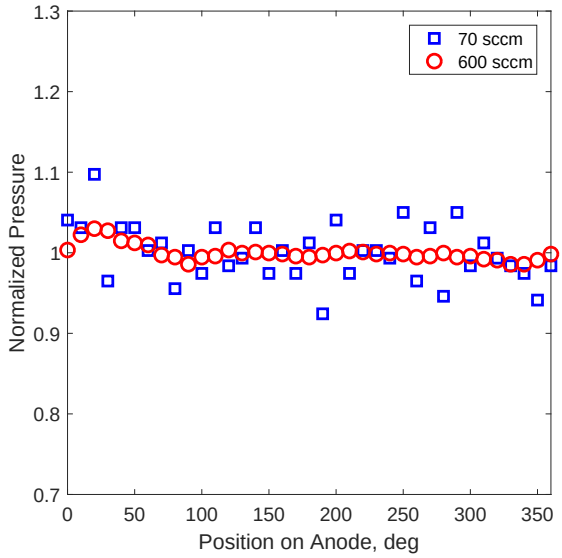

Figure 6.14: Propellant azimuthal uniformity results from the Inner channel at 70 sccm and 600 sccm.

(a spread of +7.9%/−6.8%). At 1000 sccm, the normalized pressure ranged from 0.959 to 1.027 (a spread of +2.7%/−4.1%). It should be noted that the high-flow condition here is closer to the designed maximum flow rate for this anode than for the Outer, so the increased uniformity seen here is expected. These results indicate uniform flow well within the criterion at both flow rates, and an increase in uniformity as flow rate increases.

The results from the Inner channel are presented in Figure 6.14. Again, these results are plotted on the same scale as those for the Middle and Outer channels above. At 70 sccm, the normalized pressure ranged from 0.924 to 1.097 (a spread of +9.7%/−7.6%). At 600 sccm, the normalized pressure ranged from 0.986 to 1.030 (a spread of +3.0%/−1.4%). Here, the high-flow condition is the full maximum flow rate for this anode. Once again we find the same trends: flow that is uniform within the set criterion at the low flow conditions and that increases in uniformity at the high-flow condition. It can be seen that there is a detectable high spot on this anode between 20°–50° in the high-flow condition. This was seen on both the 0°–360° map and the 360°–0° map, indicating that it is a real effect and not, for instance,

a transient issue with the flow controller or the pressure gauge. Possible causes of this spot include local geometry issues such as a hole slightly larger than the specified tolerance or a baffle out of specification. Another possibility is test rig misalignment, but because there is not a corresponding low point we deem this unlikely. However, it is important to note that though this high spot is noticeably higher than the rest of the flow, it still falls within comparable or better uniformity than the Middle and Outer anodes at high flow. Therefore, the location of this high spot will be noted and monitored during future testing, but we expect that the small variation in flow will not cause any issues with thruster operation.

A final comparison between the data sets is presented in Figure 6.15. The average data deviation $\overline{\sigma}$ was calculated by averaging the spread values for each set of data presented above:

$$\overline{\sigma} = \frac{\sigma_{max} + \sigma_{min}}{2} \tag{6.4}$$

This is then plotted against the normalized mass flow rate, which is calculated by dividing the tested mass flow rate $\dot{m}$ by the maximum design mass flow rate $\dot{m}_{max,design}$:

$$\|\dot{m}\| = \frac{\dot{m}}{\dot{m}_{max,design}} \tag{6.5}$$

Plotting in this manner helps to illustrate the effect of the artificial maximum imposed on the tested mass flow rate by the pressure gauge limitations. Data is presented such that dotted lines between each anode's conditions imply general trends; because only two data points were collected, no suggestion or implication can be made about the actual shape of the $\overline{\sigma}$ vs. $\|\dot{m}\|$ curve. We present these data with error bars representing the largest repeatability error—that is, the largest deviation between measurements at a given location for a given condition—recorded during the test. This figure further illustrates that, in all cases, flow uniformity increased as flow rate increased. It shows that all three anodes followed a similar slope across their range (regardless of what normalized flow rate the high-flow condition was

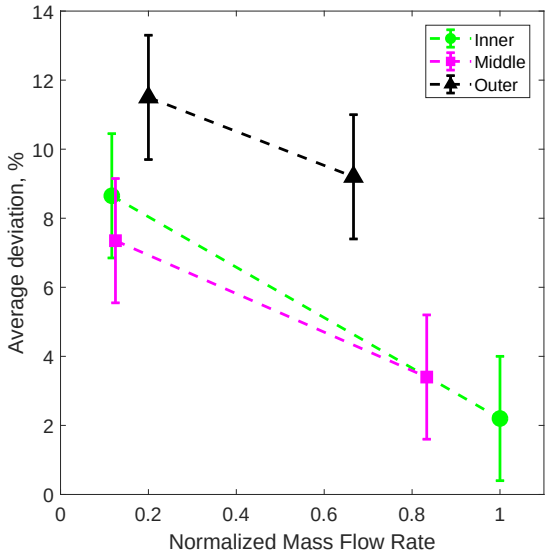

Figure 6.15: Average data deviation from the normalized pressure data for each flow condition of each anode plotted against normalized mass flow rate. The mass flow rate has been normalized by the maximum design flow rate for each channel. Error bars represent the maximum repeatability error from the test campaign to provide a measure of uncertainty.

at). This then indicates that the Outer channel may produce flow uniformity around ± 8% at its highest flow rate if the trend continues as flow rate increases.

However, this figure also illustrates that the uniformity of the Inner and Middle channels is roughly comparable, yet the Outer channel is systematically less uniform. This suggests that either the construction of the Outer anode is out of specification, that the specific design for the Outer anode is lacking (for instance, the hole spacing is improper), or that the drop-in anode concept itself does not apply well to anodes of this size. Further investigation is needed as the development of the X3 continues, but even the notably-higher non-uniformity of the Outer anode falls within the stated acceptance criterion.

6.3.2.4 Conclusions

Much like the uniform results of the magnetic field mapping, the results of this flow mapping indicate uniform propellant flow for all three anodes and do not provide any indication

of an issue that would cause asymmetric anode heating. Each anode provided flow that was satisfactorily uniform at its low-flow condition and displayed increased uniformity at increased flow. A slight high spot was found on the Inner channel between 20°–50° at the high-flow condition. Though this spot was detectable and repeatable, the variation still falls well within the acceptability criterion for uniform flow.

6.4 Thruster Disassembly and Inspection at GRC

Following the detailed magnetic field and propellant flow mapping performed at PEPL, we shipped the X3 to NASA GRC for further disassembly, inspection, and repairs. We identified a number of improvements to make to the thruster, which we detail below. These improvements were implemented before the thruster was assembled for high-power testing at GRC.

Upon disassembling the thruster, we found a large number of arc marks on the outside surface of all three channel cups. These spots often appeared as pock marks through the aluminum oxide coating on these surfaces that exposed bare metal. They typically clustered together, with as many as twenty such small marks grouped in a given location. The arcing appeared to be occurring between the channel cup, which is at anode potential during thruster operation, and the magnet bobbins, which are grounded through the thruster body. A portion of these arcs were in the direct vicinity of thermocouples mounted to the magnet bobbins facing the channel cup. On Magnet 5, one of the thermocouples had been completely destroyed by what appeared to be repeated arcing between the channel cup and the bobbin.

After removing the anodes from inside the channel cups we discovered that a significant number of the studs on the back of the anodes had cracked welds. Figure 6.16 shows a photograph of one such cracked weld. Each anode features studs serving a variety of purposes, including ones to mount the anode into its channel cup, ones for electrical power, and ones for propellant. These studs are set in through-holes in the back plate of the anode,

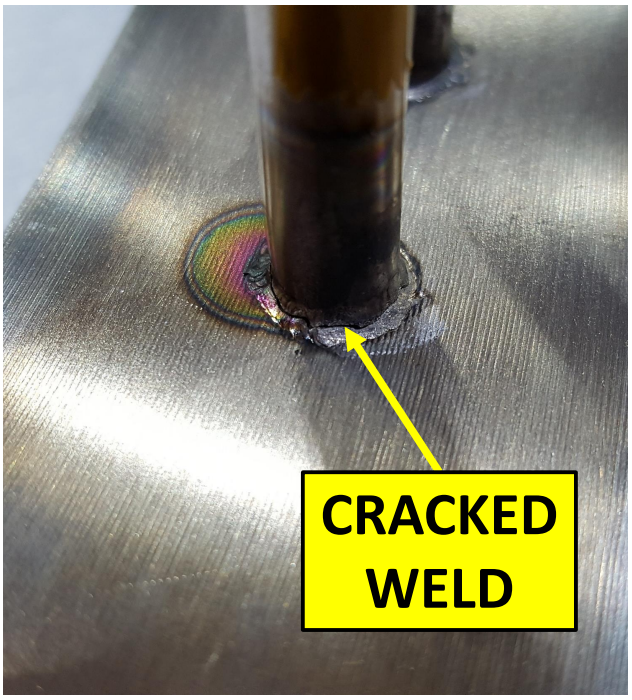

Figure 6.16: An example cracked anode weld found during the thruster inspection at GRC.

and the weld serves as the seal to keep propellant from leaking out. With these welds cracked, sometimes badly enough that a gap was visible to the naked eye, propellant was free to leak out of the backside of the anodes. The Inner anode showed no signs of cracking; the Middle channel had cracks in approximately 40% of its studs; the Outer had cracks in approximately 90% of its studs.

We performed a benchtop leak test of each of the anodes in which the anodes were sealed and pressurized with 1–2 psi of nitrogen. Leaks were identified with isopropyl alcohol, which bubbled when applied to a leaking weld. All cracked welds produced bubbles and indicated leaking. Because these studs all pass through the channel cup and into the thruster body (for mounting or for electrical and propellant inputs), this leaking propellant was free to enter the thruster chassis. Further inspection of the arc marks on the channel cups indicated that a vast majority of them clustered around locations where anode studs passed through. Leaks were likely creating localized pressure increases that caused the breakdowns. We suspect that flaws in the channel cup aluminum oxide coating may have been exacerbating

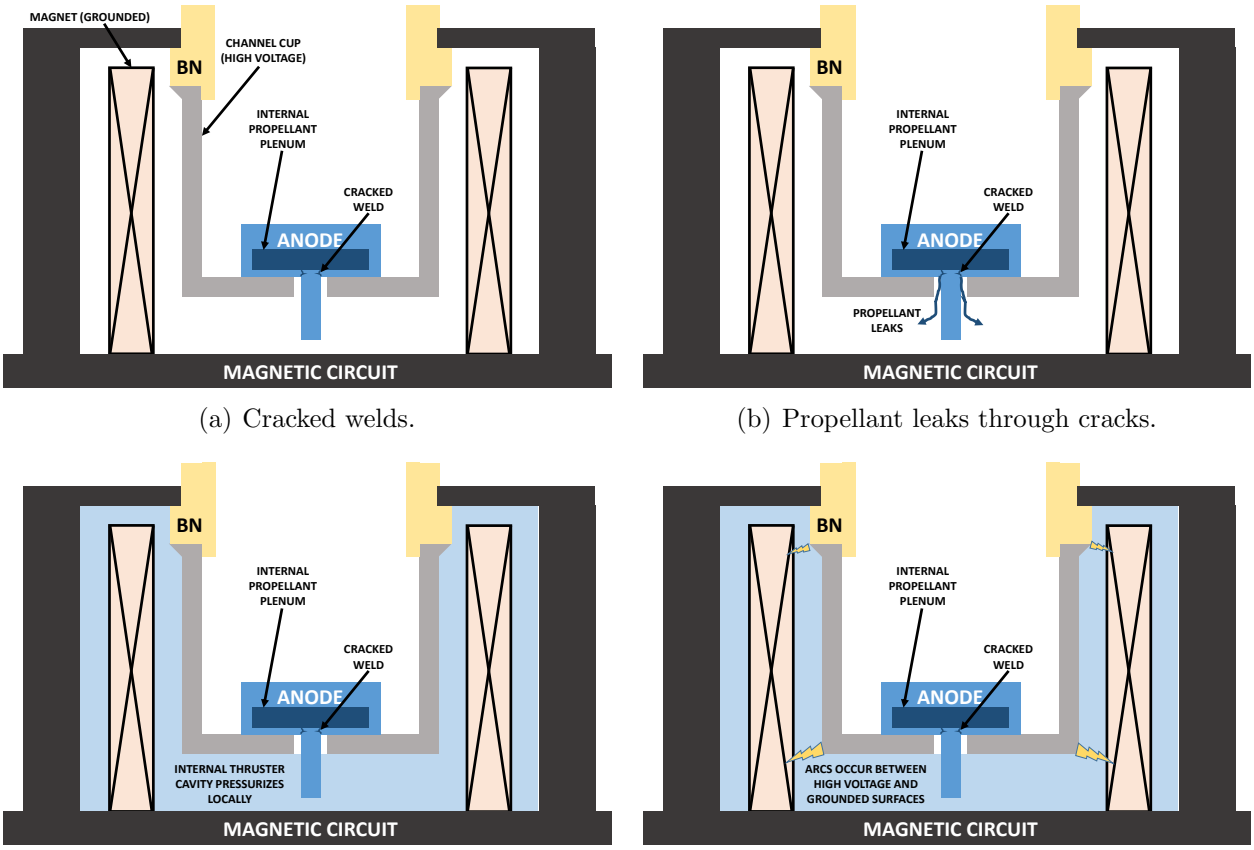

(a) Cracked welds.

(b) Propellant leaks through cracks.

(c) Internal thruster cavity pressurizes.

(d) Arcing occurs.

Figure 6.17: The suspected sequence of events in which the cracked anode studs caused arcing inside the thruster chassis between the high-voltage channel cup and the grounded thruster body. (a) The welds are cracked (b) Propellant leaks through cracks (c) The internal thruster cavity pressurizes locally due to leaks (d) Arcing occurs between high-voltage channel cup and grounded magnet bobbins.

the issue by creating areas where bare metal was exposed at high voltage. A series of notional thruster diagrams illustrating the suspected mechanism is shown in Figure 6.17. Additionally, propellant leaking from the back of the anode is not being used for thrust generation. Though it is unclear when these leaks appeared, they are a likely mechanism for some of the lower-than-predicted performance of the X3 during low-power testing.

We implemented two improvements to fix these issues. First, all anode welds were repaired by NASA GRC. The cracks were likely caused by stresses from thermal cycling of the thruster during operation. These welds were significantly undersized from what was specified

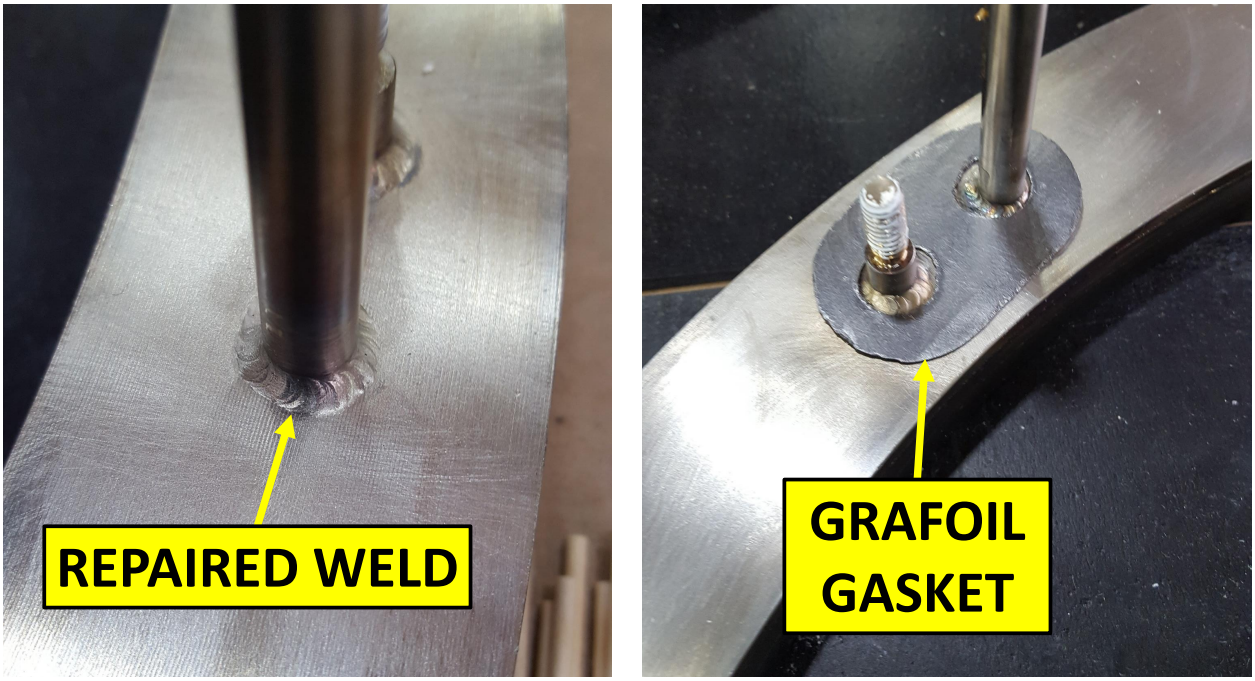

(a) Repaired weld.

(b) Grafoil gasket.

Figure 6.18: Photographs of the two improvements made to the anode welds to minimize propellant leaking. All studs were repaired in this fashion regardless of whether they were cracked.

in the thruster design. Thus, even the welds that had not yet cracked were modified to be brought to specification. Additionally, we created gaskets out of grafoil to be installed between the anode and the channel cup. These were intended to ensure the anode had a strong seal against the channel cup that would be maintained at the high anode temperatures of thruster operation. We present example photographs of both repairs in Figure 6.18.

In addition to these modification, we also removed all thermocouples from the magnet bobbins and smoothed over their former locations to minimize sharp protrusions that could induce arcing. We also crafted mica strips that were installed between the channel cups and the magnet bobbins. These strips, which spanned the full length and circumference of the magnet bobbins on both sides of each channel cup, used GRC best practices to increase the isolation between the high-voltage and grounded portions of the thruster. A similar technique was used on the NASA-173GT thruster [210]. In an effort to minimize outgassing

during the high-power test, we sent these strips to be baked out in a vacuum furnace at 500 °C for four hours before installing them in the X3.

As will be demonstrated below, these repairs were very successful. After the GRC test, which included short-duration operation to a maximum total power of 102 kW and three near-continuous hours of operation at 80 kW, no evidence of internal arcing or propellant leaks were identified. The mica sheets were mostly intact after the test except for two small spots where they appeared to arrest a breakdown before it reached the grounded magnet bobbin. Thus, these improvements will be maintained in the X3 for future operation.

6.5 Summary

Ultimately, no clear cause of the glow region was discovered. Both the magnetic field and the propellant distribution were found to be satisfactorily uniform, and nothing in the thruster disassembly and inspection provided any indication of a physical cause of the spot. This inspection did yield a number of important improvements to be made to the thruster. These included welds that required repairing as well as high-voltage isolation that required improvement. We used best practices developed at GRC in previous high-power thruster testing to solve these issues.

Further quantitative information about the spot is necessary to continue with diagnosing its cause. Possible methods include a thermal camera inside the vacuum facility, which would provide a temperature profile around the anode after thruster shut-down. A better understanding of the severity of the temperature difference will help to inform any future work to investigate its cause.

CHAPTER 7

High-Power Performance

"To conceive of an airplane is nothing, to make an airplane is little, but to try it out and fly, that is everything."

– Otto Lilenthal, German glider pioneer [211]

7.1 Introduction

The next phase of X3 characterization occurred in NASA GRC's VF-5, which features a pumping speed capable of maintaining acceptable facility background pressure during X3 operation up to 250 A discharge current and in excess of 100 kW discharge power. We describe this test in two chapters. The first (here) describes the performance of the thruster as measured with the thrust stand, and the second describes the efficiency analysis using the plasma diagnostics as well as the detailed discharge current oscillation analysis using the high-speed camera. In this chapter, we present the thrust, specific impulse, and efficiency of the X3 up to 102 kW total power in Section 7.2. We also present a comparison of operation in single-channel and multi-channel configurations. In Section 7.3 we discuss the implications of these results, including how they compare to previous X3 results and those from other NASA high-power Hall thrusters, as well as possible mechanisms for the multi-channel results we observed.

7.2 Performance Results

We successfully operated the X3 across a throttle table spanning 300–500 V discharge voltage and three current densities. Using the middle value as reference, these current densities were 0.63, 1.00, and 1.26 the reference value. All seven possible thruster configurations (I, M, O, IM, IO, MO, and IMO) were fired at each discharge voltage for the $0.63 j_{ref}$ and $1.00 j_{ref}$ conditions. Additional test points were collected at 300 V and 400 V at the $1.26 j_{ref}$ condition, bringing the total unique test conditions to 46 for this performance characterization. For all test points, all firing channels were held at the same discharge voltage and current density. We designed the test matrix to study the performance trends with both discharge voltage and current density, as well as to study how the performance varies for different channel combinations at fixed values of these parameters.

7.2.1 Thruster Operation

For all performance measurements here, the thruster was allowed to settle until the mean discharge current of each channel was stable and drifting no more than 0.01 A/min. We did not allow the thruster to come into thermal steady state during these measurements due to the length of time necessary to thermally soak the thruster (expected to be in excess of 6 hours depending on the operating point). This concession is one made previously with large high-power Hall thrusters [10]. Although we did not operate the thruster in thermal steady-state, hours-long continuous operation at 50 and 80 kW indicated thruster performance and behavior was not changing significantly during thermal soaking. We did, however, thoroughly bake out each channel at a given power level prior to any performance measurements at that power level. This incremental bakeout process was used to limit the risk from high-power operation. At each new discharge voltage/current density condition, the thruster would experience an initial period of discharge current instability associated with moisture and other contaminants outgassing from the channel walls and other thruster components. Each

channel was fired individually until this behavior passed (typically on the order of an hour at a new condition, but three to four hours the first time the thruster was fired after being exposed to atmosphere), and then we conducted performance measurements.

At each condition, we performed limited magnetic field sweeping before collecting the performance data, varying the magnetic field strength while maintaining the field shape as described previously. The anode flow rate was held constant during the sweep and was then adjusted as needed after the sweep was complete. During a sweep, thruster discharge, oscillation, and performance parameters were monitored. The optimal field strength for a given condition was that which provided the minimum discharge current. Typical sweep ranges were from 0.8 to 1.3 the same reference field strength $\boldsymbol{B}_{nom}$ used in Chapters 5 and 6. Previous field sweeps performed during the low-power characterization 2 described in Chapter 5 were more extensive, and because the sweeps during this campaign matched the trends of the previous results we did not extend our range here. The reference field strength was the same for each channel and did not vary between single- and multi-channel conditions, although the magnet coil current ratios had to be modified in multi-channel operation, as explained in Reference [205]. Optimized field strengths across all conditions did not vary by more than 30%.

Due to the complexities involved in multi-channel magnetic field optimization (where the field of each channel affects the others [205]), only minimal efforts were undertaken for multi-channel conditions. Typically, multi-channel conditions were simply operated at or near the optimum magnetic field strength found for single-channel operation at the given discharge voltage and current density. The field strength was kept constant for all firing channels for multi-channel operation. Because of this, we speculate that further performance optimization through magnetic field tuning may be possible for the multi-channel conditions, as will be discussed below.

7.2.2 Thruster Performance

Thruster performance is evaluated here using both anode and total quantities. The anode values provide an opportunity to compare thruster behavior while removing the electromagnets and cathode, neither of which has a design optimized for flight, and the total values provide insight into how the X3 thruster will fit into the XR-100 system. Anode and total efficiency and specific impulse calculations are described in Chapter 2. Thruster telemetry values used in these calculations were averaged over a 60-second period.

The uncertainty in the thrust measurement was the dominant contribution to the uncertainty in efficiency and specific impulse calculations. Because the thrust uncertainty increased at lower thrust values (due to the constant-value uncertainty from the inclination resolution), lower-power conditions typically had slightly larger uncertainties in efficiency and specific impulse as well. Average uncertainties for specific impulse were ± 40 s and for efficiency were ± 0.04. Inner channel conditions typically were closer to ± 100 s and ± 0.06 because of their higher relative thrust uncertainty. These uncertainties are reflected in the error bars on the plots below.

Figure 7.1 presents thrust as a function of discharge power. For the conditions tested, the X3 produced a maximum thrust of 5.42 N ± 0.1 N at 98.4 kW discharge power (400 V, 247 A). At 101 kW discharge power (500 V, 201 A) the X3 produced 5.03 N ± 0.1 N. At a fixed discharge voltage, the X3 operated at a similar thrust to power ratio (T/P) in each of the seven channel combinations. Average T/P values are plotted alongside the data in the figure and are compared to those of other NASA high-power Hall thrusters in Table 7.1. As shown in the table, the average T/P results are slightly improved over those of other high-power Hall thrusters, indicating that the X3 is operating as designed. These results also demonstrate that the T/P value attained is not dependent on the channel combination. This was not the result seen during previous low-power, low-current density operation of the X3 at PEPL [212]. We discuss potential reasons for this improved performance below.

Figure 7.2 shows anode and total specific impulse as a function of discharge power.

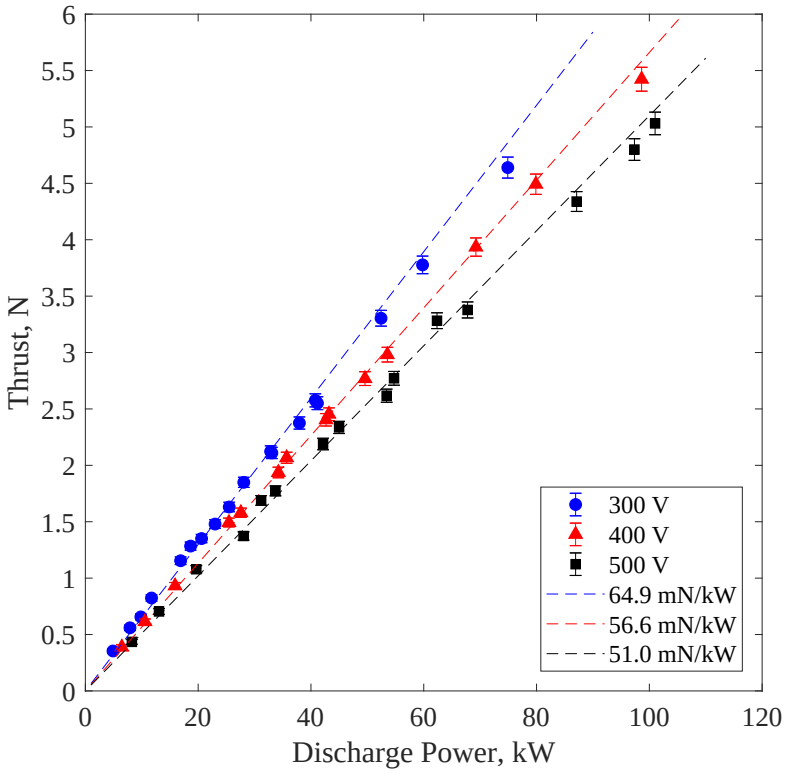

Figure 7.1: X3 thrust versus total discharge power. Error bars are plotted for all test points; for all but the highest-thrust measurements, these error bars are within the size of the marker. Lines of average thrust to power ratio for each discharge voltage are also plotted.

Table 7.1: The average thrust-to-power values for the X3 and other NASA high-power Hall thrusters for each discharge voltage tested.

V_d	NASA-457Mv1	NASA-400M	NASA-300M	X3 NHT
300 V	57.8 mN/kW	59.0 mN/kW	63 mN/kW	64.9 mN/kW
400 V	51.9 mN/kW	49.7 mN/kW	56 mN/kW	56.6 mN/kW
500 V	47.2 mN/kW	47.4 mN/kW	50 mN/kW	51.0 mN/kW

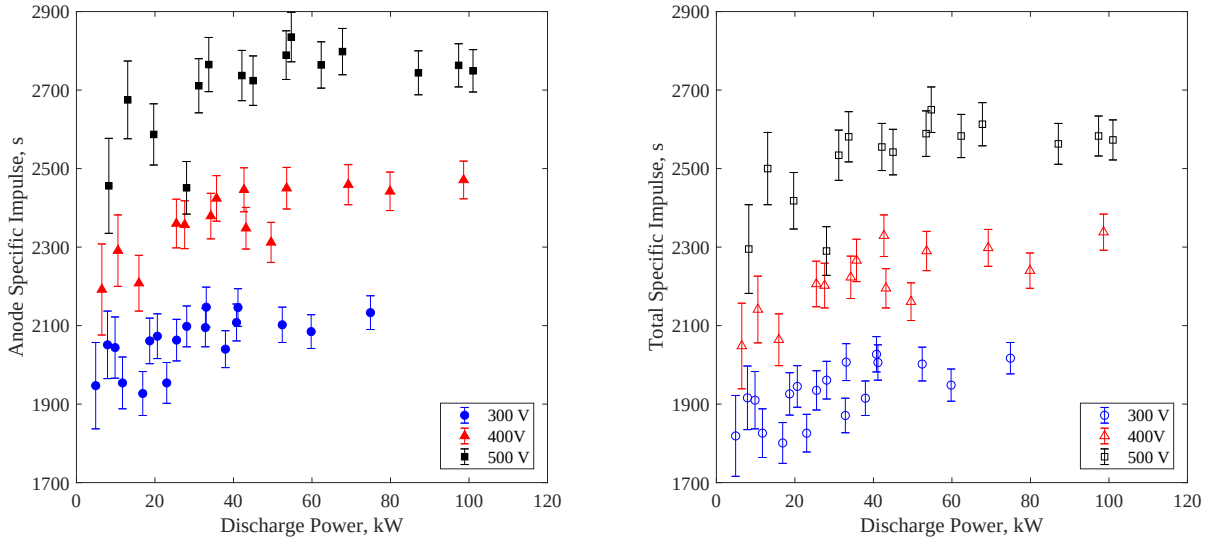

(a) Anode specific impulse

(b) Total specific impulse

Figure 7.2: X3 specific impulse as a function of discharge power for different discharge voltages: (a) anode and (b) total values. Error bars indicate the statistical uncertainty in the measurement for each point.

Table 7.2: Comparison of anode specific impulse ranges at different discharge voltages from the X3 and other NASA high-power Hall thrusters.

V_d	NASA-457Mv1	NASA-400M	NASA-300M	X3 NHT
300 V	1750–2100 s	1700–2100 s	1900–2200 s	1930–2150 s
400 V	2100–2400 s	2000–2600 s	2200–2600 s	2190–2470 s
500 V	2400–2750 s	2500–3000 s	2500–2900 s	2480–2840 s

We found that the anode specific impulse for the conditions tested ranged from 1930–2150 seconds (± 50 seconds) for 300 V, 2190–2470 seconds (± 50 seconds) for 400 V, and 2480–2840 seconds (± 60 seconds) for 500 V, and that except for a low-power drop off, specific impulse for a given discharge voltage were roughly constant. This trend was true regardless of channel combination. Because of the large number of channel combinations for the X3, these different configurations are not noted in these figures for clarity. The X3's values match those found for other high-power Hall thrusters, as shown in the anode specific impulse ranges presented in Table 7.2.

Table 7.3: Comparison of anode efficiency ranges at different discharge voltages from the X3 and other NASA high-power Hall thrusters.

V_d	NASA-457Mv1	NASA-400M	NASA-300M	X3 NHT range	X3 NHT avg.
300 V	0.50–0.57	0.55–0.59	0.60–0.67	0.62–0.71	0.66
400 V	0.53–0.60	0.60–0.65	0.55–0.73	0.63–0.71	0.66
500 V	0.46–0.65	0.66–0.71	0.60–0.73	0.58–0.72	0.68

Table 7.3 presents the anode efficiency ranges and averages for the X3 versus discharge voltage, as well as ranges for other high-power Hall thrusters. Anode efficiency is used for this comparison because total efficiency values were not published for all other thrusters. It can be seen that the average efficiency of the X3 trends upwards with increasing discharge voltage, a trend shared with the other thrusters. In addition, the efficiency at a given discharge voltage increases from the NASA-457Mv1 to the NASA-400M to the NASA-300M. This is due to the fact that each subsequent thruster benefited from lessons learned from the prior designs. For example, the NASA-400M implemented an improved magnetic field design based on the development of the NASA-457Mv1 and NASA-173M that was intended to improve the efficiency of its operation [3, 79].

During this test, the X3 operated at a peak total efficiency of 0.64 ± 0.03 at 400 V and 0.67 ± 0.03 at 500 V. In general, the total efficiency was approximately 0.05 lower than anode efficiency regardless of operating configuration, as can be seen in the throttle tables presented in Appendix C. For reference, the NASA-300M demonstrated peak total efficiencies of 0.67 at 400 V and 0.66 at 500 V. Thus, it is apparent that the X3 is operating at the state of the art, and continued magnetic field optimization may even further increase these values.

7.2.3 Comparison of Single-Channel and Multi-Channel Operation

7.2.3.1 Performance

With three discharge channels operating in close proximity, there is the expectation that coupling between the channels could effect performance. Indeed, previous experiments with the X2 6-kW two-channel thruster demonstrated increased thrust of up to 11% when firing both channels simultaneously as compared to the sum of each channel firing individually. A recent test campaign by Georgin [187] and Cusson [204] investigated this phenomenon in more detail and found that the thrust increase was replicated when flowing propellant through the non-firing channel in single-channel mode. They attributed this effect, and the changes observed in the plasma that they found to be causing the increase in thrust, to be due to the change in the neutral pressure field close to the thruster caused by the flow from the adjacent channel. Beal identified similar trends in a cluster of 200-W Hall thrusters, finding that a single thruster could couple normally with a cathode on the other side of the two-thruster cluster if propellant was flowing through the intermediate thruster [213]. Plasma plume measurements suggested that this extra neutral flow was likely improving electron transport across field lines via collisional effects. Unfortunately, thrust was not measured in that study, so it is unclear whether there was a performance effect due to the increased local neutral pressure.

We operated the X3 in its single- and three-channel configurations for seven different combinations of discharge voltage and current density. Figure 7.3 presents the thrust produced by the X3 at each of its IMO-configuration points alongside the summed thrust from the I, M, and O configurations at the same conditions. Error bars reflect the thrust measurement uncertainty for each point, typically around 2%. We found that the summed I, M, and O conditions often produced slightly more thrust than the IMO condition, but that in general the measurements matched to within their uncertainty. A discussion of these results

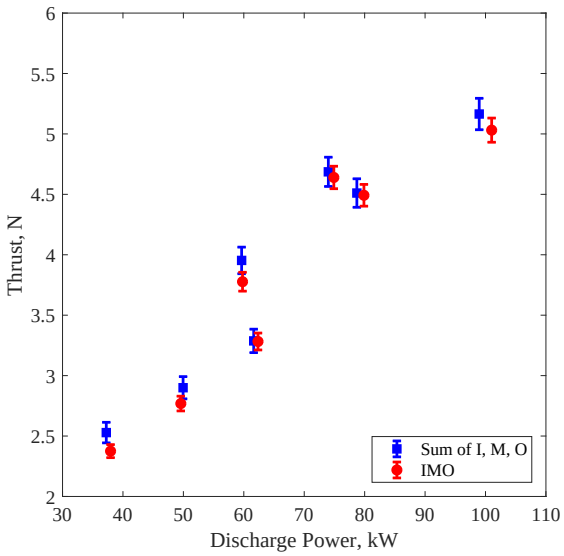

Figure 7.3: Comparison of three-channel thrust and summed single-channel thrusts at the same discharge voltage and current density for the X3.

and their significant in the context of these previous results by Georgin, Cusson, and Beal, is presented in Section 7.3 below.

7.2.3.2 Oscillatory Behavior

Discharge current oscillations have been shown to affect Hall thruster performance [48] and their effects have been proposed to have roles in thruster processes such as anomalous transport [49, 50] and cathode erosion [51]. Work has shown that the magnetic field strength is a major driver of oscillation strength and character in unshielded Hall thrusters [48]. Characterizing discharge current oscillations was an important part of understanding operation of the X3.

We collected peak-to-peak (P2P) and root-mean-square (RMS) measurements of the discharge current oscillations during this performance mapping campaign. These measurements were made by the current guns and oscilloscopes described above in Chapter 4 and read by the data logger during telemetry measurement cycles. The values presented here are aver-

aged over the same 60 second period as the telemetry used in the performance calculations. For simplicity, only the P2P values are presented. The RMS values showed similar trends between conditions but in all cases were a smaller percentage of the mean discharge current. Figure 7.4a presents the P2P values normalized by their respective channel's mean discharge current. Single- and multi-channel conditions are indicated by closed and open markers, respectively. Figures 7.4b–d show the average values for each channel in single- and multi-channel operation for each discharge voltage tested. The error bars on the bar charts represent the standard deviation of the values for each case.

The figures show that P2P values for all three channels were less than 40% for single-channel operation at all discharge voltages. The Outer channel's oscillation strength typically was unchanged or grew slightly in multi-channel operation but remained below 50% for all but a single condition (at 500 V). The Middle channel's oscillation strength grew slightly more in multi-channel operation, rising to an average of about 50% for 300 V and 400 V and closer to 75% for 500 V. The Inner channel experienced the largest difference in oscillation strength between single-channel and multi-channel operation, rising from a single-channel average less than 40% the mean values (comparable to other two channels) to a multi-channel average approximately 70% at 300 V and 400 V and an average approaching 100% at 500 V.

7.2.3.3 Cathode Coupling

Another metric to compare the operation of the X3 in single-channel and multi-channel modes is the cathode-to-ground voltage (V_{cg}). Hall thruster cathode-to-ground voltage is typically between 5 and 30 V below facility ground, and is a metric of how well the thruster is coupling to the cathode [214–216]. Values of V_{cg} that are smaller (less negative) are indicative of improved coupling, and these smaller values allow for more of the thruster discharge voltage to be available for ion beam acceleration. Substantial differences in V_{cg} between single- and multi-channel conditions for the X3 would indicate that the thruster is coupling to the cathode differently between these conditions. Figure 7.5 presents V_{cg} for

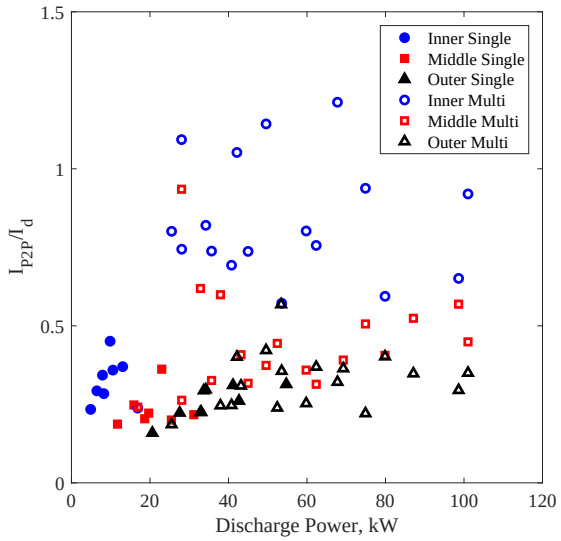

(a) P2P versus discharge power for all test conditions.

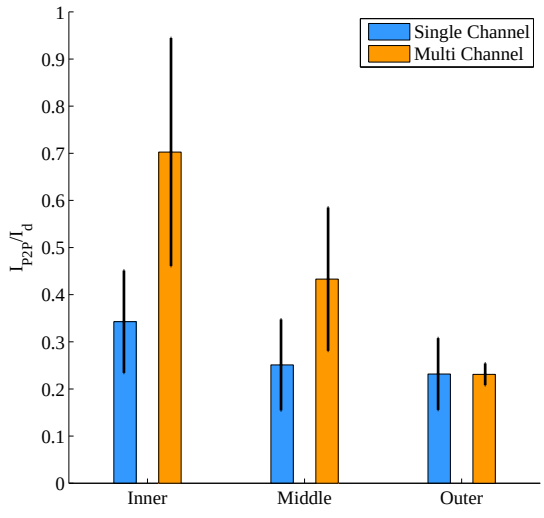

(b) 300 V average P2P values for each channel in single- and multi-channel operation.

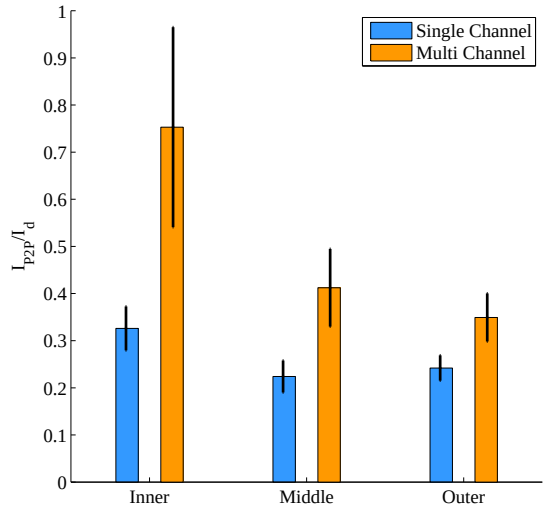

(c) 400 V average P2P values for each channel in single- and multi-channel operation.

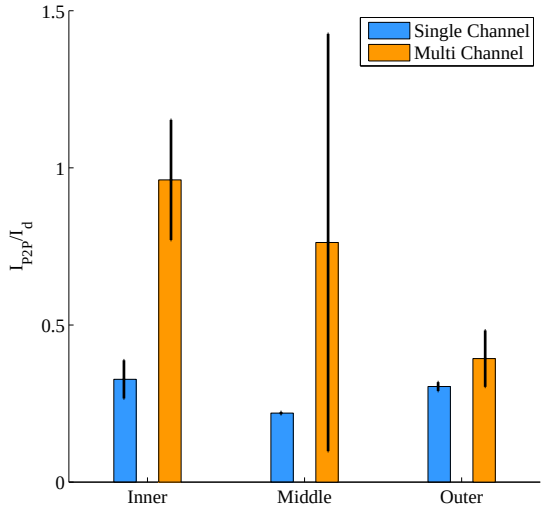

(d) 500 V average P2P values for each channel in single- and multi-channel operation.

Figure 7.4: Peak-to-peak discharge current oscillation values normalized by the mean discharge current for that channel: (a) versus discharge power for all test conditions and average values for each channel in single- and multi-channel operation for (b) 300 V operation, (c) 400 V operation, and (d) 500 V operation, respectively. Error bars on the bar charts represent standard deviation of data.

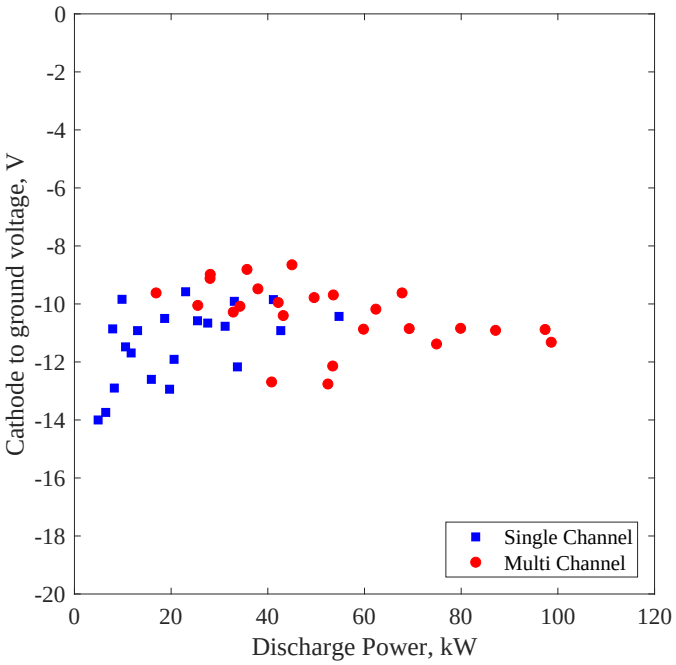

Figure 7.5: Cathode-to-ground voltage versus discharge power for all test conditions. Single- and multi-channel conditions are noted by square and circle markers, respectively.

both single- and multi-channel X3 operation as a function of discharge power. The TCFF was held at 7% of the anode flow for all test points here except one; this test point, at 5%, featured a V_{cg} of –10.9 V, no different than the rest of the conditions. The data show that V_{cg} varied between –8 and –14 V across all conditions and power levels tested here, demonstrating no difference between single- and multi-channel operation and demonstrating no trend with discharge power.

7.3 Discussion

7.3.1 Performance

The X3 was designed to leverage the physical insight and lessons learned from previous NASA single-channel high-power Hall thruster development. The results presented here indicate

that this effort was successful. The X3 demonstrated T/P, efficiency, and specific impulse values that are comparable or in excess of that family of high-power Hall thrusters and that showed similar trends with discharge voltage and power. The thrust was approximately linear for a given discharge voltage, and average T/P values matched those of the NASA-300M. Specific impulse was generally constant for a given discharge voltage, except at low powers. This trend is likely attributable to the fact that Hall thruster performance decreases at low current densities. The lower-specific impulse conditions for the X3 were at the lowest current density, $0.63j_{ref}$. The NASA-300M showed similar trends in specific impulse with discharge power (for fixed discharge voltage). All of these results indicate that the NASA-style Hall thruster scaling as developed by NASA GRC is applicable to the NHT configuration with equal success as was achieved with single-channel Hall thrusters, regardless of channel combination. These results differ from those of low-power characterization 1 from Chapter 5, which indicated that the Middle and Outer channels were under-performing expected values. We speculate that the improved performance here is a product of one or more factors. These include improved magnetic field settings as a result of a more thorough modeling effort [205], a number of thruster modifications made as part of the preparation for this risk reduction test, and the improved background pressures in VF5.

In addition to the high performance across the throttle table, these results represent a number of new achievements for laboratory Hall thruster operation. To date, the NASA-457Mv1 thruster demonstrated the highest total power operation (96 kW), the highest discharge current operation (112 A), and the highest thrust recorded (3.3 N), all reported by Manzella [11]. Here we demonstrated the X3 at different conditions to 102 kW total discharge power and 247 A total discharge current, and recorded a maximum thrust of 5.42 N (at 99 kW, 400 V).

The X3's demonstrated performance and extended Hall thruster power, discharge current, and thrust operating envelopes have important implications for the application of Hall thrusters to forthcoming missions. As discussed in Chapter 3, the need is apparent for elec-

tric propulsion systems operating at power levels in excess of 300 kW and specific impulses of 1500–3000 seconds, and that modeling work suggests that these systems should consist of individual thrusters on the order of 50–100 kW to minimize both system cost and mass. Our results demonstrate for the first time NHT operation at 100 kW. Although the X3 has yet to be throttled up to its full design power of 200 kW, these results have further implications beyond simply demonstrating 100-kW operation. The NASA-457Mv1 100-kW operation was at 845 V discharge voltage, providing a total specific impulse of 3460 seconds, and a total efficiency of 0.58. The 100-kW operation of the X3 here was at 400 V and 500 V discharge voltage and demonstrated total efficiency of 0.63 at both points. This lower-voltage operation provided total specific impulses of 2340 s at 400 V and 2570 s at 500 V, values squarely within the 1500-3000 second range identified by mission planners as ideal for cargo tugs and crew transport to destinations such as Mars [31, 133]. With these results, the X3 has demonstrated that Hall thrusters—and more specifically NHTs—are a viable propulsion option for upcoming high-power space missions of many types.

7.3.2 Multi-channel Operation Effects

7.3.2.1 Performance

There remain open questions regarding how the channels of NHTs interact with and affect one another while operating simultaneously. These data provide a certain amount of insight to these potential mechanisms. One possible manifestation of channel interaction would be in the performance: adjacent channels firing simultaneously could perhaps cause an increase or decrease in thruster performance due to localized effects. However, these data suggest that this is not occurring significantly with the X3, and that for a given discharge voltage and current density, each channel combination produces comparable performance. This is in stark contrast to previous work with the X2 6-kW NHT, which showed a thrust boost of 5–11% in dual-channel mode, a boost that would be well outside of the uncertainty of the measurements here. There are a number of possible explanations for this behavior:

1. We operated the X3 here in a constant-discharge current manner, adjusting the mass flow to achieve the target discharge current. This was selected due to the range of background pressures experienced during typical X3 operation. By throttling to target discharge currents, any background ingestion experienced by the thruster is accounted for and more direct performance comparisons can be made. Georgin and Cusson ran the X2 in a constant-mass flow manner for their work and did not report on thruster ingestion or changes in discharge current due to this effect, so it is unclear what role this may have had on their results. Liang, however, found the 9% thrust boost in his tests for constant mass flow rate (and background pressure) accompanied negligibly-small discharge current changes due to ingestion, so we suspect that our controlling for this does not account fully for the X3's behavior.

2. For single-channel operation, we only powered the magnets for the channel that was firing. This is different than work with the X2, where all four magnets of the thruster were left on throughout the test. Because we tuned the field strength at each condition, it is possible that we have obscured a trend that would be present if we operated the X3 such that magnetic fields were kept constant throughout I, M, O, and IMO operation.

3. The X3's cathode features downstream neutral injectors, as discussed in Chapter 4. These injectors were flowing at least small amounts of xenon during all M and O operation and during some I operation. Both the X2 work and Beal's work with thruster clusters demonstrates that neutral flow plays a large role in the behavior of multiple Hall thruster discharges in close proximity. Additional work by Brown has shown that near-cathode neutral injection can stabilize Hall thruster discharges [106]. A possible explanation for the X3's lack of apparent thrust boost is that the cathode neutral injectors are acting in a similar fashion to the neutral propellant that both research teams flowed through their non-firing anodes. The enhanced neutral pressure in front of the X3 due to these injectors may be affecting the discharge in the same

way, thus actually raising the single-channel performance to match that of the IMO condition, which experiences the multi-channel enhanced neutral pressure similar to the X2.

4. The X2's channels are scaled differently than those of the X3. It is possible that NASA-style Hall thruster scaling is less susceptible to these effects, or that another difference in design of the thrusters (unrelated to channel scaling) makes the X3 impervious to the effects observed on the X2.

Uncertainty of the measurements aside, it appears that the IMO condition produced less thrust than the sum of the I, M, and O conditions for certain cases, opposite the trend found with the X2. However, we expect that this is a product of magnetic field settings and not some kind of multi-channel loss mechanism at work. Because the single-channel conditions generally went through more rigorous magnetic field optimization than what was done for the IMO condition, it is not surprising that these conditions were producing slightly more thrust. We suspect that additional magnetic field tuning would bring the thrust of the IMO conditions up to match the sum of the single-channel thrusts. Further work, including more detailed magnetic field optimization, mapping the pressure field in front of the thruster with the injectors flowing, and operating the X3 in a similar manner as work with the X2, is necessary to fully explore this behavior.

7.3.2.2 Oscillatory Behavior

Another mechanism of communication between channels is through the oscillations. As we show in Figure 7.4, the peak to peak oscillations vary between single- and multi-channel operation for the Inner channel and to a lesser extend the Middle channel, yet stay roughly constant for the Outer. Potential explanations for this include a background pressure effect, a channel-coupling effect, or a magnetic field effect. There were large changes in facility background pressure between the single- and multi-channel conditions (up to 2x for the Outer

channel between O and IMO operation and up to 6x for the Inner channel between I and IMO operation), though previous work with the H6MS and the HERMeS thrusters [63, 111] has demonstrated that for those thrusters, oscillation strength was not strongly affected by changes in background pressure (though this was not the case for the H6MS with an externally mounted cathode). The cathode P2P value typically matched the largest P2P value from whichever discharge channel(s) were firing. When the Inner channel is operating in any multi-channel condition, it is thus typically coupling to a cathode with a larger P2P value. For instance, at a typical $1.0 j_{ref}$ condition, the Outer channel would be operating at 110 A discharge current and approximately 30–40 A P2P and the Inner would be at 26 A discharge current. What effects there may be from a 26-A discharge coupling with a cathode oscillating at 40 A P2P are unclear. Finally, the magnetic fields changed between single- and multi-channel operation. This effect is explored in work by Cusson [205]. Compared to single-channel operation, the Outer channel's electromagnets operate at nearly the same settings in multi-channel operation. However, the magnet current ratios for the Inner and Middle channels change significantly between these conditions due to the way that magnetic flux is shared among the magnetic circuit of the thruster.

To explore this phenomenon, we operated the Inner channel alone at 500 V, $0.63 j_{ref}$ at the IMO magnetic field condition (that is, with all six electromagnets on, set to the same conditions used during IMO thruster operation, with with no propellant or discharge power applied to the Middle and Outer channels). The resulting P2P strength is presented in Figure 7.6, compared against the results from single-channel operation (with only the Inner channel magnets on) and from IMO operation (with plasma discharge from all three channels). The result is striking. When the Inner channel is operated in the IMO-field condition, it exhibits oscillations that are much stronger than when it is operated in the I-field condition. This IMO-field condition was at the same facility background pressure as the I-field condition, eliminating this as a potential mechanism for the increased oscillation strength. This condition was also without the other channels firing, eliminating any cross-

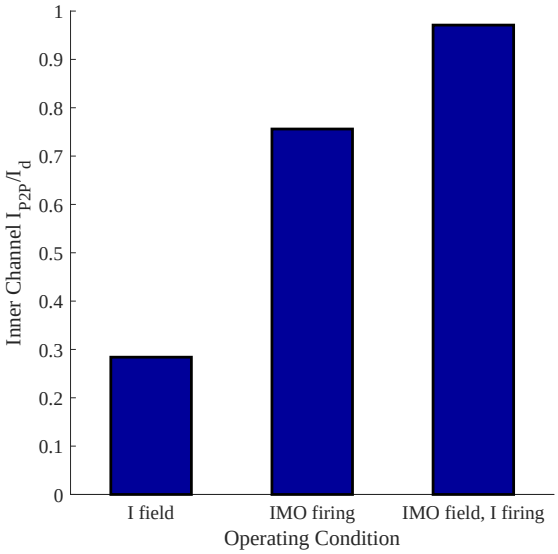

Figure 7.6: The discharge current P2P oscillation strength (normalized by the mean discharge current) for the Inner channel operating at 500 V, $0.63 j_{ref}$ in the I and IMO configurations, compared against results from firing only the Inner channel with the IMO magnetic field (at the same field strength as the other two conditions).

channel coupling mechanism and the cathode-oscillation coupling as mechanisms.

This is only a single test point, but the implications are still clear. As explained by Cusson [205], the magnetic field topology in a NHT such as the X3 can be tuned between single- and multi-channel operation such that the same plasma lens shape is present for all conditions. However, if there are slight differences in the plasma location in a given discharge channel between magnetic field configurations, it is possible that the plasma-wall interaction is playing a role in the discharge oscillations. We found at times during multi-channel operation that the Inner and Middle channels had areas of deposited carbon on the discharge channel walls that were glowing orange, indicating that plasma was impinging on the walls in a location different than during single-channel operation when these bands were established. At times of intense burn-off, the discharge current oscillations grew very large. Though we made an effort to minimize or burn off these spots before taking performance measurements, this change in plasma location could still be contributing to the change in

oscillation character through these wall effects.

If the increased oscillations were solely due to a transient wall effect (e.g., wear-band burn off), it is expected that the oscillations would settle over time. Near the end of the test campaign at GRC we operated the X3 in a small number of longer-duration firings in the IMO configuration: two at 50-kW total discharge power (400 V, $0.63 j_{ref}$) and one at 80-kW total discharge power (400 V, $1.0 j_{ref}$). During these firings the X3 was operated at fixed settings for 2–3 hours. An example of the change in P2P strength with time during one of the 50-kW runs is plotted in Figure 7.7. The data demonstrate that there is a certain level of reduction of oscillations for the Inner channel over the first 90 minutes of thruster operation before they level off to a value around 0.85 of the mean discharge current. However, the seemingly steady-state value of the oscillations for the Inner channel is still over double that of the Middle and Outer channels, which for this particular condition are each at around 0.4 of their respective mean discharge currents. This indicates that although there is likely a certain amount of plasma-wall interaction and transient wear-band burn-off artificially inflating the oscillation values in multi-channel operation there is still another mechanism causing these oscillations to be larger that does not go away with time. For reference, the Inner operating in the I configuration at this discharge voltage and current density exhibited a I_{P2P}/I_d value of 0.293, well below the steady-state value seen here.

There are a number of possibilities for mechanisms causing this phenomenon. For instance, it is possible that the different field ratios for the Inner and Middle channels during multi-channel operation resulted in a change in I-B characteristic [48], and that a change in field strength on those channels would have reduced the oscillations. The need is apparent for further work characterizing these oscillations with magnetic field strength and operating condition, as well as further understanding of the shape and location of the plasma inside the discharge channel during various X3 operating conditions. Ultimately, although the Inner and Middle channels experienced elevated multi-channel oscillations, none of the X3's P2P oscillations fell outside of the range typically seen by Hall thrusters [178].

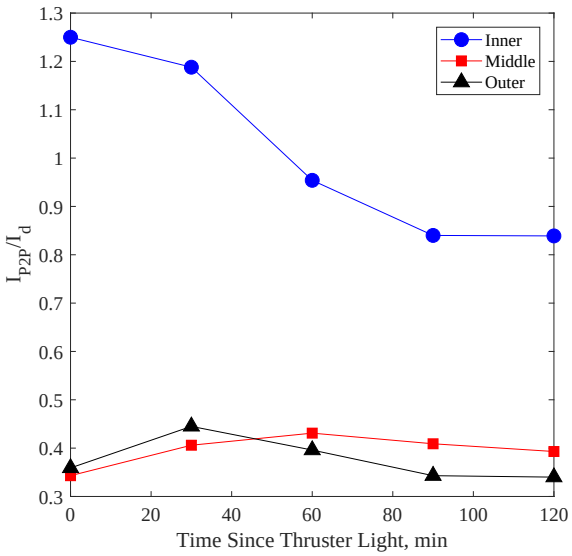

Figure 7.7: The discharge current P2P oscillation strength (normalized by the mean discharge current) for all three channels as a function of time during a long-duration firing at 50-kW total discharge power.

7.3.2.3 Cathode Coupling

One more potential mechanism through which the channels of an NHT may affect each other or differ in operation at a given condition is through coupling to the cathode. However, the X3 data suggest that cathode coupling was consistent across all conditions and power levels tested, indicating that cathode coupling is not a function of channel combination or power level. One possible explanation for this is related to the discussion of multi-channel performance above. The increased neutral flow near the cathode (via the external cathode injectors) may be responsible for enhancing cathode coupling in a manner similar to that described by Beal [213]. Because of the way the channels are scaled, the larger the channel, the more auxiliary flow needed during operation to provide the proper TCFF. This increasing neutral flow field in front of the thruster may be enhancing the cathode coupling for larger channels and overcoming potential negative effects such as distance between channel and cathode. Further work, including firing the larger channels without auxiliary flow, may help identify this mechanism.

7.4 Summary

We successfully measured the performance of the X3 for a range of conditions spanning total power levels from 5 to 102 kW. These conditions consisted of discharge voltages from 300 to 500 V and current densities that were 0.63, 1.00, and 1.26 of a reference value. The seven channel combinations of the thruster were throttled across this range of settings. For each test point, we directly measured thrust using a high-power inverted-pendulum thrust stand, and from those thrust measurements and thrust telemetry, we calculated specific impulse and efficiency values. We collected measurements of the discharge current oscillations at each point to assess thruster stability.

Our results demonstrated that a three-channel 100-kW class NHT can offer comparable or even improved performance over high-power single-channel thrusters. The X3 demonstrated total efficiencies ranging from 0.54–0.67 and total specific impulses from 1800–2650 seconds, experiencing the peak efficiency at 500 V discharge voltage. Additionally, the results indicate that none of the available channel combinations were significantly over- or under-performing the rest, demonstrating comparable specific impulse and efficiency at a fixed discharge voltage and current density. The sum of the single-channel thrusts generally was not significantly different than the three channels operating together for a given discharge voltage and current density, a surprising result that differed from previous NHT testing.

Discharge current oscillation measurements showed that the oscillation strength did not vary significantly between single- and multi-channel operation for the Outer channel. However, the Middle and Inner channels both showed increased oscillation strength during multi-channel operation, with the Middle increasing from approximately 25% in single-channel operation to approximately 50% in multi-channel operation and the Inner increasing similarly from approximately 25% to approximately 75%. We performed a limited exploration of this phenomenon and found that the Inner channel exhibited elevated oscillation strength while operating on its own but with the IMO magnetic field settings (that is, all six electromagnets on as they would be for IMO thruster operation). Though these results are not conclusive

and a more extensive study is warranted, they provide insight into the mechanisms that may be causing this behavior.

The X3 demonstrated new levels of Hall thruster power (102 kW), thrust (5.4 N), and discharge current (247 A). These results demonstrate the capability of Hall thrusters and NHTs for cargo and crew transport applications and show that the X3 represents a significant milestone on the roadmap to Mars.

CHAPTER 8

High-Power Diagnostics

"The light emanating from a space shuttle launch is different in color, quality, and intensity from any other kind of light. Photographs and videos can only approximate it, can only serve as a souvenir to the odd sensation, the combination of beauty and near-painfulness of that specific brightness in the sky."

– Margaret Lazarus Dean [217]

8.1 Introduction

This chapter presents the detailed diagnostic data from testing at GRC. In Section 8.2 we discuss the average efficiency results of the X3 and trends with discharge voltage and current density. We also compare these results and trends to those from the NASA-300M 20-kW Hall thruster. Next, we study in more detail how each channel combination of the X3 operates in Section 8.3. This includes accounting for cross-channel ingestion effects as well as facility ingestion effects. Finally, we provide detailed high-speed diagnostics results in Section 8.5. These include typical high-speed camera results, atypical results, and a discussion of the implications of these results. We then demonstrate that the high-speed discharge current analysis captures the same results, validating the high-speed camera analysis.

8.2 Average Phenomenological Efficiency Results

The far-field plasma diagnostics described in Chapter 4 were used to calculate the four limited-diagnostics utilization efficiencies as described in Chapter 2. These efficiencies provide insight into overall efficiency of the thruster via the plasma processes at work within the discharge. For the X3, these quantities can also be used to further study the way the various channel combinations of the thruster differ (or not) in operation.

8.2.1 On Not Correcting for Facility Ingestion and NHT Effects

As discussed above in Chapter 2, it is possible to correct the anode efficiency for facility ingestion. However, in practice this is more of an academic exercise than one that uncovers hidden phenomena in thruster operation. This is because an increase in anode flow rate produces an increase and facility background pressure, and thus in ingestion, meaning that the ingestion stays roughly the same fraction of the anode flow rate. As such, it is expected that $\eta_{facility}$ will be roughly constant across all conditions, and thus represent a linear offset in the data. Indeed, in calculating $\eta_{facility}$ using the background flow model, we find that the average value across all test conditions during the high-power campaign was 0.013 with a standard deviation of 0.001, indicating that the value did not vary much. The slight variation in values was likely due to the exact facility conditions for the day. Additionally, because this correction is not typically undertaken in the literature, comparison with other thrusters will be more complicated with this facility ingestion correction applied. As such, we simply note the η_{fac} value but do not correct anode efficiency η_a for the following sections studying trends with discharge voltage and current density, and for the comparisons to the NASA-300M thruster.

We also choose not to explicitly correct for NHT ingestion effects for these average comparisons. The thruster was not operated in the perfectly controlled manner discussed in Chapter 2, meaning that background pressure-based effects on thruster oscillation mode and

acceleration region location are likely captured artificially in an η_{NHT} calculation. Additionally, as we discussed in Chapter 2, this cross-channel ingestion will be present on-orbit and thus is a fundamental component of device operation and should not be removed or corrected for. However, in the section below comparing the performance of the various channel combinations, it is necessary to at least extract and study this factor to understand its role in thruster operation. Thus, for the first section of this chapter we will not account for these factors, but in the second section we will.

8.2.2 Trends with Discharge Voltage

We begin the phenomenological efficiency analysis with a study of how the efficiencies trend with discharge voltage. To do so, the values are averaged over all discharge currents and all channel combinations for a given discharge voltage. That is,

$$\overline{\eta_x}\left(V_d\right) = \sum_j \sum_C \eta_{x,jC}\left(V_d\right), \tag{8.1}$$

where $\overline{\eta_x}$ is the average utilization efficiency x (e.g., anode, charge, voltage, etc.) for discharge voltage V_d, j is the current density (e.g., $0.63j_{ref}$, etc.), and C is the channel combination (e.g., I, M, etc.). This averaging technique obscures trends with discharge current density and channel combination, but these will be studied separately in detail below.

Figure 8.1 presents these averaged values versus discharge voltage. As can be seen, the anode efficiency increases with increasing discharge voltage as previously shown in Chapter 7. The utilization efficiency with the lowest magnitude is η_{bd}, which matches results seen with other thrusters [79], and it trends upwards with increasing discharge voltage, contributing to the increase in anode efficiency. Divergence-weighted mass utilization efficiency, voltage utilization efficiency, and charge utilization efficiency are all greater than 0.90 for all conditions. Voltage utilization efficiency increases with increasing discharge voltage, an expected result because the loss voltage in a Hall thruster is typically invariant with discharge voltage

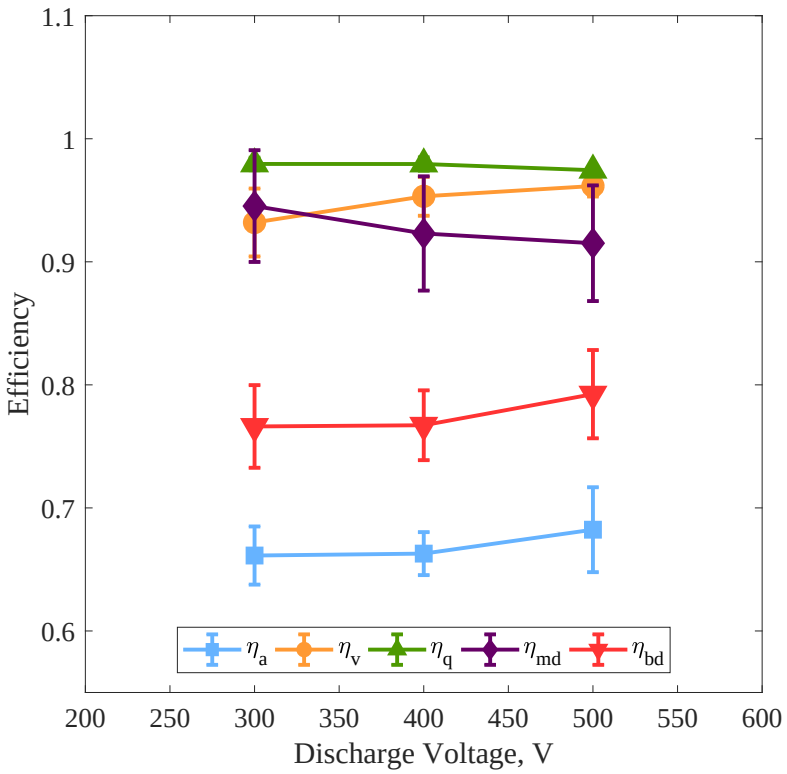

Figure 8.1: The average utilization efficiencies versus discharge voltage for the X3. Error bars represent the standard deviation of the averaged data, which for all cases was smaller than the statistical uncertainty (which was omitted for clarity). Error bars on η_q are smaller than the markers.

and thus represents a smaller percentage at higher values of V_d. Charge utilization efficiency decreases slightly with increasing V_d, indicating an increase in the population of multiply-charged ions at higher discharge voltages. This is a trend seen with the NASA-400M [76] and the NASA-173M [79] thrusters. Finally, η_{md} decreases with increasing discharge voltage, but overall the increase in η_v and η_{bd} is greater than this decrease and anode efficiency still increases with V_d.

8.2.2.1 Comparison with NASA-300M

To continue with the comparisons to other high-power Hall thrusters that were made for the performance metrics in Chapter 7, we compare the X3's phenomenological efficiency values and trends to those of the NASA-300M, which is the only of the NASA high-power Hall thrusters to have its plume thoroughly studied in this manner. These results are presented by Huang in Reference [108]. The NASA-300M is also the most state of the art of the series of thrusters, incorporating all of the lessons learned from the previous thrusters (lessons applied to the design of the X3 as well). In the NASA-300M study the plume was thoroughly interrogated using a full suite of plasma diagnostics, including not only an RPA, LP, and WFS like those used on the X3 but also a radially-swept Faraday probe. This provided the researchers with information on the $I_B \cos\theta$ term of the efficiency analysis which was missing from X3 analysis, and thus they were able to calculate all five utilization efficiencies separately. To make the comparison to the X3 data, we took the published data and calculated values for η_{md} and η_{bd} in the following manner:

$$\eta_{md,300M} = \eta_{m,300M}\sqrt{\eta_{d,300M}} \tag{8.2}$$

and

$$\eta_{bd,300M} = \eta_{b,300M}\sqrt{\eta_{d,300M}}, \tag{8.3}$$

where $\eta_{b,300M}$, $\eta_{m,300M}$ and $\eta_{d,300M}$ are the current utilization, mass utilization, and divergence utilization reported for the NASA-300M. We then averaged these values for discharge voltage in a similar fashion as the X3 data.

First we present anode efficiency, which indicates how the thrusters are performing relative to each other. As can be seen in Figure 8.2, the anode efficiencies derived from thrust measurements for both thrusters track closely together with increasing discharge voltage. The X3 performed approximately 0.01 higher at 300 V and 500 V, while the NASA-300M

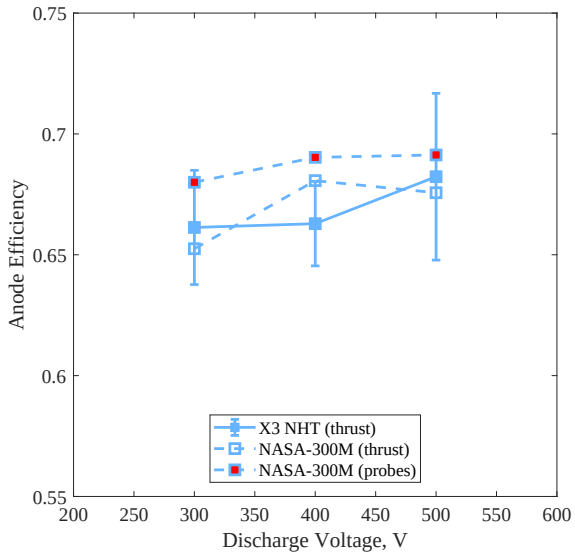

Figure 8.2: Anode efficiency as a function of discharge voltage for the X3 and the NASA-300M. Anode efficiency was calculated for the NASA-300M using both the thrust stand (like the X3) and with the plasma probes. Error bars on the X3 data reflect the standard deviation of the averaged points.

performed approximately 0.02 higher at 400 V. For both thrusters, anode efficiency increased for increasing discharge voltage. Also plotted in Figure 8.2 are the probe-derived anode efficiencies for the NASA-300M. As can be seen, these values generally show the same trend and slope with discharge voltage but are on average 0.03 higher than the thrust-derived values. Within the uncertainty of the X3 measurements (± 0.03), the thrust-derived anode efficiency values for both thrusters are the same. This can be taken as a confirmation that the X3 is operating as designed. Based on the scaling laws developed by Manzella in Reference [11], we subsequently expect that the various utilization efficiencies should also match between the thrusters because those scaling techniques are intended to ensure that the plasma (and thus its internal processes) is kept similar as power is increased. Because the anode efficiency values are very similar, any deviations in the various utilization efficiencies will have to cancel out.

We start by studying charge utilization efficiency, which we plot for both thrusters in

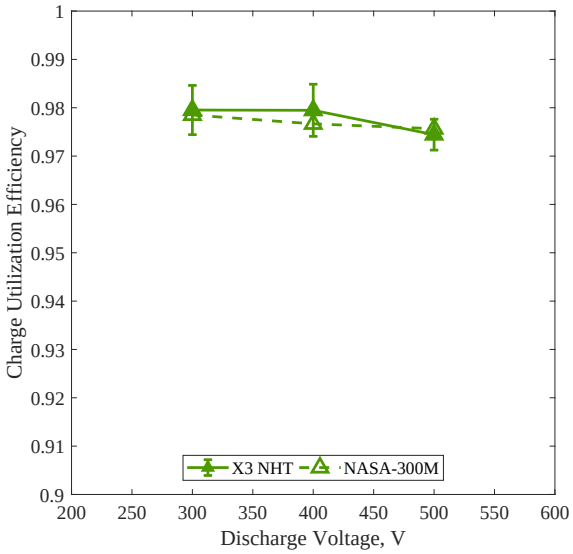

Figure 8.3: Charge utilization efficiency for both the X3 and the NASA-300M as a function of discharge voltage. Error bars on the X3 data reflect the standard deviation of the averaged points.

Figure 8.3. The results for both thrusters are nearly identical in magnitude and feature a similar negative slope with increasing discharge voltage. These values also match historical ones such as those reported by Hofer on the NASA-173M [79]. This indicates that the relative populations of charge states of xenon ions is similar between the thrusters at a fixed discharge voltage. The negative slope indicates that higher discharge voltages are producing more multiply-charged xenon ions. Overall, this result indicates that the X3 is not producing significantly different populations of multiply-charged ions than the NASA-300M.

Next, Figure 8.4 shows the voltage utilization efficiency as a function of discharge voltage for both the X3 and the NASA-300M. Here, we find that the X3 marginally outperformed the NASA-300M by 0.005–0.01 at each discharge voltage. However, the uncertainty of the X3 values were between ±0.01 and ±0.02, making the values statistically the same to the NASA-300M values. The trend with increasing discharge voltage is also the same, with voltage utilization efficiency increasing for both thrusters. This is to be expected, as discussed above. The similarity between these data for the two thrusters indicates that they are accelerating

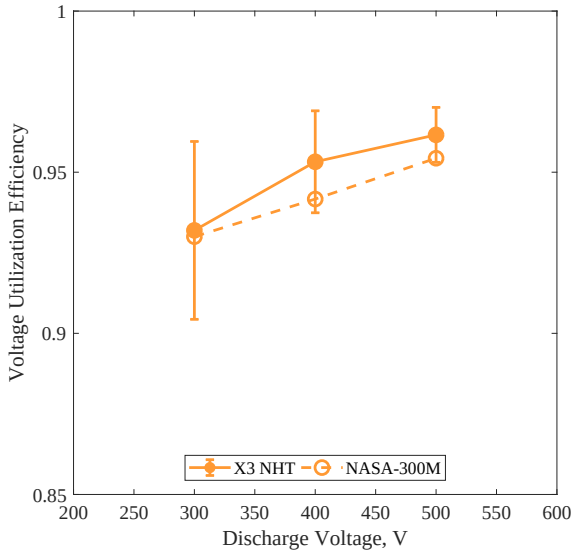

Figure 8.4: Voltage utilization efficiency for both the X3 and NASA-300M as a function of discharge voltage. Error bars on the X3 data reflect the standard deviation of the averaged points.

ions (as well as coupling to their respective cathodes) with similar efficacy.

Divergence-weighted current utilization efficiency is presented for both thrusters in Figure 8.5. Here we find that once again the X3 and NASA-300M data track closely. Because of the nature of the calculation for η_{bd} for the X3, it features a rather large uncertainty of ±0.11, but the mean values are always within 0.02 between the two thrusters for a given discharge voltage. Both thrusters show an increase with increasing discharge voltage as well. This indicates both that the thrusters are converting discharge current to beam current equally efficiently and that the method of calculating η_{bd} for the X3 with the limited plasma diagnostic suite is valid and yields informative information about thruster behavior.

Finally, Figure 8.6 presents divergence-weighted mass utilization efficiency versus discharge voltage for both thrusters. Here we find a more striking difference between the two thrusters. Though the two thrusters show a similar downward trend in η_{md} with increasing discharge voltage, the X3's average efficiencies are consistently between 0.035 and 0.05 lower than those of the NASA-300M. This difference is not technically statistically significant, as

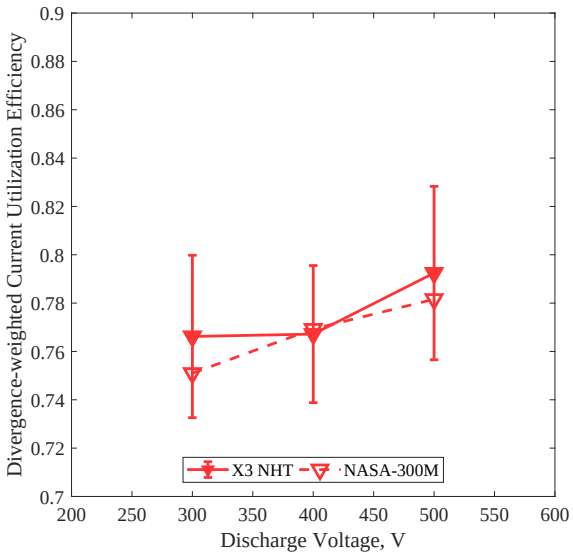

Figure 8.5: Divergence-weighted current utilization efficiency for both the X3 and NASA-300M as a function of discharge voltage. Error bars on the X3 data reflect the standard deviation of the averaged points.

the η_{md} values for the X3 also carry an uncertainty of ±0.11, but the difference is still significantly larger than those observed on any other efficiency. This result initially seems at odds with the similarity in all other efficiencies between the two thrusters. However, if we return to the anode efficiencies in Figure 8.2, we are reminded that the probe-calculated anode efficiency for the NASA-300M was higher than the thrust-calculated value. The NASA-300M's larger η_{md} is counteracted by the X3's slightly larger η_v to result in probe-calculated anode efficiencies about 0.03 larger than the thrust-calculated ones. It is worth noting that the data for the NASA-300M produced surprisingly large values of η_m, at times slightly in excess of unity. The uncertainties on the values for the NASA-300M were ±0.04, the largest of any parameter calculated during that study. Thus, it is possible that with the large uncertainties on the values from both thrusters that there is not actually a difference between the X3 and the NASA-300M in how efficiently xenon atoms are converted to ions.

To study this discrepancy further, the component of η_{md} that is dependent on the WFS probe data, α_m, is plotted in Figure 8.7 as a function of discharge voltage. The results match

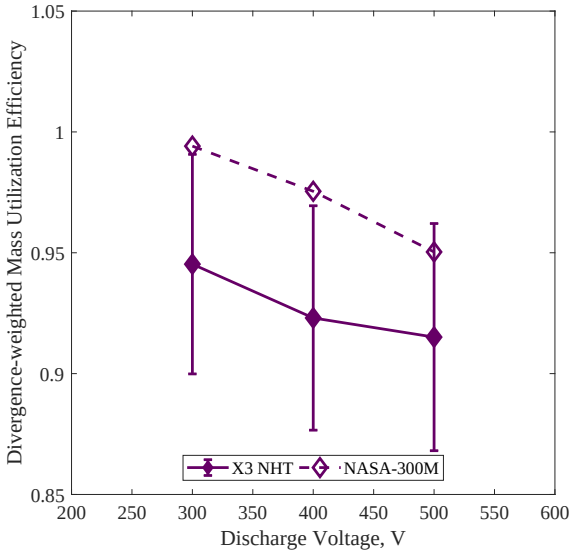

Figure 8.6: Divergence-weighted mass utilization efficiency for both the X3 and NASA-300M as a function of discharge voltage. Error bars on the X3 data reflect the standard deviation of the averaged points.

closely in trend and value for both thrusters, a result that is unsurprising based on the similarities in η_q, which is calculated from WFS data as well. This indicates that, if there is a difference between the two thrusters, the cause is likely a difference in the ratio of I_b to $\dot{m}_a$, as shown in Equation 2.19. As shown in Figure 8.5, however, the divergence-weighted current utilization efficiency is very similar between the two thrusters. This indicates that they are converting discharge current to beam current equally well, but that the X3 perhaps requires slightly more anode flow rate to produce a given discharge current (and thus are operating at different $I_d/\dot{m}_a$ ratios, which unfortunately cannot be verified due to the lack of thruster telemetry in Reference [108]). However, the mis-match between the NASA-300M probe-calculated and thrust-calculated anode values and the un-physical nature of some of the η_m values suggests that those values are potentially suspect. This, combined with the large uncertainty in the values for both thrusters, leads us to be unable to draw any significant conclusions about the difference in η_{md} between the X3 and the NASA-300M.

A final comparison to make to the NASA-300M is that of current fractions of charged

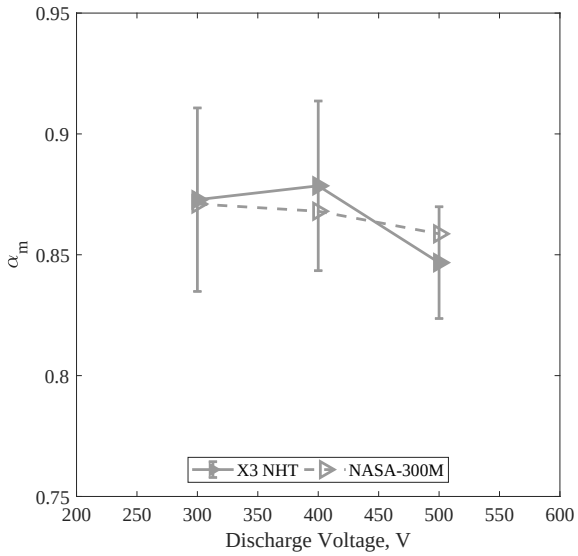

Figure 8.7: The α_m component of the divergence-weighted mass utilization for both the X3 and NASA-300M as a function of discharge voltage. Error bars on the X3 data reflect the standard deviation of the averaged points.

species. The similarity in η_q between the two thrusters indicates that the relative populations should be similar as well, but it is still insightful to study trends in relative populations with V_d. These results for the X3 and the NASA-300M are plotted in Figure 8.8. Work such as Hofer's in Reference [79] has shown that in general Hall thrusters experience an increase in multiply-charged species populations with increasing V_d. However, for the X3 (solid lines and markers in the figure) we find that the population of Xe^+ actually increases from 0.755 at 300 V to 0.770 at 400 V before dropping to 0.712 at 500 V. This trend is mirrored in the Xe^{2+} current fraction, which decreases from 0.218 at 300 V to 0.195 at 400 V before increasing to 0.238 at 500 V. The Xe^{3+} population, which had an average current fraction less than 0.05 for all V_d, rose monotonically with discharge voltage from 0.021 at 300 V to 0.038 at 500 V. The Xe^{4+} current fraction was typically less than 0.01 and is thus not reported here.

We find that the NASA-300M produced magnitudes and trends very similar to the X3, as shown with the dashed lines and open markers in Figure 8.8. The trends in the current fractions of all three species were very similar, with a rise in Xe^+ at 400 V and a decrease

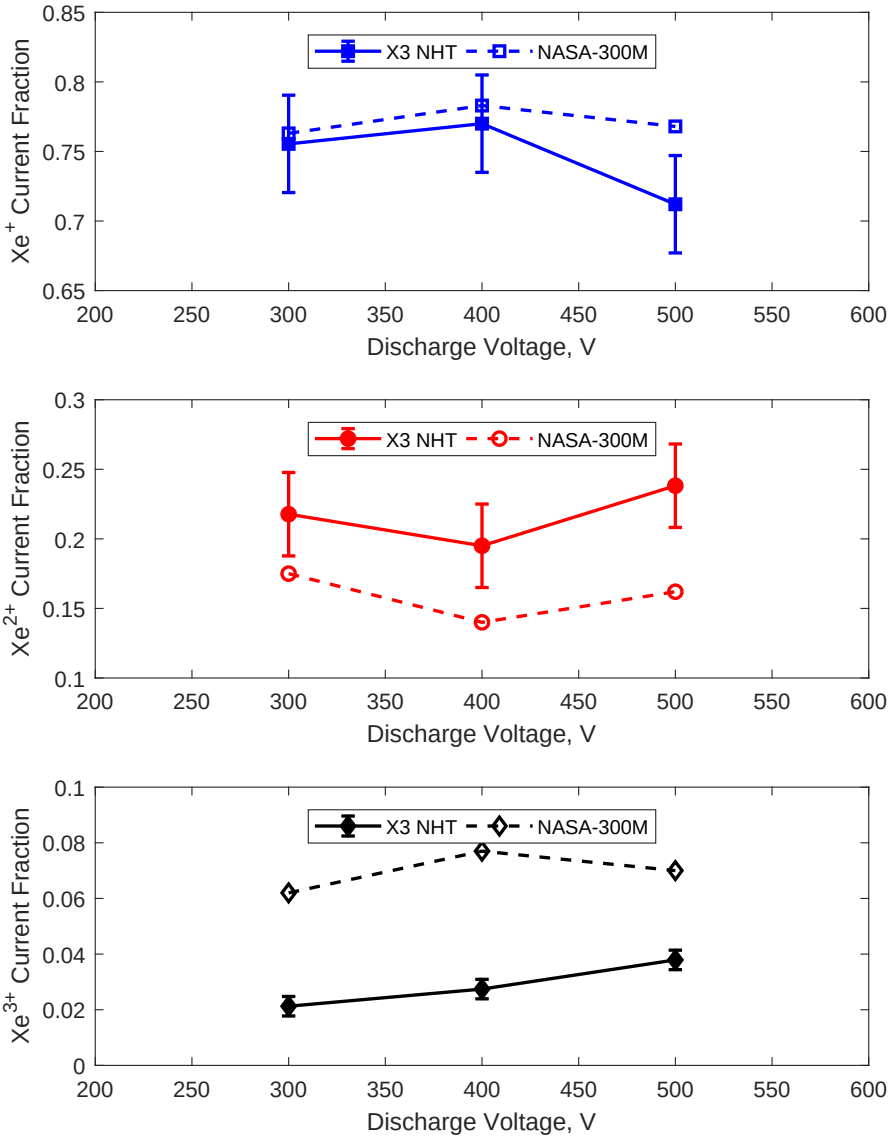

Figure 8.8: Average current fractions of Xe^{+}, Xe^{2+}, and Xe^{3+} for the X3 and the NASA-300M at three different discharge voltages. Error bars represent the statistical uncertainty in the current fractions for the X3 data. The NASA-300M results had very similar uncertainties, which are suppressed on the plot for clarity.

at 500 V, corresponding to a decrease in Xe^{2+} at 400 V and an increase at 500 V. The X3 and the NASA-300M produced nearly identical current fractions of Xe^{+} at 300 and 400 V, though at 500 V the X3 decreased farther than the NASA-300M. However, the magnitudes of the Xe^{2+} and Xe^{3+} populations are different between the two thrusters: the X3 is producing more Xe^{2+} and less Xe^{3+} than the NASA-300M. The X3's average current fractions for Xe^{2+} are 0.043 to 0.078 larger in magnitude than the NASA-300M, which correspond to similar differences in the opposite direction for Xe^{3+}. This, in turn, explains the X3's slightly higher magnitudes of η_q as shown in Figure 8.3: higher charge states are responsible for more losses, so by producing more ions in a lower charge state, the X3 is operating more efficiently than the NASA-300M. However, as the difference in η_q demonstrates, this is responsible for a very small difference in performance between the two thrusters.

Aside from some reservations about divergence-weighted mass utilization, the X3 has been found to have similar magnitudes and trends for all utilization efficiencies. This suggests that the X3 is performing as it was designed to, and that the Manzella-type scaling laws serve to keep the plasma properties and processes similar for thrusters ranging in power from 20 kW to 100 kW and sizes ranging from approximately 300 mm to 800 mm outside diameter. This also lends initial credence to the idea that the X3 does not suffer from any major changes in plasma properties and processes between single- and multi-channel operation, since the X3 average values here contained both types of operation and were not significantly different than the single-channel NASA-300M. Before analyzing that question more rigorously we investigate how the various efficiency values trended with changing discharge current density.

8.2.3 Trends with Current Density

When developing the NASA-457Mv1 50-kW class Hall thruster, Manzella created a simplified model of Hall thruster operation to study the effects of changes in different thruster design parameters, including channel length, magnetic field strength, and propellant density [11]. This model suggested that increased propellant density should produce increased efficiency

by moving the ionization region upstream of the acceleration region within the discharge channel. This work highlighted the balance involved in optimizing the propellant density in a Hall thruster. Too low of a propellant density decreases ionization: Manzella's work suggested that at too low a flow rate, 20% of the propellant was not ionized at all. But at a certain propellant density, all ions are being created upstream of the Hall region, and further increases in propellant density increase collisional losses without further improvements in ionization and acceleration. Additionally, increased collisions increase electron mobility to the anode (and thus electron current), providing another loss mechanism. Ultimately, this balance suggests that with increasing propellant (and thus current) density, there should be an increase in either η_v or η_m (or both). Here, the η_m changes would be reflected in η_{md}.

The X3 was throttled over a total of three different current densities: 0.63, 1.0, and 1.26 the middle reference value j_{ref}. For both $0.63j_{ref}$ and $1.0j_{ref}$, full or nearly-full sets of points spanning all channel combinations were collected for all three discharge voltages. To study the effect of increasing the current density, we averaged the 300 V, 400 V, and 500 V data together across all test points for each current density as such:

$$\overline{\eta_x}(j) = \sum_V \sum_C \eta_{x,VC}(j), \tag{8.4}$$

where x, j, and C hold their same definitions as in Equation 8.1 but where we are summing over discharge voltages V for fixed current density j. Summing in this manner will average out trends with V_d as discussed above. However, because these trends were similar for both current densities, comparisons for these terms are still valid. Summing in this manner over many conditions provides more statistically-meaningful values than taking only a particular set. There are too few conditions at $1.26j_{ref}$ to draw meaningful conclusions so these points are excluded from the present study.

Figure 8.9 shows the average of anode efficiency and each utilization efficiency for both $0.63j_{ref}$ and $1.0j_{ref}$. We find that anode efficiency, on the far left of the figure, increases by about 0.02 for the higher current density. Charge, voltage, and divergence-weighted

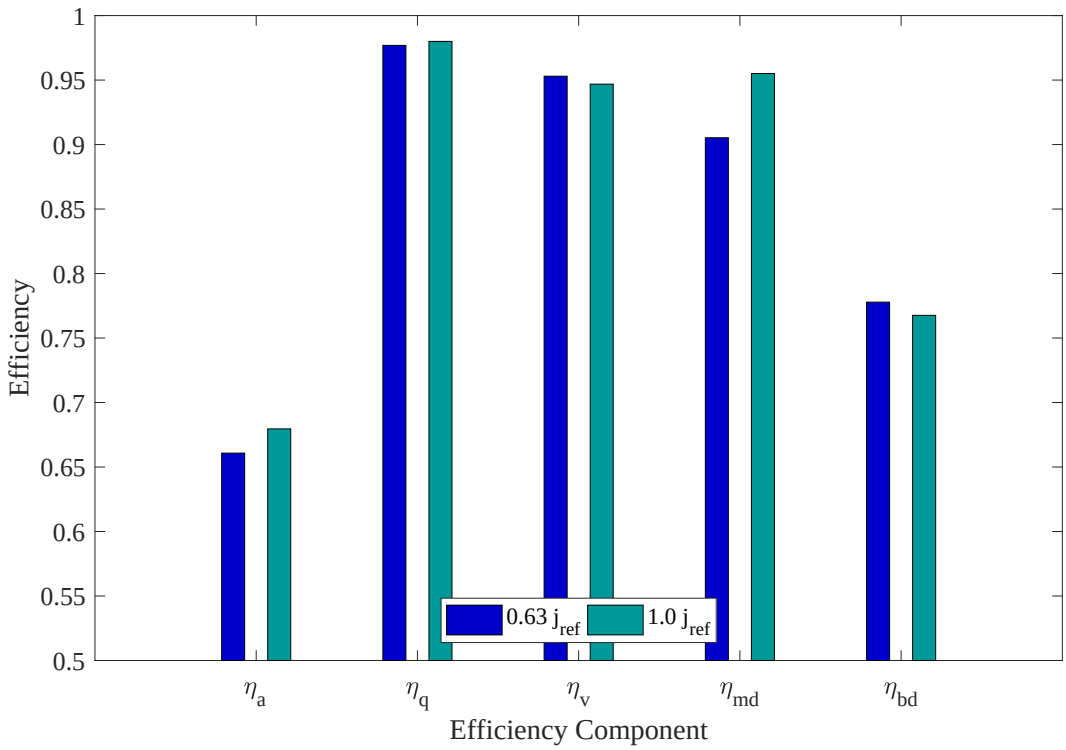

Figure 8.9: Anode and utilization efficiencies for two current densities, averaged across all discharge voltages.

current utilization all change by less than 0.007 between the two current densities. However, divergence-weighted mass utilization shows an increase of 0.04. Although this difference is within the statistical uncertainty of the values, the closeness of the other parameters indicates that the difference is nonetheless important. Furthermore, the lack of change in divergence-weighted current utilization efficiency suggests that the divergence itself is not changing.

If we re-write our expression for η_{md} in Equation 2.19 as such:

$$\eta_{md} = \left(\frac{m_{Xe}}{e}\right)\left(\frac{I_b \cos\theta}{I_d}\right)\left(\frac{I_d}{\dot{m}_a}\right)\alpha_m, \tag{8.5}$$

where we have substituted discharge current I_d into the expression by means of the divergence-weighted current utilization, we can see a change in η_{md} for near-constant η_{bd}

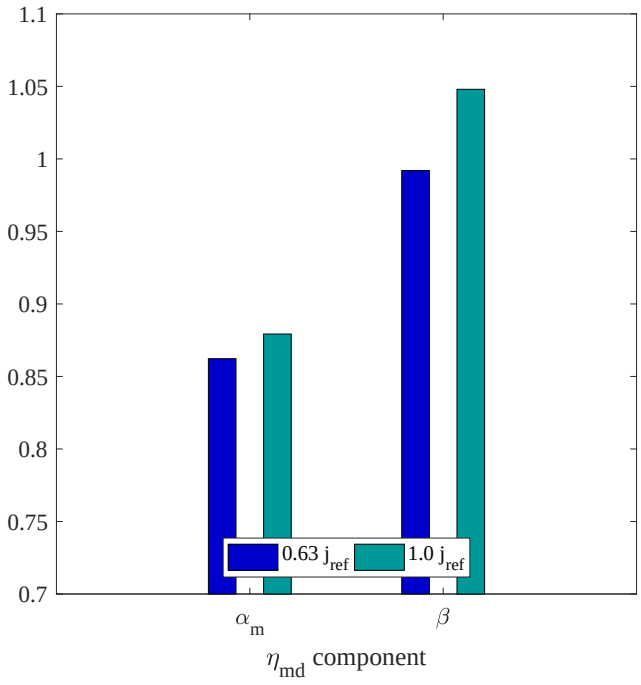

Figure 8.10: The α_m and β components of η_{md} for two current densities, averaged across all discharge voltages.

will be due to increases in α_m or $I_d/\dot{m}_a$, which we define here as β:

$$\beta = \frac{I_d}{\dot{m}_a}. \tag{8.6}$$

We can determine β directly from the thruster telemetry. Figure 8.10 presents α_m and β for the $0.63j_{ref}$ and $1.0j_{ref}$ current densities. The data demonstrate that both parameters increase with increased current density, but that the majority of the difference comes from the increase in β.

We can study this further by investigating the change in average charge species current fraction between the two current densities, as plotted in Figure 8.11. We find that the $1.0j_{ref}$ condition produces on average approximately 0.03 more singly-charged xenon than the $0.63j_{ref}$ condition. This corresponds to the slight improvement in η_q seen between the conditions, as well as the improvement in α_m.

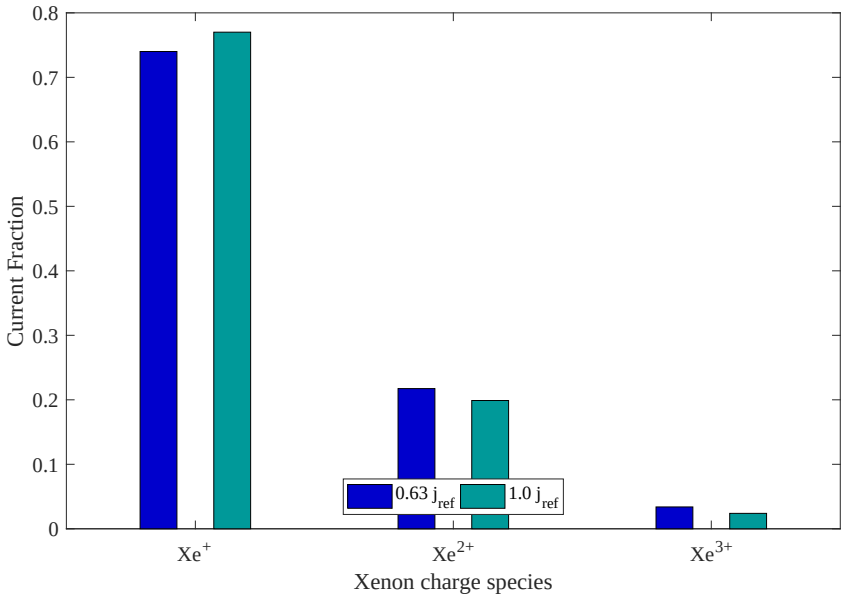

Figure 8.11: Average current fractions of Xe^{+}, Xe^{2+}, and Xe^{3+} for the X3 at two different current densities, averaged across all discharge voltages.

Ultimately, these results indicate that the X3's increased performance at higher current densities is due to an increase in how efficiently the thruster is converting neutral xenon propellant to ion current. This is the expected result based on the previous work discussed at the beginning of this section.

8.3 Comparison of Channel Combinations

Finally, we analyze more rigorously how the various channel combinations of the X3 operate with respect to each other. The previous two sections have demonstrated that, on average, the X3 is producing the expected trends and magnitudes of the various utilization efficiencies. However, it is still possible that there are differences between channels that on average cancel out to produce the results seen above. For instance, a certain channel could be operating more efficiently than average, and another less efficiently, such that the average value appears nominal. First, we compare the values without accounting for any potential NHT cross-

channel ingestion effects. Then we explicitly take these effects into account and study their impact on this analysis.

8.3.1 Channel Combination Efficiency Comparison

For this study, we averaged the corrected efficiency results for $0.63j_{ref}$ and $1.0j_{ref}$ conditions at 400 V discharge voltage. As was shown, there is variation in the values with current density. However, we elected to do the averaging here to provide a more complete set of data. The probe data for a number of test points was incomplete, due to issues such as dropped data points or electronics errors. Averaging together two sets at the same discharge voltage allows for a more complete picture of channel-to-channel variation while not complicating the comparison by averaging the expected change in utilization efficiencies with discharge voltage.

First, Figure 8.12 presents the average anode efficiency for 400 V operation for each of the seven channel combinations. The data seem to indicate that the IM condition produced an anode efficiency approximately 0.02 better than the other channel combinations. However, it should be noted that the IM condition was missing $0.63j_{ref}$ data and thus only represents the $1.0j_{ref}$ condition. As was discussed in the section above, the $1.0j_{ref}$ condition produced higher average η_a, meaning that this higher value is likely due to that. Ultimately, even with this higher value for the IM condition, η_a for each channel combination is the same to the uncertainty of the measurements and there do not appear to be any systematic trends. The similarity in η_a across channel conditions suggests that the various utilization efficiencies will be similar as well. To investigate this further, next we present each of the utilization efficiencies in turn and, when applicable, divide each into its relevant sub-components for continued study.

First, Figure 8.13 shows the variation of η_q between channel combinations. The data indicate small variation between combinations, although the previous sections have demonstrated that in general η_q varies little across conditions and relatively large changes in current

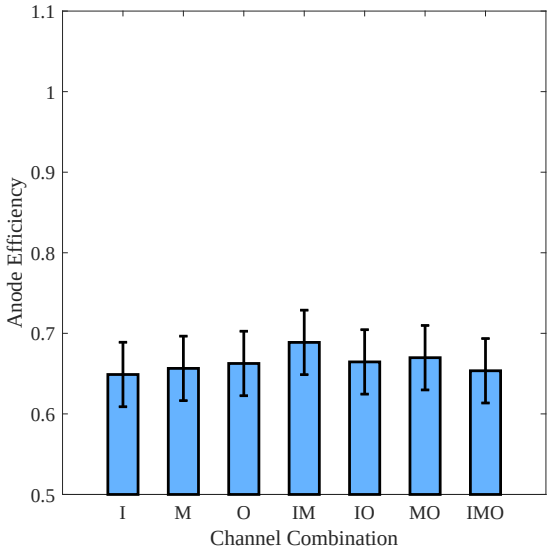

Figure 8.12: Anode efficiency plotted as a function of channel combination for 400 V operation.

fractions. These data suggest that the thruster is producing more singly-charged xenon in the MO and IMO configurations. Indeed, if we plot the average current fractions of Xe^{+}, Xe^{2+}, and Xe^{3+} as Figure 8.14 shows, we find that the MO and IMO configurations are producing more singly-charged xenon ions and fewer doubly-charged ions. The populations of triply-charged ions remains approximately the same and a small portion of the beam population.

One caveat to the WFS data that should be noted here concerns the charge exchange (CEX) correction applied to the data. We applied the full CEX correction as detailed by Shastry [105] and described above in Chapter 4 to all data. However, Shastry recommends that to minimize the CEX effects and the correction necessary, researchers should keep the "pz" factor less than or equal to 2, where p is the chamber background pressure in units of 10^{-5} Torr-Xe and z is the distance between the thruster and the probe in units of meters. As described previously, the X3 is a large, high-power thruster, which required us to place the WFS far downstream of the thruster to limit probe heating and probe/thruster interaction,

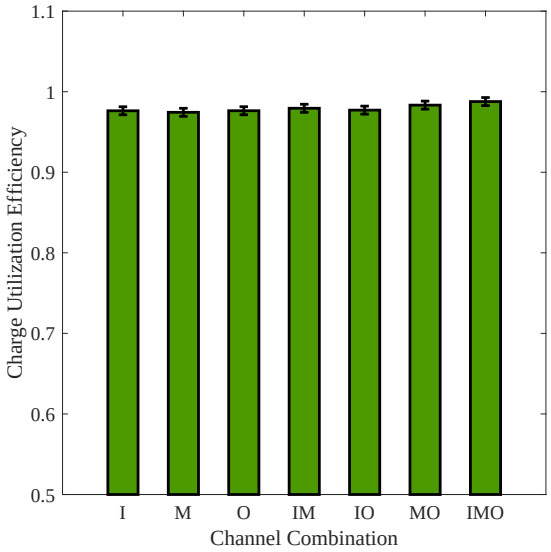

Figure 8.13: Charge utilization efficiency plotted as a function of channel combination for 400 V operation.

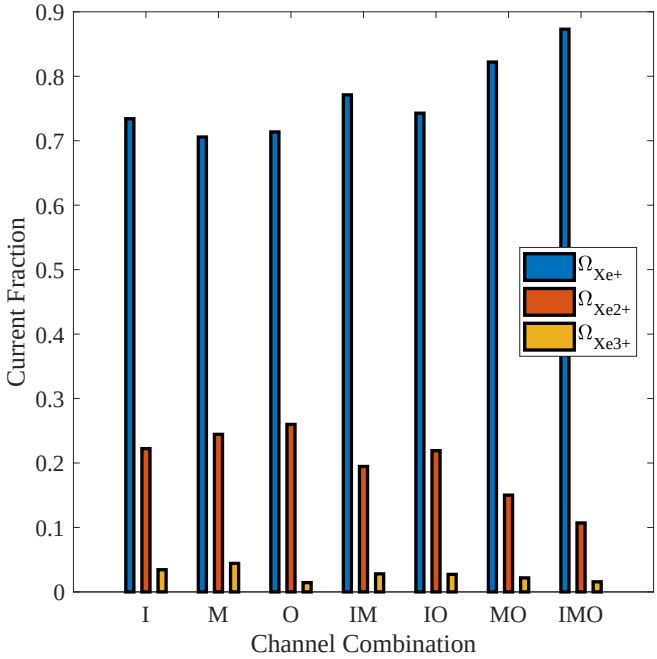

Figure 8.14: The average current fractions of Xe^{+}, Xe^{2+}, and Xe^{3+} at 400 V discharge voltage for each X3 channel combination.

Table 8.1: Average pz values for each channel combination for 400-V operation.

Channel Combination	I	M	O	IM	IO	MO	IMO
$\overline{pz}$	4.6	10.4	16.5	16.5	19.6	24.2	27.2

resulting in a z of 8.7 meters. The high-power nature of the X3 also results in higher background pressures, often in excess of $1{\times}10^{-5}$ Torr. As a consequence, the pz factor for these WFS traces was always in excess of the suggested limit of Shastry, sometimes an order of magnitude larger. The average pz values ($\overline{pz}$) for the 400-V test conditions are presented in Table 8.1. Although they exceed Shastry's recommendation, the effect of these elevated values of pz is unclear. It seems more likely that at higher background pressure, there will be more processes to deplete the beam of ions that would not be captured by the model (e.g., symmetric collisions that are ignored in the model), meaning the CEX correction would not correct the data enough. However, the higher values of Ω for the MO and IMO conditions (which feature the highest values of $\overline{pz}$ due to their large background pressures) if anything appear over-corrected compared to the other conditions (which would result in artificially high populations of Xe^+). Thus, the role of the CEX correction model on these results is unclear.

Next we turn to voltage utilization efficiency, as plotted in Figure 8.15. We find that η_v varies approximately 0.03 between conditions. The O, MO, and IMO configurations appear to produce slightly lower η_v than the other conditions. One potential explanation for this is that the Outer channel's lower efficiency occurs in both single- and multi-channel operation and thus brings down the average η_v for any condition that channel operates in. However, this would not explain the high η_v for the IO condition, which is on par with that for the I condition. Thus, is is also possible that this is simply a product of magnetic field settings; further optimization of the field, especially for the multi-channel conditions as described in Chapter 7, may bring these values closer together.

Another possible explanation for the variation in η_v seen here is variations in either V_{cg}, V_p, or both. The data for these conditions are plotted in Figure 8.16. We find that there is

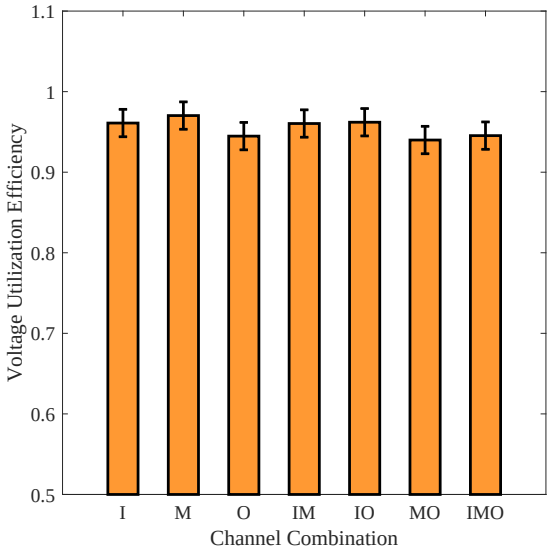

Figure 8.15: Voltage utilization efficiency plotted as a function of channel combination for 400 V operation.

variation of up to 4 V in V_{cg} among conditions and up to 5 V in V_p among conditions. The variation in V_{cg} is less systematic, but V_p for the O, MO, and IMO conditions is the highest of the set. However, the increase in V_p for these sets is not coupled with similar increases in the magnitude of V_{cg}, resulting in overall loss parameters that are no more than 1-2 V higher than the rest. One notable exception is the IM condition, which features small values of both V_p and V_{cg}, though this smaller loss parameter does not translate to a detectably larger η_v. The overall variation in the loss parameter is approximately 5 V, which is just over 1% of the discharge voltage. These results indicate that the mechanism by which η_v is lowered for certain conditions is related to the effectiveness at accelerating ions through the available potential, not in a difference in the available potential itself.

It is interesting to note that the channel combinations that included the Outer channel had larger plasma potentials V_p than the combinations that did not. The differences were relatively small and close to the measurement uncertainty, meaning that this may not be representative of a physical trend. However, modeling work done by Mikellides of the X3 in

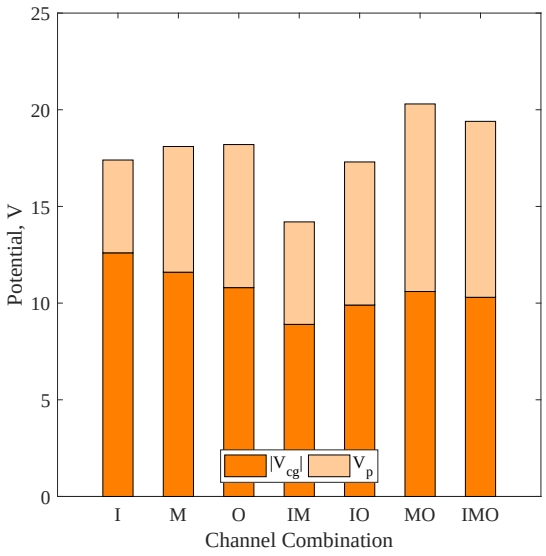

Figure 8.16: The average magnitudes of V_{cg} and V_p for each channel combination at 400 V discharge voltage.

JPL's two-dimensional axisymmetric Hall thruster simulation code Hall2De suggested that the Outer channel was likely to produce higher plasma potentials than the other two channels [218]. Although this difference is strongest near the thruster, at 3.5 thruster diameters downstream (the extent of the simulations), the simulations predicted a difference of over 10 V between the Outer channel's plasma potential and that of the other two channels. Here, at approximately 10 thruster diameters downstream, we typically measured plasma potentials a few volts higher for conditions that included the Outer channel. Mikellides's modeling suggested that this phenomenon was ultimately due to the magnetic field in the near-plume region of the Outer channel, which is distinctly stronger there than in the other two channels. This stronger magnetic field imposes a higher plasma resistivity, which in turn causes increased electron heating and a higher electron temperature. Through the electron pressure, these higher electron temperatures correlate with higher plasma potentials. Both the simulation and experimental results are preliminary, and near-field diagnostics are certainly necessary to verify this result. Nonetheless, the initial correlation between the

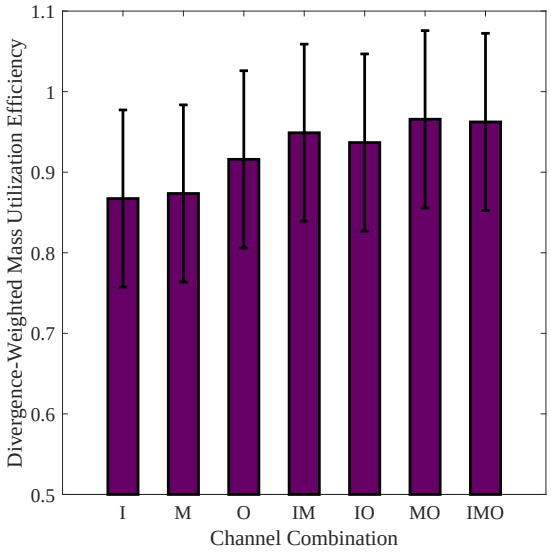

Figure 8.17: Divergence-weighted mass utilization efficiency plotted as a function of channel combination for 400 V operation.

simulations and the laboratory results is encouraging.

Next, Figure 8.17 presents the divergence-weighted mass utilization as a function of channel combination. Here, a difference among the combinations is immediately apparent, although it is within the uncertainty of the values. It appears that the single-channel conditions (in particular the I and M conditions) operate at values of η_{md} nearly 0.10 lower than the multi-channel conditions. This is a striking result considering the parity in η_a between conditions and in all average utilization efficiencies between the X3 and the NASA-300M. The clustering of the conditions is particularly stark: the I and M conditions are within 0.006 of each other and the IM, IO, MO, and IMO conditions are within 0.03 of each other, yet the two groups differ on average by more than 0.08. Only the O condition falls between the two clusters. This is an interesting result that may provide more insight into the difference between the two populations.

To study this phenomenon further, we can split η_{md} into the α_m and β components as we did in the previous section. These are plotted for each channel configuration in Figure 8.18.

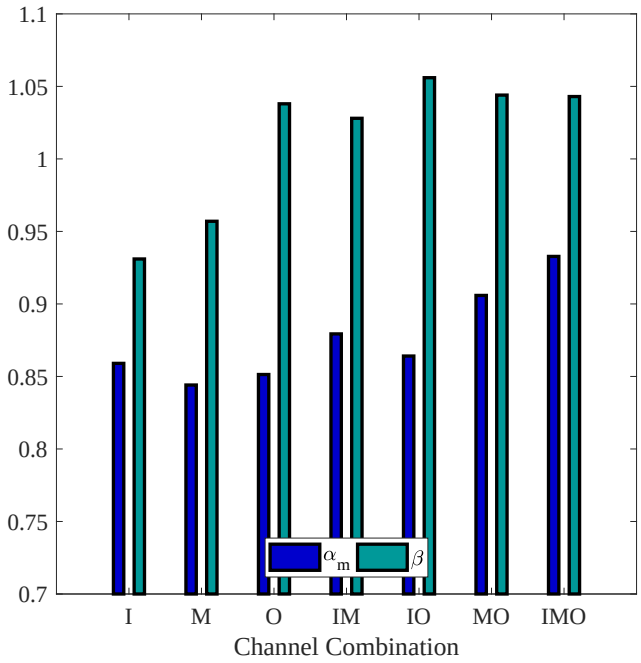

Figure 8.18: The α_m and β components of η_{md} for each channel configuration at 400 V discharge voltage.

The α_m is relatively consistent around 0.86 for all but the MO and IMO conditions, where it increases to 0.91 and 0.93 respectively. This trend reflects the trend seen in η_q and Ω in Figures 8.13 and 8.14, respectively, which indicated that the MO and IMO conditions were producing more singly-charged ions. The variation in α_m does not account for the trend seen in η_{md} since the multi-channel conditions are not as a whole higher than the single-channel ones. However, β shows a trend more like that seen in η_{md}, where the I and M conditions are at a noticeably lower value than the rest of the configurations. This indicates that the change is not related to multiply-charged species effects but is instead related directly to the conversion of anode mass flow rate to discharge current.

Finally, the divergence-weighted current utilization efficiency is presented in Figure 8.19. We find that the values are nearly identical at 0.79 for the I and M cases before dropping across the remaining configurations to 0.73 for the IMO condition. This decrease in η_{bd},

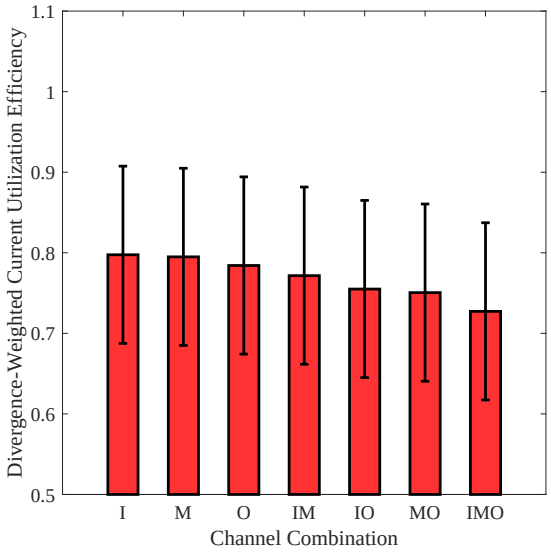

Figure 8.19: Divergence-weighted current utilization efficiency plotted as a function of channel combination for 400 V operation.

coupled with the trends in the other parameters, notably η_{md}, explains how the anode efficiency could be nearly invariant for all channel combinations. This result implies that either the divergence is getting larger for the multi-channel conditions or there is an increase in electron current. However, η_{md} is also weighted by $\cos\theta$ but trends in the opposite direction. This implies that this observed change in η_{bd} is not due to a change in divergence angle but instead is due to changes in electron current.

8.3.2 Calculation of NHT Effects

8.3.2.1 Procedure

The analysis of NHT cross-channel ingestion effects had three major steps: power-correction, facility ingestion calculation and removal, and calculation of η_{NHT}. The various operating conditions of the X3 were not always perfectly power-matched across a set (that is, target discharge currents and voltages were not always perfectly achieved). To facilitate more accurate calculations of η_{NHT}, we power-corrected the flow rates of the thruster. This process

was straightforward: we used the value of β (as defined in Equation 8.6) to predict the required flow rate at the target discharge current. Maximum corrections were 1.81 mg/s for the Inner channel, 3.35 mg/s for the Middle channel, and 2.84 mg/s for the Outer channel, with average corrections much smaller (between 0.29–0.50 mg/s for all three channels). These corrections allow us to compare channel conditions as though they operated at exactly the same discharge power. Though the value of β was shown to change between conditions, these power corrections are small enough that we assume β is constant for the correction range.

Using these power-corrected flow rates ($\dot{m}_{pc}$), we used Frieman's background flow model as described in Chapter 2 to calculate the neutral ingestion to the thruster at each condition. These ingestion flow rates $\dot{m}_{facility}$ were then subtracted from the $\dot{m}_{pc}$ values to isolate the "space-like" flow rates. Using these flow rates, we calculated the apparent ingestion due to NHT effects. At a given set of configurations at a given operating condition (discharge voltage and current density), we compared multi-channel flow rates to the sum of their constituent single-channel flow rates. Following Equation 2.28, we use the difference in these flow rates to calculate η_{NHT} for each condition.

If ingestion of vacuum facility neutrals were the only mechanism of ingestion, the two flow rates should match and η_{NHT} should be equal to unity. If the multi-channel flow rates are larger than the single-channel values it indicates that there is a loss mechanism at play during multi-channel operation that is requiring more propellant for the same discharge power. This would result in values of η_{NHT} less than unity. Finally, if the predicted result is true—that propellant is being ingested between channels in multi-channel operation—then η_{NHT} should be greater than unity.

8.3.2.2 Results

Indeed, we find that for all multi-channel conditions, η_{NHT} is greater than unity. Table 8.2 presents values for all multi-channel operating conditions tested. We found that η_{NHT} varied between a minimum of 1.016 and a maximum of 1.128. We found that the various

Table 8.2: Values of η_{NHT} for various X3 operating conditions.

V_d	j	IM	IO	MO	IMO
300 V	$0.63 j_{ref}$	1.016	1.057	1.070	1.095
300 V	$1.0 j_{ref}$	1.060	1.026	1.023	1.030
400 V	$0.63 j_{ref}$	–	1.052	1.048	1.062
400 V	$1.0 j_{ref}$	1.051	1.035	1.034	1.040
500 V	$0.63 j_{ref}$	1.070	1.031	1.128	1.056
500 V	$1.0 j_{ref}$	1.057	1.030	1.029	1.041
	Average	1.051	1.039	1.055	1.054
	Std. Dev.	0.020	0.013	0.039	0.023

multi-channel combinations demonstrated very similar average values of η_{NHT}, with the IO condition producing slightly smaller values than the others. This is an interesting result with physical meaning. Of all the multi-channel combinations, IO is the only one where the firing channels are not directly adjacent to each other. We expect that η_{NHT} and the cross-channel ingestion is a local process, and the significant increase in distance between the channels in the IO configuration (over double that of adjacent channels) likely results in less neutral ingestion. This is reflected in these values, and provides a hint at the difference in the γ_{ij} geometry factors for the channels that are farther apart.

The mass flow rates involved in the calculation of η_{NHT} are significantly larger than those ingested from the facility. For the 400 V, $1.0 j_{ref}$ condition, for example, the background flow model predicts the thruster is ingesting 1.2 mg/s for the IMO condition, yet the total difference in flow rates between that condition and the sum of the I, M, and O conditions (with their respective facility ingestion subtracted) is 7.5 mg/s. Ingestion of the facility neutrals accounts for only 16% of this flow rate deficit.

These results also have an important effect on the phenomenological efficiency analysis performed in the above section. Figure 8.17 showed that the divergence-weighted mass utilization efficiency of the multi-channel configurations was notably higher than for the single-channel configurations. When we correct these values by removing η_{NHT}, we find that η_{md} for these conditions is much more alike across all conditions. These results are presented

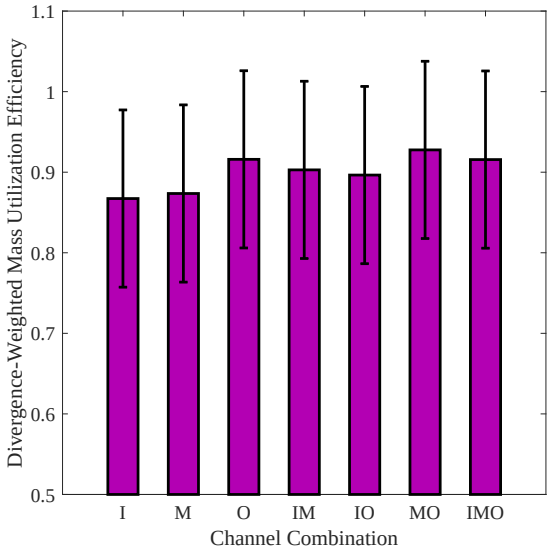

Figure 8.20: Divergence-weighted mass utilization efficiency, corrected for NHT ingestion effects, plotted as a function of channel combination for 400 V operation.

in Figure 8.20. Whereas the uncorrected values varied by between 0.08 and 0.10 between single- and multi-channel, with the correction, variation is between 0.04 and 0.06.

As discussed in Chapter 2, the η_{NHT} effect will be present in space because it is related to ingestion of neutrals emanating from the device itself, not those remaining due to finite facility pumping speed. However, calculation of η_a here incorporates these $\eta_{NHT} > 1$ values. That is, the η_a values for multi-channel operation are inflated by η_{NHT}. This must be kept in mind when comparing results with state-of-the-art SHTs.

Further work is needed to fully characterize this effect on the performance of the X3 and NHTs in general. Although the X3's performance appeared comparable to other high-power Hall thrusters, the present calculation indicates that a portion of the η_a value was likely due to NHT effects. This implies that the X3 may not be operating as efficiently as state of the art SHTs in multi-channel modes once these effects are controlled for. Our calculated η_{NHT} values indicate that the X3 and other NHTs should be capable of operation at anode efficiencies approaching 0.05 in excess of the state of the art when operating multiple

channels. At a minimum, further magnetic field optimization is needed for multi-channel operation of the X3, and it is entirely likely that a more detailed study of the fields would increase thruster performance.

8.3.3 Analysis Limitations

This experiment was not designed with measuring η_{NHT} as a primary goal, and as such, the calculations done above are notional at best. As discussed, there are likely a number of thruster behavior differences being captured in the η_{NHT} term, and it is unclear from this test how many of those effects will be present in space. A more detailed experiment, where parameters are properly controlled with the express purpose of studying this effect, will help differentiate between ground-test facility effects and those inherent to the NHT geometry. The values reported here should be taken as maximums, because they likely capture other effects as well. It is also likely that the plume structure and thus the ingestion mechanism may change with decreasing facility background pressure, further changing how the effects actually appear in space.

8.4 Fixed Faraday Probe Results

In addition to the diagnostics detailed above which contributed to the phenomenological efficiency analysis of the X3, the probe suite included the four fixed Faraday probes (FFPs) as described in Chapter 4. Here, we use the standard deviation between the four probes as a characterization of the relative beam asymmetry. Because this test represented the first high-power test of the X3, it was important to verify the approximate symmetry of the beam even if detailed plume maps could not be collected with a swept Faraday probe.

Figure 8.21 presents the standard deviation between the four FFPs for each condition, plotted as a function of discharge power. Single-channel conditions are plotted as blue circles and multi-channel conditions as red squares. Standard deviation varied between 0.7 and

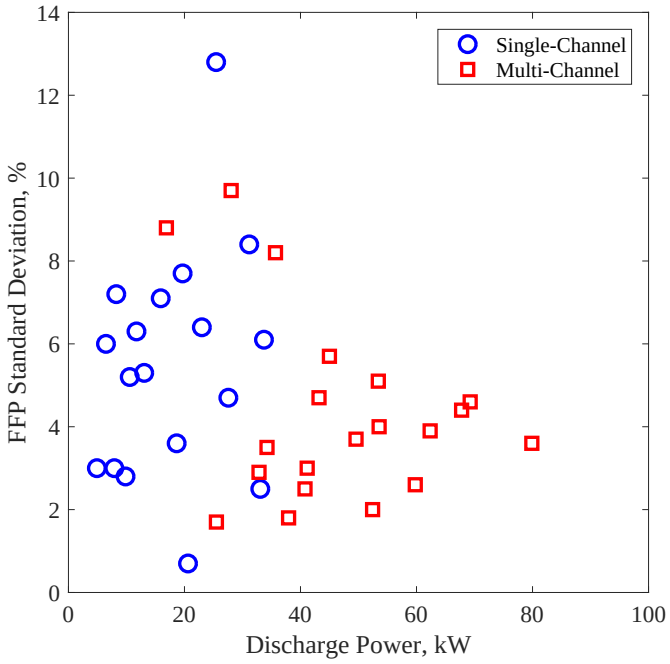

Figure 8.21: The standard deviation of the four fixed Faraday probe measurements at each test condition plotted against total discharge power.

12.8% for single-channel operation and between 1.7 and 9.7% for multi-channel operation. The average standard deviation was close between single-channel operation at 5.5% and multi-channel operation at 4.3%. The variation in standard deviation does not appear to trend with discharge power, discharge voltage, or channel combination. We suspect that the variation seen, including some of the outlier conditions such as the 12.8% condition, may be a result of how the beam was coupling to the facility day-to-date. Visual observation of the X3 during operation at PEPL showed that the plume would occasionally couple to a certain portion of the facility, on some occasions actually bending downstream of the thruster due to the influence of the facility. We had no visual access to the downstream portion of the beam at GRC, but it is possible that a similar phenomenon was happening that may cause variation in the FFP data. Not shown here is that the probe reading highest and lowest randomly varied from day to day, another indication that the variation was random and facility-related, as opposed to systematic and thruster-related.

8.5 High Speed Diagnostics

8.5.1 High Speed Camera

At each test condition, we collected high-speed camera (HSC) data using the setup described in Chapter 4. We used these data to investigate a number of questions regarding operation of the X3. In particular, for each discharge channel at each test condition we looked at:

1. Is the channel operating in a breathing mode or a spoke mode?
2. What are the peak frequencies?
3. What are the widths of those peaks?
4. What is the probability distribution function of the global channel oscillation?

For channels operating simultaneously, we looked at:

1. Are the channels operating in the same mode?
2. Are the channels oscillating at similar strengths?
3. Do the peak frequencies match between channels?
4. Do the channels correlate in time? (That is, are they oscillating together?)
5. Is there a phase-delayed correlation between channels? (That is, are they oscillating together but with a phase delay?)

Finally, for the same channel operating in different conditions we looked at:

1. Does the mode change?
2. How do the peak frequencies and widths change?
3. How do the relative strengths of the peak frequencies change?
4. How do the probability distribution functions change?

Answers to these questions provide insight into how the thruster oscillations are changing across the throttle table. Because Hall thruster operating mode has been correlated to performance [48, 60], trends in these data can be compared to the performance reported above in Chapter 7. Differences in mode between channels or conditions may correspond with changes in performance or thruster operation.

8.5.1.1 Typical Results

Here we present a subset of typical results from the HSC study and explore the features and trends. To do so, we present single-channel and three-channel results for each channel at the 400 V, $1.0 j_{ref}$ condition and compare how features change between conditions.

First we present results from the Inner channel. Figure 8.22 shows the two-dimensional Fourier transform analysis results plotted as power spectral densities (PSDs) for both single-

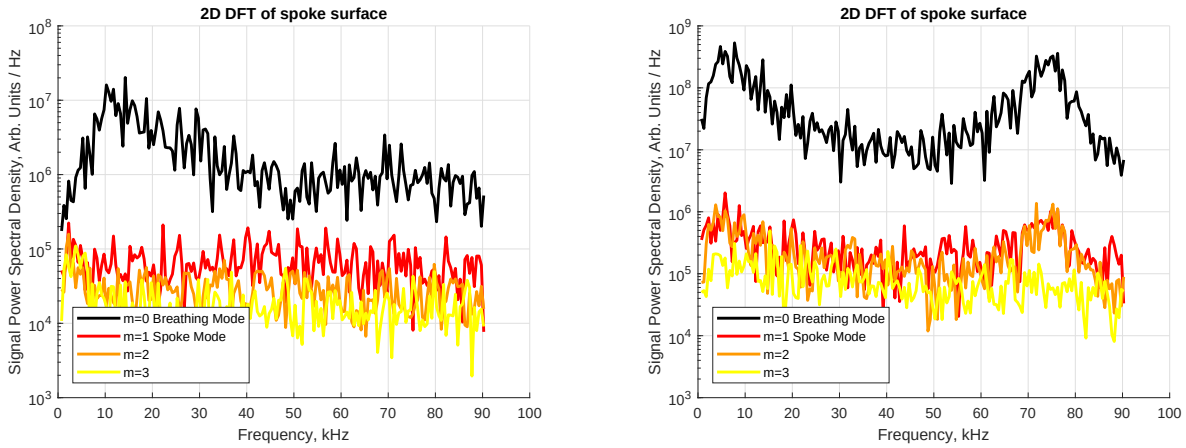

(a) I (b) IMO

Figure 8.22: Inner channel 2D Fourier transform results for (a) single-channel and (b) three-channel operation.

and three-channel operation. The black trace on both plots represents the m=0 mode or the breathing mode; the red, orange, and yellow traces represent spoke modes of increasing m. In both single-channel and multi-channel operation, it is apparent that the m=0 breathing mode is dominant over the spoke modes. The spoke modes are completely flat for single-channel operation, and reflect the same peaks seen in the m=0 mode for three-channel operation, but at a much lower strength. The m=0 PSD features a low-frequency peak at 14 kHz for single-channel operation, which shifts to 6 kHz for three-channel operation. The low-frequency peak is an order of magnitude stronger for multi-channel operation, which corresponds with the change in P2P discharge current as measured with the current guns, which went from $0.36I_{d,I}$ in single-channel operation to $0.59I_{d,I}$ in multi-channel operation. Both low-frequency peaks are generally the same shape and feature a broad distribution. The higher-frequency peak was similar in frequency between the two conditions (67 kHz in single-channel operation and 73 kHz for multi-channel operation) but changed significantly in relative strength: for single-channel operation, that peak is barely distinguishable and is approximately an order of magnitude lower in strength, but in multi-channel operation the peak has grown to be approximately the same strength as the low-frequency peak.

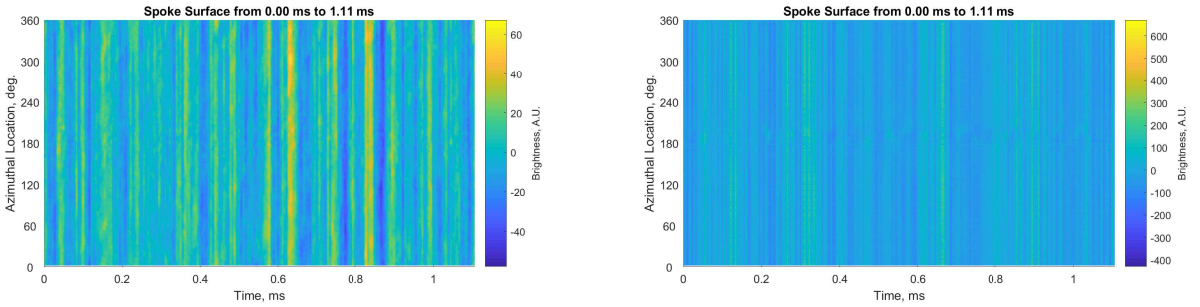

(a) I (b) IMO

Figure 8.23: Inner channel spoke surface ($\phi - t$) plots for (a) single-channel and (b) three-channel operation.

Figure 8.23 shows the spoke surface plots or $\phi - t$ plots for the same conditions. These plots reinforce the fact that the thruster is predominantly in a global or breathing mode: the sharp vertical lines on the plots indicate the entire discharge channel is getting brighter and darker in unison. Unlike the example plot in Figure 4b of Sekerak [48], both of these oscillations appear to vary in strength with time. Sekerak's example breathing mode was very sinusoidal and featured the same peak strength in every cycle; these oscillations vary in strength and are not evenly spaced temporally.

As Figure 8.24 shows, these apparently randomly-varying oscillations translate to Gaussian probability distribution functions (PDFs) for both cases. As discussed by Huang, a sinusoidal oscillation will appear as a double-peaked PDF like that shown in Figure 14 of Reference [63]. The Gaussian nature of the PDF and the random variation in oscillation strength and period are likely related to the large width of the low-frequency peak. Huang's work with the magnetically-shielded NASA HERMeS thrusters demonstrated similar trends: conditions with relatively sharp peaks in their PSDs demonstrated distinctly even spokes on their $\phi - t$ diagrams and a double-peaked PDF (c.f. Figures 8, 14, and 18 in Reference [63]), whereas conditions with broader peaks exhibited Gaussian PDFs and random-appearing global oscillations (c.f. Figures 7, 13, and 17 in Reference [63]). A wide peak on the PDF corresponds to strong oscillations at many frequencies; a Gaussian distribution can be constructed from the combination of many sine waves.

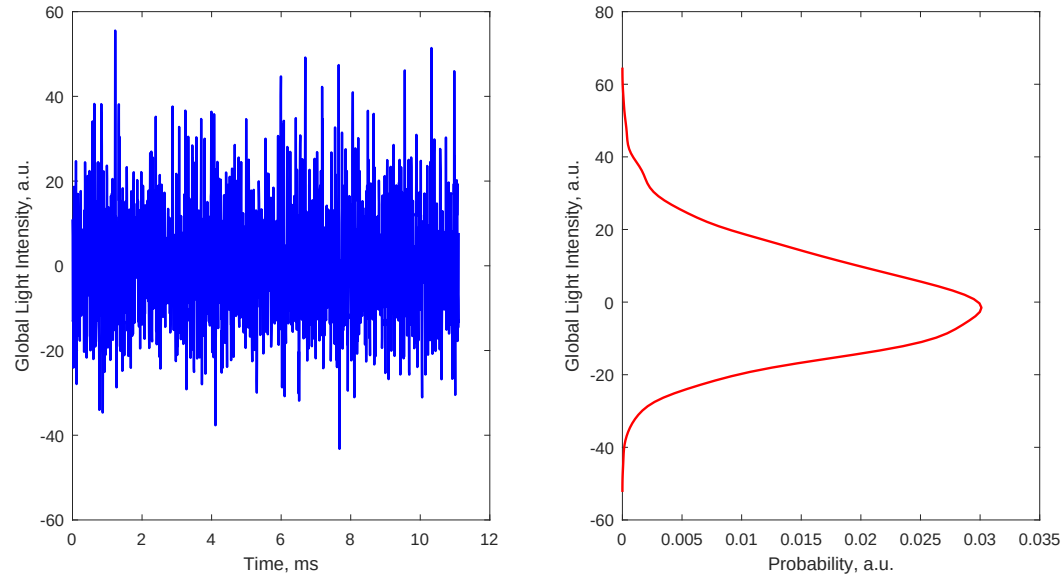

(a) I

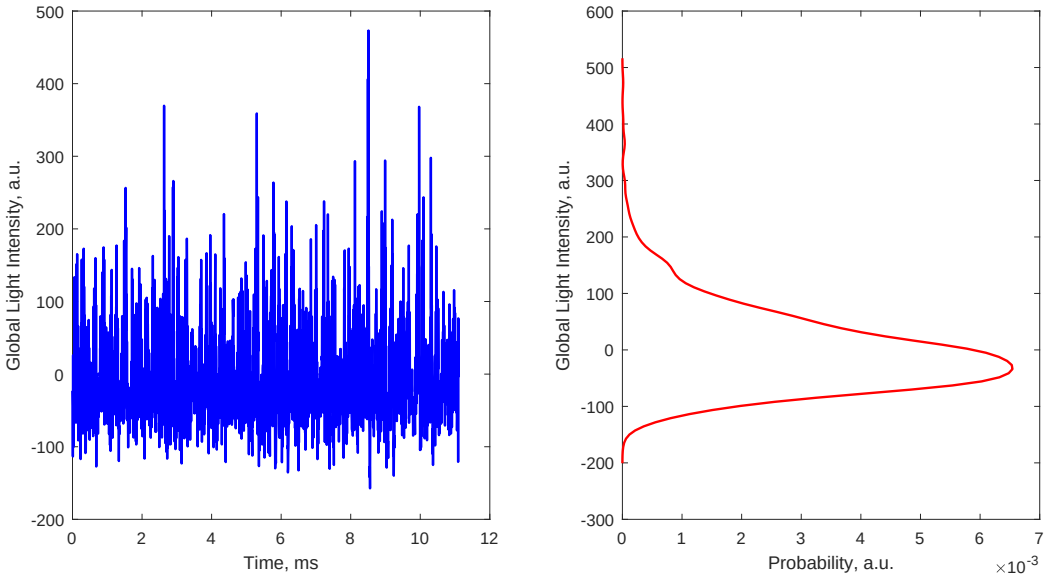

(b) IMO

Figure 8.24: Inner channel global oscillation probability distribution functions for (a) single-channel and (b) three-channel operation.

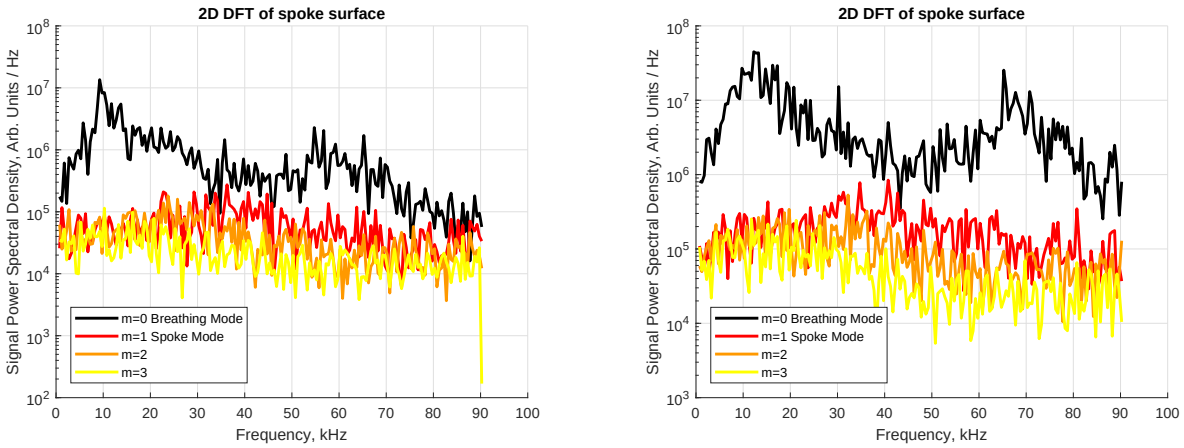

(a) M (b) IMO

Figure 8.25: Middle channel 2D Fourier transform results for (a) single-channel and (b) three-channel operation.

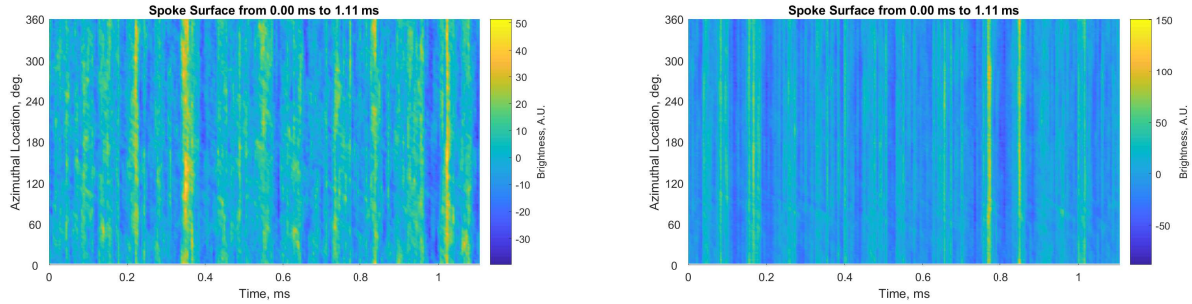

(a) M (b) IMO

Figure 8.26: Middle channel spoke surface ($\phi - t$) plots for (a) single-channel and (b) three-channel operation.

Next we turn to the Middle channel. As shown in Figures 8.25, 8.26 and 8.27, the results are largely the same as for the Inner. In both single-channel and three-channel operation, the Middle is strongly in a m=0 or breathing mode. The low-frequency and high-frequency peaks are roughly constant in location between the two conditions: the low-frequency peak shifted from 11 kHz in single-channel operation to 14 kHz in three-channel operation, and the high-frequency peak shifted from 57 kHz to 67 kHz from single- to three-channel operation. Once again, the high-frequency peak gains in strength in three-channel operation, though not equal to the low-frequency peak like for the Inner channel. The spoke surfaces in Figure 8.26 also

appear very similar to those for the Inner channel, indicating that the Middle channel is also operating in a mode characterized by strong breathing at a randomly-appearing strength and period. This is further supported by the PDFs in Figure 8.27, which show similar Gaussian-like shapes for both conditions. Unlike for the Inner, both spoke surfaces for the Middle channel show faint diagonal striations as well, indicating very faint spoke features. These striations are focused in the 60–120 degree location and appear to die out instead of propagating around the entire channel. They travel in the counter-clockwise direction (a negative slope on the spoke surface), which is the $\vec{E} \times \vec{B}$ direction for the Middle channel. Work on the H6 thruster demonstrated that spokes propagate in the ExB direction and change directions when the magnetic field direction is reversed on that thruster [60].

Finally, we turn to the Outer channel in Figures 8.28, 8.29, and 8.30. And once again, the trends are largely the same as for the Inner and Middle channels. A strong m=0 breathing mode dominates, with low- and high-frequency peaks. The low-frequency peak shifts from 9 kHz in single-channel operation to 15 kHz in multi-channel, whereas the high-frequency peak stays at 55 kHz for both cases. Unlike the Inner and Middle channels, the high-frequency peak loses strength relative to the low-frequency peak for three-channel operation. The spoke surfaces in Figure 8.29 demonstrate very similar features as those for the other channels as well. Here, there are localized spoke features in two locations around the channel visible for both cases, in the 60–120 degree and 240–300 degree locations (roughly opposite from each other on the thruster face). These spoke features propagate in the clockwise direction (positive slope on the $\phi - t$ diagrams), opposite of the Middle channel. However, because the direction of the magnetic field switches between the Middle and Outer channels, the spokes here are still propagating in the $\vec{E} \times \vec{B}$ direction. Like for the Middle channel, in both conditions these spokes remain local and die out before propagating fully around the channel. They are superimposed over the global oscillation of the discharge, and their relative weakness is reflected in the PSDs. The PDFs for the Outer channel, as shown in Figure 8.30, are also very similar to those for the other two channels, which is to be expected.

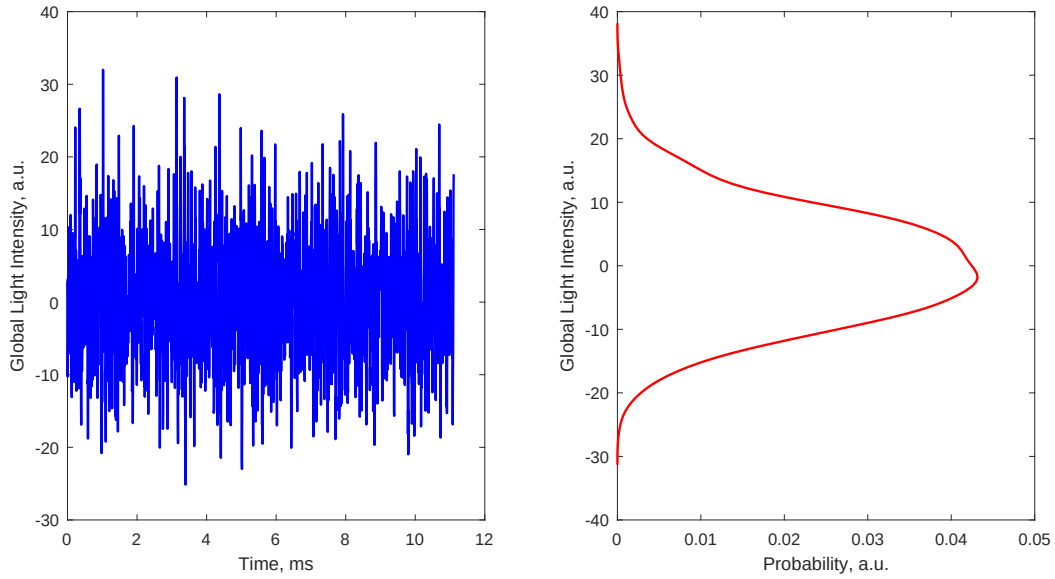

(a) M

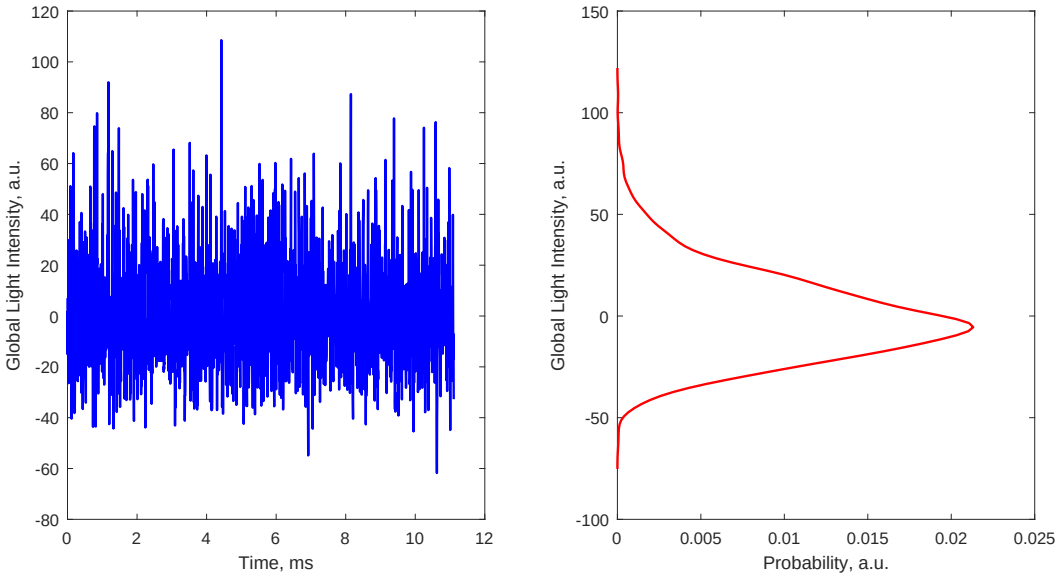

(b) IMO

Figure 8.27: Middle channel global oscillation probability distribution functions for (a) single-channel and (b) three-channel operation.

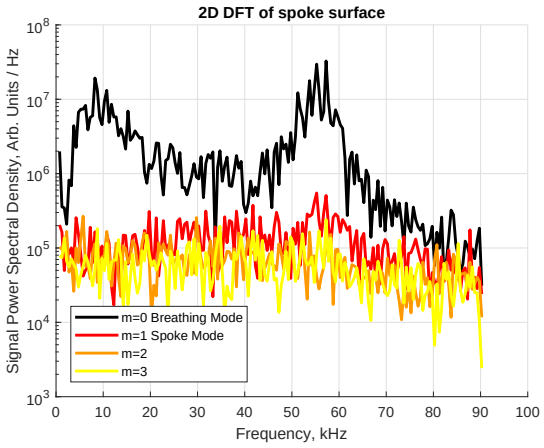

(a) O

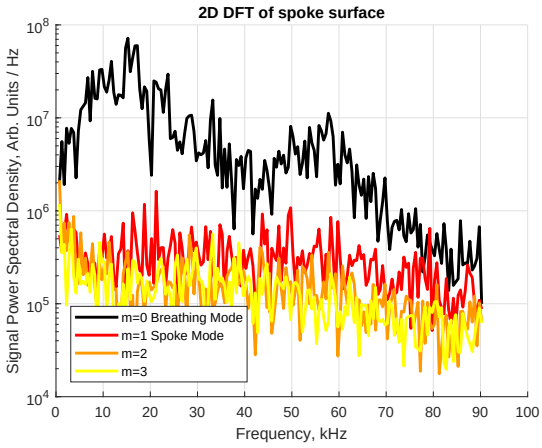

(b) IMO

Figure 8.28: Outer channel 2D Fourier transform results for (a) single-channel and (b) three-channel operation.

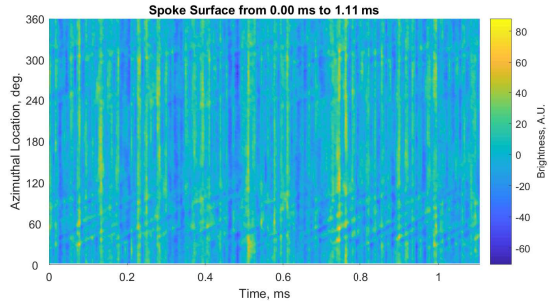

(a) O

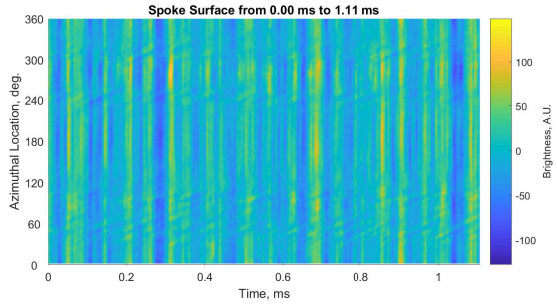

(b) IMO

Figure 8.29: Outer channel spoke surface ($\phi - t$) plots for (a) single-channel and (b) three-channel operation.

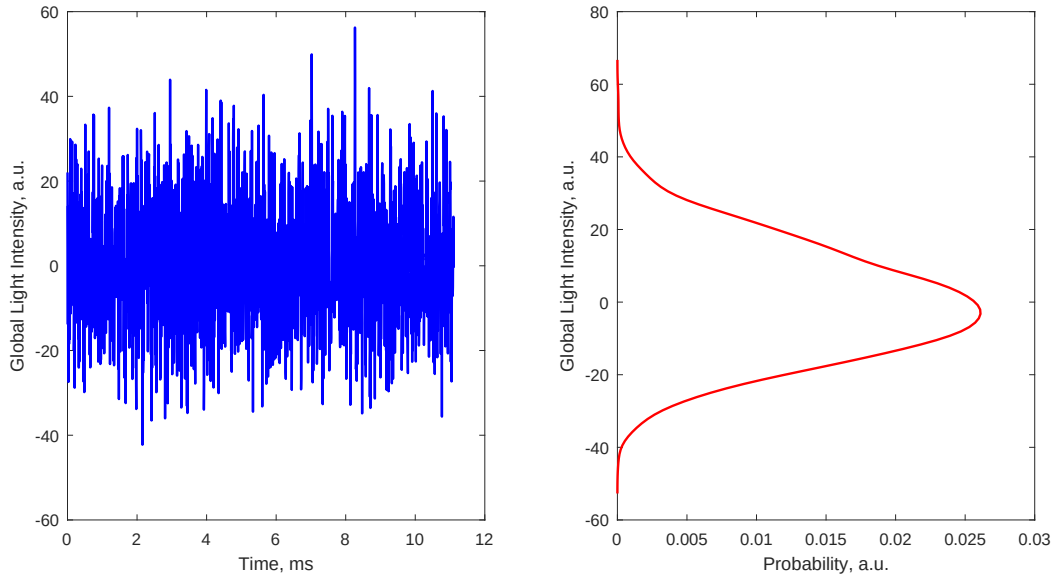

(a) O

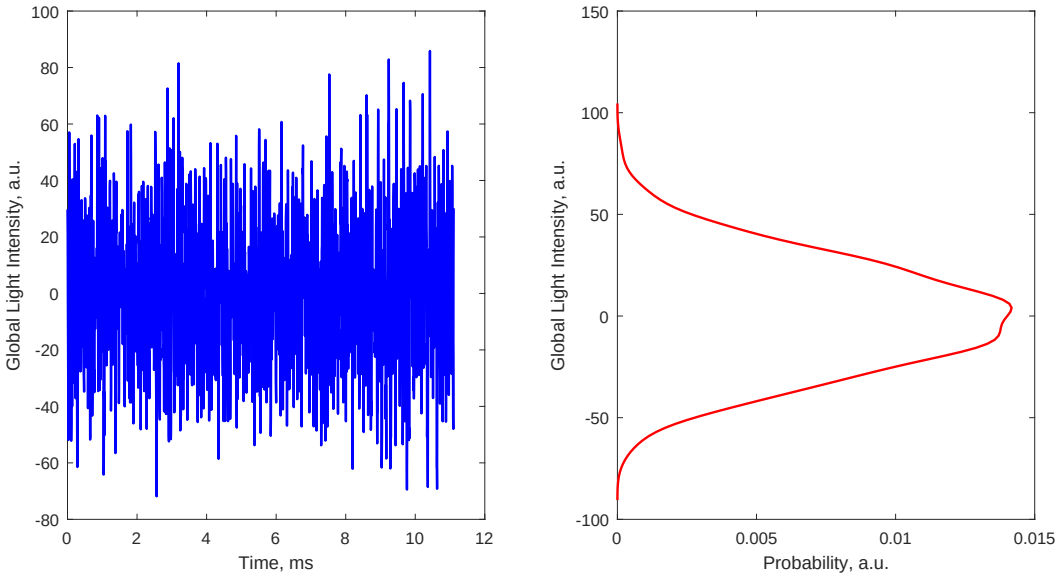

(b) IMO

Figure 8.30: Outer channel global oscillation probability distribution functions for (a) single-channel and (b) three-channel operation.

All evidence suggests that all three channels, whether operating separately or together, are operating in a very similar mode.

8.5.1.2 Oscillation Peak Frequencies

As demonstrated in the above example data, the oscillations of all three channels of the X3 are typically characterized by a low-frequency peak on the order of 10 kHz and a high-frequency peak on the order of 50-80 kHz. To understand how these frequencies vary by channel, we present a histogram of the values of both peaks across all conditions. Histograms are useful for these situations because they will illustrate whether there are multiple clusters of conditions. Two distinct groupings of frequencies on these plots could potentially indicate that, for instance, the most common frequency was different for single-channel and multi-channel operation. A plot of the low-frequency peaks for each channel across all conditions is presented in Figure 8.31. These peaks varied in location from 6–20 kHz, with most conditions for all three channels falling into the 10–15 kHz range. The Inner demonstrated slightly more spread in peak location than the other two channels, but the difference between the three channels is minimal. There are no indications of multiple populations on any channels.

However, the high-frequency peak shows clear variation between the channels, as shown in Figure 8.32. The Inner channel high-frequency peak showed the largest spread of values, ranging from 50–95 kHz, with most conditions clustered between 65–85 kHz. The Middle channel featured a peak between 65–70 kHz for a vast majority of its conditions, with a small group of conditions between 40–60 kHz and a single outlier in excess of 100 kHz. This very-high-frequency result may be an artifact of the fitting technique. The high-frequency peak of the Outer channel was very tightly clustered between 45–60 kHz for all test points. The Middle channel's results indicate that the possibility of multiple populations, with the 50–60 kHz conditions detectably separate from the strong cluster of 65–70 kHz conditions, but otherwise the channels again show single populations for the high-frequency peak.

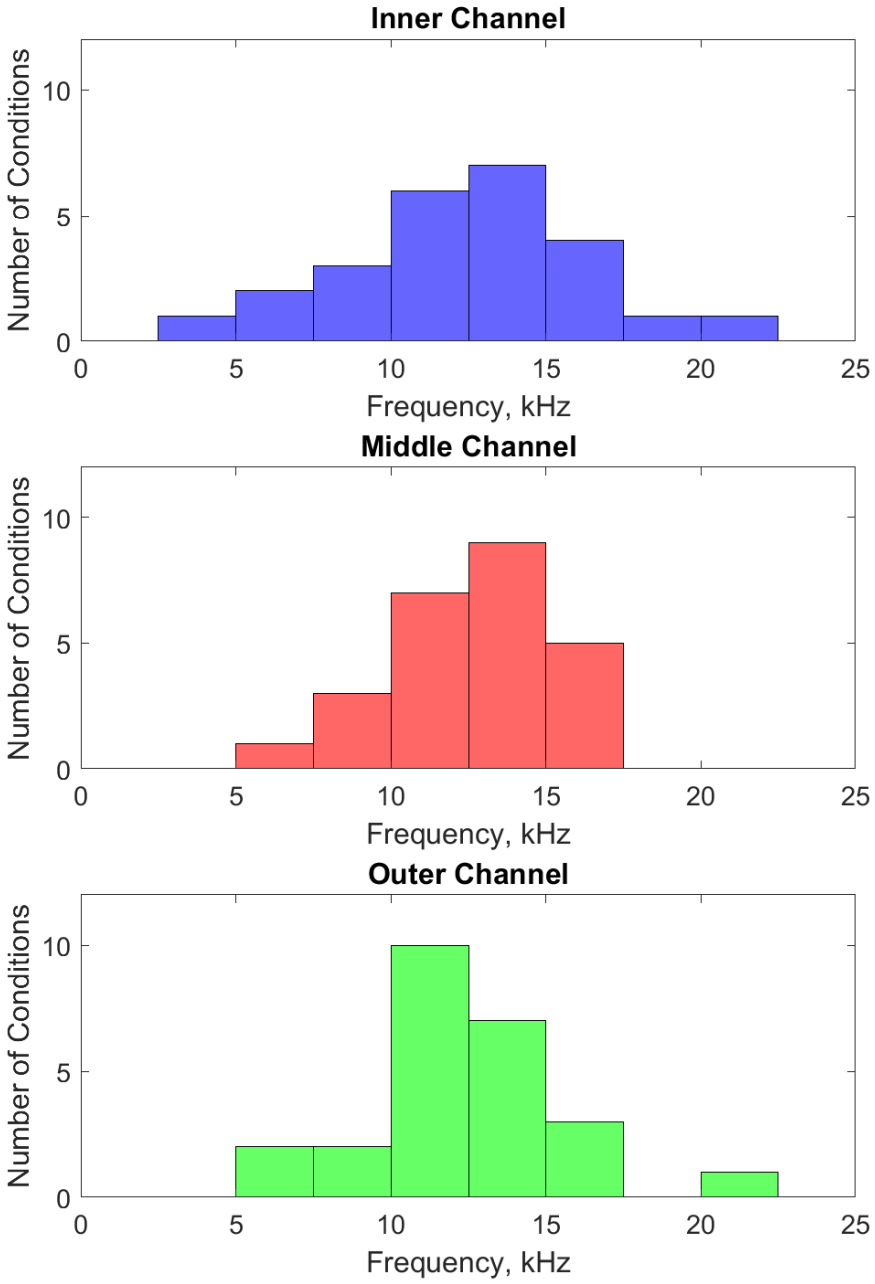

Figure 8.31: Histograms of the location of the low-frequency peak for each channel of the X3 for all test conditions.

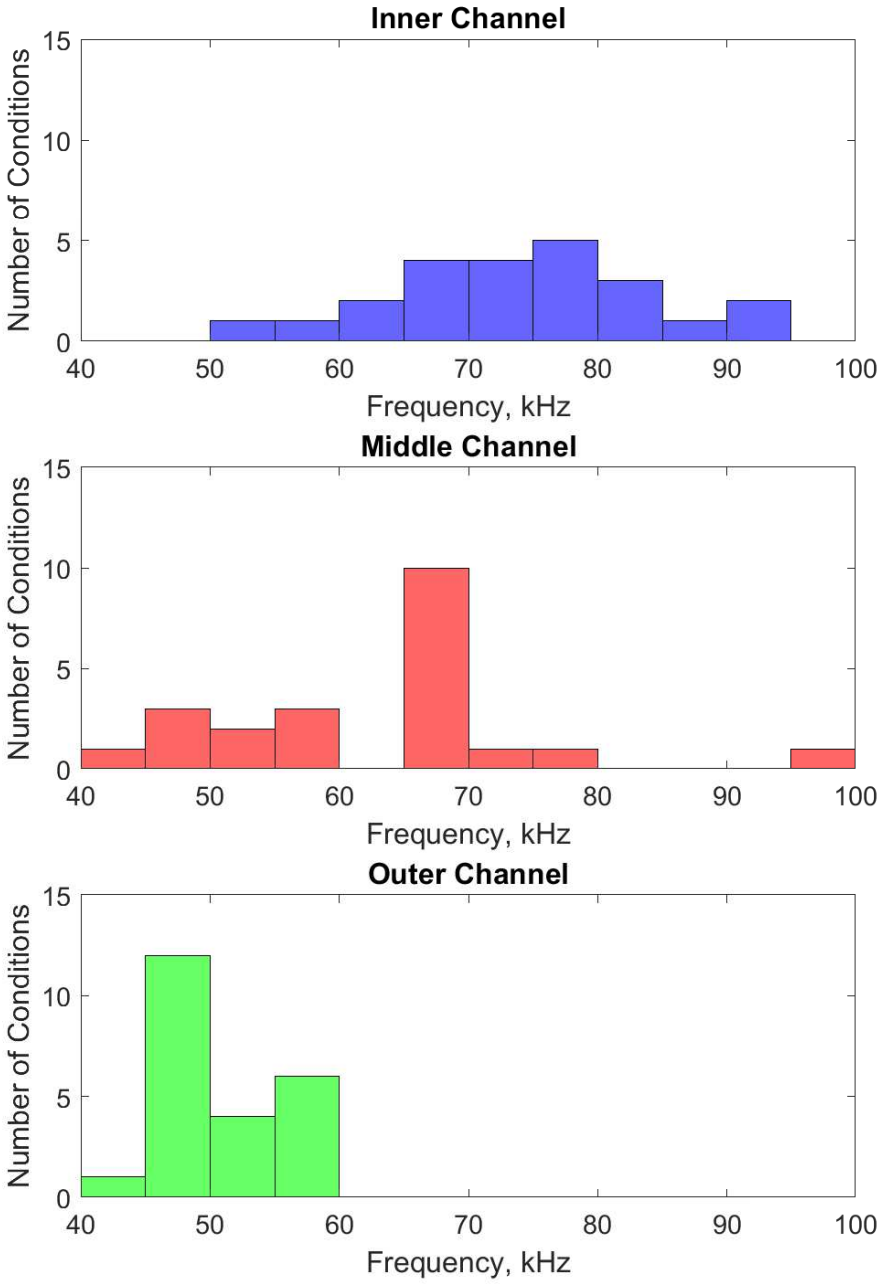

Figure 8.32: Histograms of the location of the high-frequency peak for each channel of the X3 for all test conditions.

8.5.1.3 Correlation to Oscillation Strength

As shown in Chapter 7, the P2P oscillation strength of the Inner and Middle channels was notably larger during multi-channel operation than during single-channel. A number of characteristics of the discharge current oscillations changed between single- and multi-channel operation, including the location and width of the low- and high-frequency peaks, as well as the ratio of the strength of these peaks (high:low). To assess whether the change in oscillation strength (defined, as in Chapter 7, as the ratio of P2P of the discharge current to the mean discharge current, I_{P2P}/I_d) is correlated with any of these features, we performed a correlation study similar to those performed by Huang previously on both HSC data and RPA data on the HERMeS thrusters [63, 219]. By plotting I_{P2P}/I_d against the frequency, width, and strength of the low-frequency and high-frequency peaks we can determine whether any parameters correlate with oscillation strength.

First, Figure 8.33 presents the oscillation strength I_{P2P}/I_d as a function of both the low-frequency and high-frequency peak strengths. This can be viewed as a validation of this technique, as we expect stronger discharge current oscillations to correspond to stronger features on the HSC PSDs. If the HSC analysis is properly capturing the oscillations, increasing peak strength should correspond to increasing I_{P2P}/I_d. Indeed, we find that both plots generally show this trend. There is a certain amount of variation, but this is to be expected because of the statistical noise associated with the HSC analysis technique. In general, the low-strength oscillation conditions ($I_{P2P}/I_d < 0.5$) cluster at lower peak strengths and the high-strength oscillation conditions ($I_{P2P}/I_d > 0.5$) appear at higher peak strengths. This indicates that the HSC technique is properly capturing the strength of the discharge current oscillations, and further correlation analysis should identify trends if present.

Figure 8.34 presents the location of the low-frequency and high-frequency peaks versus I_{P2P}/I_d for all conditions of X3 operation. Oscillation strength shows no trend with the low-frequency peak (left of the figure). This agrees with the histograms in Figure 8.31, which showed the conditions clustering together in single populations centered around frequencies

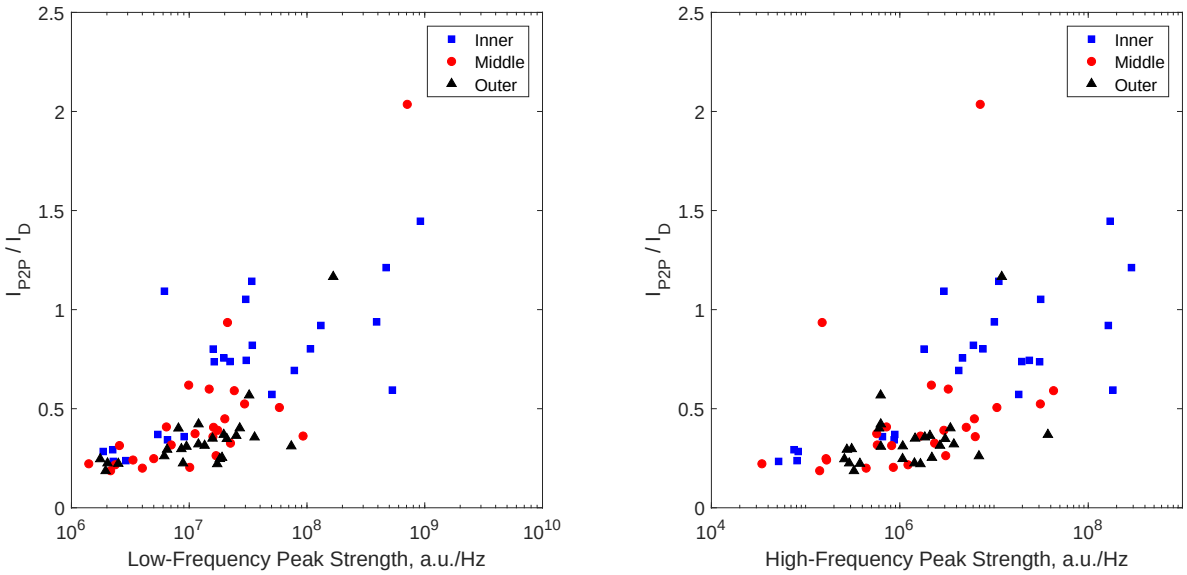

(a) Low-frequency peak strength

(b) High-frequency peak strength

Figure 8.33: Discharge current oscillation strength (I_{P2P}/I_d) for each channel plotted as a function of the strength of the (a) low-frequency and (b) high-frequency peaks of the HSC PSDs.

of 10–15 kHz. The data demonstrate that the strong-oscillation conditions ($I_{P2P}/I_d > 0.5$) fall within the same range of frequencies as the weak-oscillation conditions ($I_{P2P}/I_d > 0.5$), and furthermore that at a given frequency, both strong- and weak-oscillation conditions exist for the same channel. The high-frequency peak shows a similar lack of clear trend. For the Inner channel (blue squares on the plot), it appears that the weak-oscillation conditions cluster in the 60–70 kHz range and the strong-oscillation conditions in the 70–90 kHz region; however, the fact that some of the strongest and weakest oscillations occur around 70 kHz for that channel go against this trend and suggest that there may be no correlation. The Middle channel demonstrates a similar overlap in strong and weak conditions. The Outer channel, which did not experience a significant change in oscillation strength between conditions, shows the smallest cluster of conditions, as was shown in the histogram in Figure 8.32.

The peaks detectably changed widths among conditions. Figure 8.35 presents the oscillation strength as a function of the width of both the low- and high-frequency peaks.

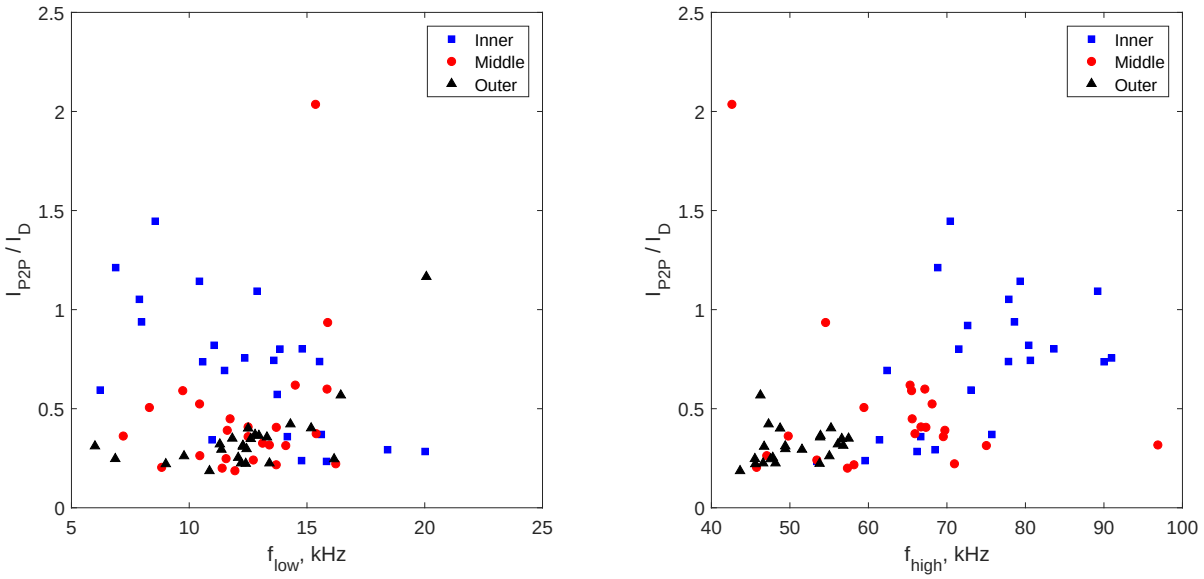

(a) Low-frequency peak

(b) High-frequency peak

Figure 8.34: Discharge current oscillation strength (I_{P2P}/I_d) for each channel plotted as a function of the location of (a) the low-frequency peak and (b) the high-frequency peak.

The low-frequency peak width ranged from 1–10 kHz and the high-frequency peak width from 5–45 kHz. All channels experienced approximately the full range of widths for both peaks. As can be seen, neither peak width shows any trend with oscillation strength. The Inner appears to have two clusters for the high-frequency peak width, one between 5–15 kHz and another between 30–40 kHz. However, both of these clusters contain high-strength oscillations, indicating that there is no trend between oscillation strength and peak width. The Middle and Outer channels are both relatively evenly dispersed along the range of peak widths.

As shown in Figure 8.36, we also study the ratio between the strength of the two peaks. If for instance the high-frequency peak is a cathode-related (something that needs further study to confirm), the high-frequency peak being on the same order as the low-frequency peak would represent cathode oscillations being equal in strength to the thruster breathing mode, which could potentially result in stronger oscillations. However, here, as previously, there is clearly no trend between the oscillation strength I_{P2P}/I_d and the ratio of high-frequency

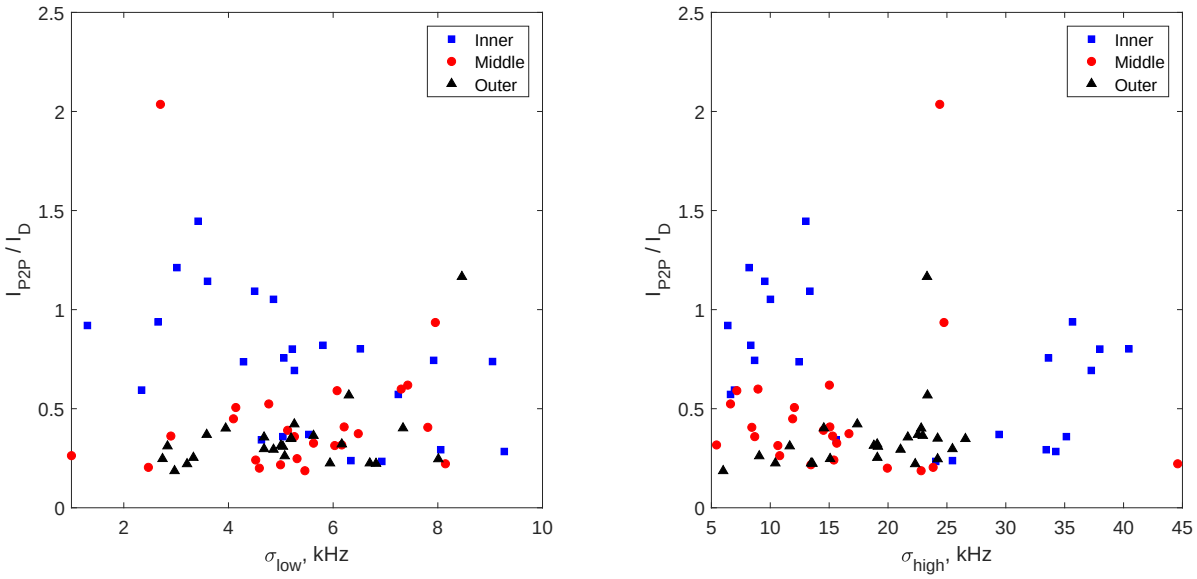

(a) Low-frequency peak width

(b) High-frequency peak width

Figure 8.35: Discharge current oscillation strength (I_{P2P}/I_d) for each channel plotted as a function of the width of the (a) low-frequency peak and (b) the high-frequency peak.

peak strength to low-frequency peak strength. For all three channels, the high-frequency peak ranges from being roughly as strong as the low-frequency peak to being two orders of magnitude weaker. And for all three channels, both high-strength and low-strength I_{P2P}/I_d values occur through the range of magnitudes.

Ultimately, this analysis did not identify any clear correlations between the discharge current oscillation strength I_{P2P}/I_d and features of the HSC PSDs (other than peak strength, which was physically expected and seen as a validation of the technique). This indicates that the cause of the dramatic change in oscillation strength for the Inner and Middle channels between single-channel and multi-channel operation, as presented in Chapter 7, is a process unrelated to those causing the changes in the features of the PSDs captured in this analysis.

8.5.1.4 Correlation Between Channels

We analyzed the cross-correlation between the global oscillations in the channels at each multi-channel condition to study whether channels were oscillating in sync with each other

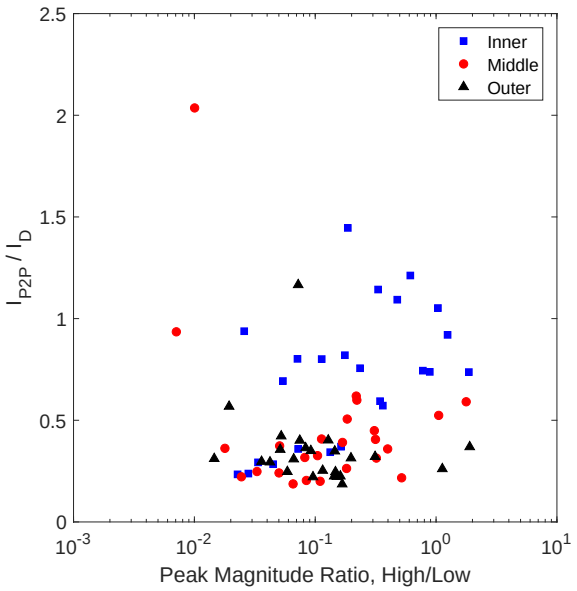

Figure 8.36: Discharge current oscillation strength (I_{P2P}/I_d) for each channel plotted as a function of the ratio of the strength of the high-frequency and low-frequency peaks of the HSC PSDs.

(either in-phase or with a phase delay between channels). The non-sinusoidal nature and the general lack of agreement between peak locations among channels operating together suggests that the channels are likely not oscillating with any sort of correlation between each other, and our results showed that to be true. Even for conditions where the peaks on the PSDs appeared in to be at the same frequency, the brightness of the discharge channels showed no correlation at any phase delay. This indicates that, at least for the thruster conditions tested, the oscillations between channels do not interact in any significant or meaningful way. A study by McDonald showed that for the X2 thruster in certain operating conditions the low-frequency peak of the outer channel would "bleed over" and appear in the inner channel's PSD trace (though the inner channel's peak never appeared in the outer channel's PSD trace) [195]. No concrete explanation was found for this behavior on the X2, and it is unclear from this work whether this type of phenomenon would appear for the X3 under certain operating conditions. However, the operating conditions tested here were optimized for performance, indicating that any condition where this type of channel

cross-talk was present would be an off-nominal condition.

8.5.1.5 Atypical Conditions

The vast majority of operating conditions across the entire range of X3 operation reported here had HSC results very similar to those we presented above. However, three test conditions exhibited HSC data that was, in various ways, atypical. These conditions will be discussed briefly here.

Figure 8.37 shows the HSC results from the Inner channel operating in the IMO configuration at 300 V, $0.63 j_{ref}$. The Middle and Outer channels demonstrated typical HSC results for this condition, but the Inner shows a distinctly different oscillation mode. In the Fourier transform results, four distinct peaks can be seen at approximately 16, 32, 48, and 64 kHz, with a fifth peak faintly discernible around 80 kHz. The four higher-frequency peaks represent the second, third, fourth, and fifth harmonics of the sharp 16-kHz peak. The appearance of multiple peaks, as well as of harmonics of the low-frequency breathing mode, has been seen before in Hall thruster HSC studies (c.f. the top left of Figure 4.1 of Sekerak [60]). A strong, sharp breathing-mode peak is typically associated with strongly sinusoidal oscillations, and the spoke plot in Figure 8.37b shows this to be true for this condition. Unlike typical spoke surface plots, this result shows very periodic brightness and darkness in the discharge channel. The Fourier transform results indicate that the $m > 0$ spoke modes also have considerable strength and content, but the spoke surface plot seems to indicate that these may simply be an artifact of the analysis technique: there are distinctly no signs of any spoke-like features propagating azimuthally in the channel. Interestingly, the PDF from this condition does not show the two-peak shape that Huang associates with sinusoidal oscillations in the HERMeS TDU1 thruster [63].

At this condition, the P2P value was nearly four times the mean discharge current. The oscillations are so large that it is likely that the discharge is truly extinguishing on the low side of the sinusoid. In fact, the sharp bottom edge of the global light intensity (the left-hand

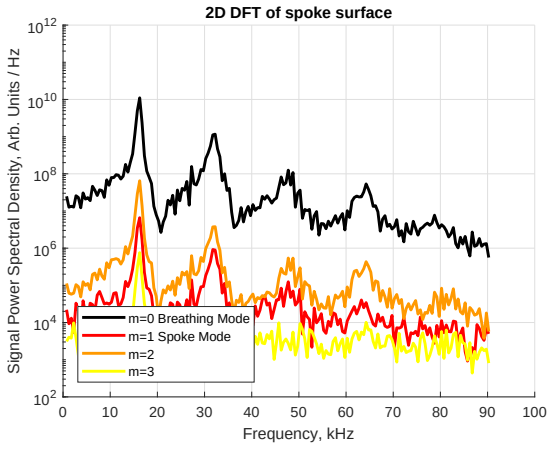

(a) 2D Fourier transform

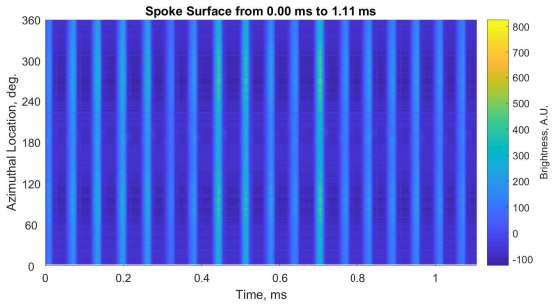

(b) Spoke surface ($\phi - t$) plot

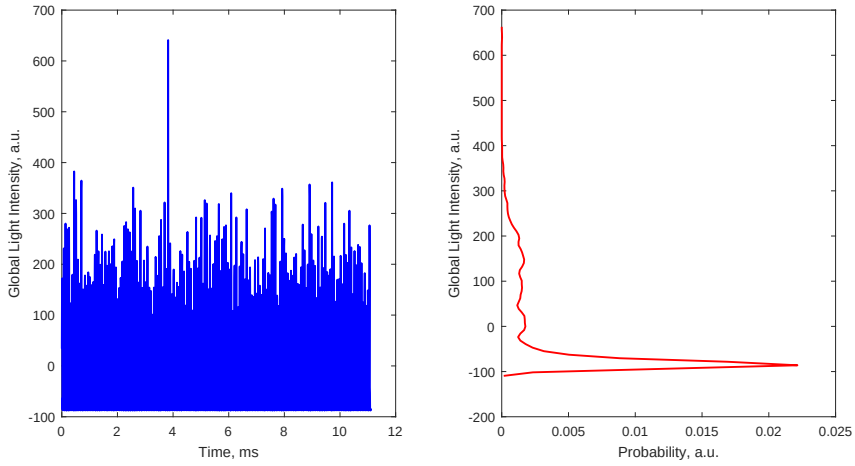

(c) Probability distribution function

Figure 8.37: HSC results for the Inner channel in the IMO configuration at 300 V, $0.63 j_{ref}$: (a) the 2D Fourier transform, (b) the $\phi - t$ diagram, and (c) the PDF.

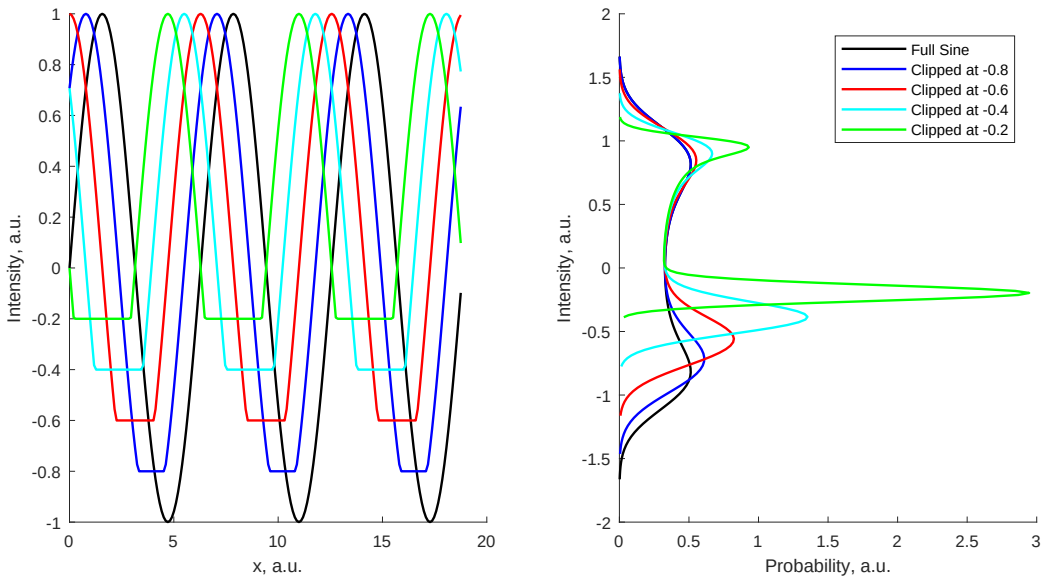

Figure 8.38: Sine waves clipped at various magnitudes (left) and the resultant probability distribution functions (right). A Hall thruster discharge that is oscillating in a global mode heavily enough that it is extinguishing can be thought of as a clipped sine wave.

plot of Figure 8.37c) seems to corroborate this. The discharge extinguishing can be thought of as a sinusoid clipping, and the PDF of a clipped sinusoid increasingly skews to a single peak with a long tail. This is illustrated in Figure 8.38, which shows the PDFs that result from sine functions clipped at various magnitudes. Ultimately, we suspect that this aggressive oscillation was due to a hysteretic effect related to how the channel was lit, perhaps related to the channel struggling to couple with the cathode. As discussed in Appendix A, the X3 experienced a small number of oscillation anomalies that appeared to be related to the lighting of a channel when other(s) were already firing. We returned to this firing condition later in the test campaign and the Inner channel experienced typical oscillations for IMO operation and exhibited typical HSC data, unlike what is shown in Figure 8.37.

The next atypical condition was for the Middle channel operating in the MO configuration, the results of which we present in Figure 8.39. Here, the Middle channel exhibited oscillations that were very similar in character to those exhibited by the Inner channel above. The Fourier transform uncovers a sharp breathing mode peak and at least one visible har-

monic (at approximately 15 and 30 kHz, respectively) and shows some features in the $m > 0$ spoke modes. The spoke surface plot shows distinct, periodic oscillations without any sign of spoke-like features. Much like the Inner channel above, the global light intensity and subsequent PDF indicate the discharge may be extinguishing during the period of the oscillation, though perhaps not to the degree of the Inner channel seen in the previous example. Indeed, the P2P value for this condition was approximately two times the mean discharge current. Once again, we suspect this may be a hysteretic effect related to the coupling with the cathode for this channel. Later in the same test day, we operated the thruster in the IMO configuration at the same discharge voltage, current density, and magnetic field conditions, and the Middle channel exhibited typical oscillations and HSC results unlike those shown here.

The final outlier in the HSC analysis was the Outer channel operating in the IMO configuration at 500 V, $1.0 j_{ref}$. This condition was different than all others observed in this test campaign. As shown in the temporal plot of the global light intensity in the left of Figure 8.40c, the discharge was characterized by occasional, somewhat periodic spikes of the discharge brightness that were several times the typical level of oscillation. As shown in the spoke surface, which captures one of the spikes at around 0.7 ms, the oscillations in the channel were otherwise very typical. These surges of brightness occurred at approximately 1 kHz, which is reflected by the increase in the Fourier transform results at very low frequencies. These surges did not appear in the other two channels operating at the same time, and their source is unclear from these data. Other than the surges, the oscillations appear typical in character as demonstrated by the Fourier transform, spoke surface, and PDF results. The PDF has a long tail due to the spikes but is otherwise very Gaussian in shape.

For all three of these outlier cases the thruster performance did not vary significantly or fall out of family from other test points at similar conditions. More work is necessary to fully understand the mechanisms behind these oscillations and the reasons that the oscillation mode would occasionally change, but it is nonetheless encouraging to see that thruster

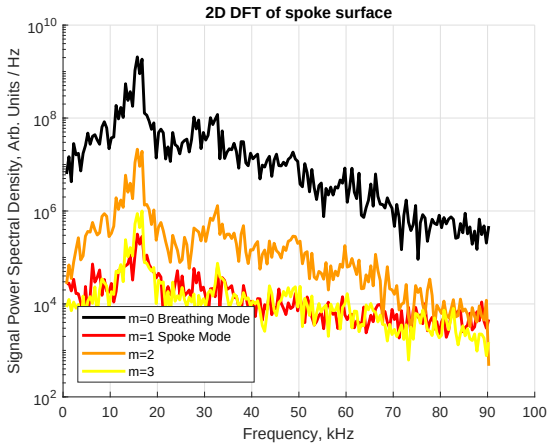

(a) 2D Fourier transform

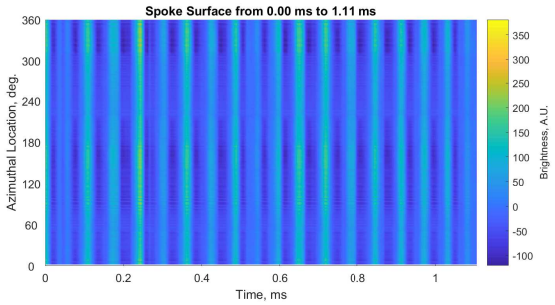

(b) Spoke surface ($\phi - t$) plot

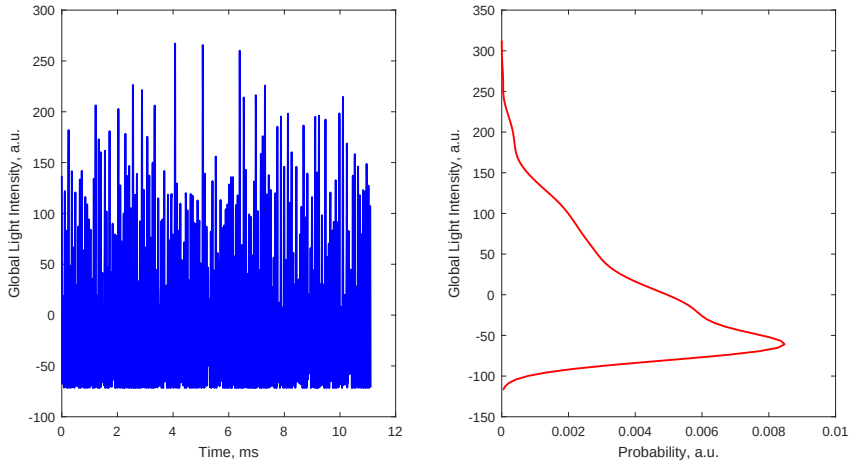

(c) Probability distribution function

Figure 8.39: HSC results for the Middle channel in the MO configuration at 500 V, $0.63 j_{ref}$: (a) the 2D Fourier transform, (b) the $\phi - t$ diagram, and (c) the PDF.

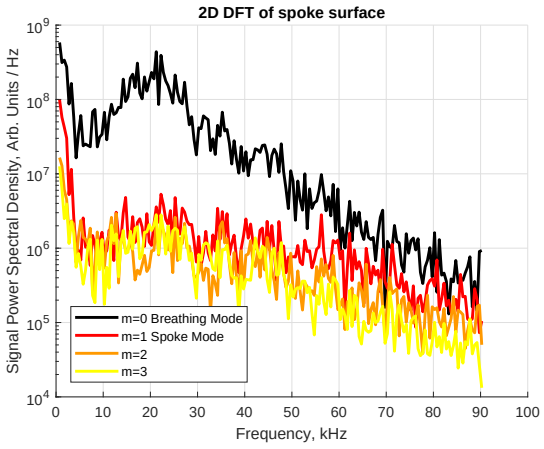

(a) 2D Fourier transform

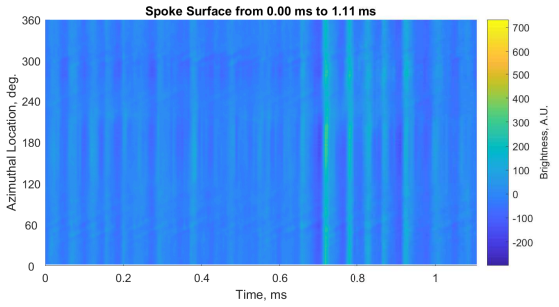

(b) Spoke surface ($\phi - t$) plot

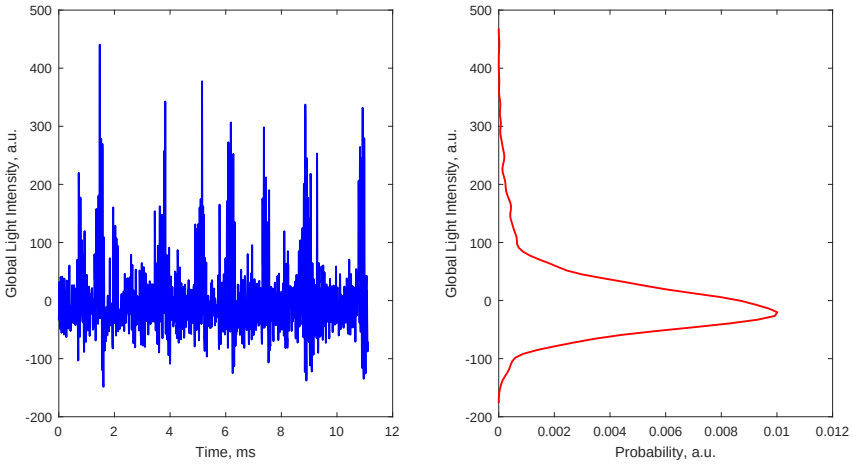

(c) Probability distribution function

Figure 8.40: HSC results for the Outer channel in the IMO configuration at 500 V, $1.0 j_{ref}$: (a) the 2D Fourier transform, (b) the $\phi - t$ diagram, and (c) the PDF.

performance is not adversely affected when the oscillations do change.

8.5.1.6 Discussion

The HSC analysis indicated that the X3 was operating in a similar oscillation mode not only for all channel combinations at a given discharge voltage and current density but in fact for practically all conditions tested. The fact that the thruster operated in the same oscillation mode across all channel combinations and conditions has important implications both for thruster design and for ground testing.

To date, no HSC measurements have been performed on a Hall thruster larger than the 20-kW NASA-300MS [12]. The analysis presented here thus represents not only the first detailed HSC study of the X3 but also the first of a Hall thruster operating in excess of 20 kW and the first on discharge channels the size of both the Middle and Outer channels. The Middle channel of the X3 is approximately the same diameter as the NASA-457M, which never had any HSC studies performed on it. The Outer channel represents the largest Hall thruster discharge channel ever built. It is a significant result that these three discharge channels demonstrated nearly identical oscillation modes. These data indicate that scaling Hall thruster discharge channels using the NASA/Manzella scaling [11] results in oscillations that are of the same mode and character at a given discharge voltage/current density/magnetic field strength condition regardless of channel diameter or discharge power. This then implies that the mechanism driving these oscillations is borne from a feature (or features) of the thruster that are constant across all channels with this scaling technique, features such as current density, channel length, magnetic field topology, etc. Though more work is needed to identify the mechanism behind these oscillations, our results here indicate that oscillations stay of a similar mode and frequency for Hall thruster discharge channels up to approximately 80-cm diameter and channel power levels up to 55 kW (the maximum power the outer channel was tested to). No new oscillation modes have been identified in these larger channels or in multi-channel operation, implying that ongoing work in understanding the physical

mechanisms driving oscillations in 6–12 kW thrusters [12, 60, 61, 63, 193] will directly relate to the oscillations in NHTs, high-power thrusters, and large-diameter thrusters.

These results also have significant implications for ground testing. Previous work by Huang with the 12.5-kW NASA HERMeS thruster demonstrated that for that lower-power, magnetically-shielded thruster, oscillation characteristics did not change significantly with increasing facility background pressure, except for a single 800-V condition [63]. Though our experiments were not as rigorous as Huang's study, a given discharge channel did not see significant changes in oscillation mode between single- and multi-channel operation. There was a change in oscillation strength on the Inner and Middle channels that has not been fully explained by this work, but as discussed in Chapter 7, it seems as though these are not driven by changes in facility background pressure. For a given discharge voltage/current density test point, each channel experienced a significant variation of facility background pressure—approximately a two-fold increase in pressure for the Outer channel between single- and three-channel operation and nearly a six-fold increase for the Inner channel. The fact that these large changes in facility background pressure, which are of similar span as a number of background pressure studies in the literature [63, 111], do not appear to affect the oscillation mode of the X3, implies that there are not any significant background-pressure effects to attempt to remove from the data.

Large variations in background pressure across conditions are an unfortunate consequence of the large throttling ranges of NHTs. Maintaining a constant facility background pressure through downstream injection of propellant is likely prohibitively expensive for higher-power thrusters, which are already prone to higher propellant costs due to their large flow rates. Additionally, artificially increasing the pressure moves the test environment farther from space-like conditions, which goes against general ground-testing philosophy. Though the pressures experienced in VF5 during higher-current operation of the X3 were in excess of what is currently suggested for Hall thruster test conditions [111, 220], the fact that the oscillation character did not vary significantly between these higher-pressure conditions and

those at lower pressures is encouraging for future characterization and flight-qualification of the X3 and other high-power Hall thrusters in testing facilities with finite pumping capability.

8.5.2 High-Speed Discharge Current Analysis

Additionally, at each performance test point we collected high-speed discharge current measurements using the current guns and oscilloscopes described in Chapter 4. These measurements were triggered by the HSC data acquisition and thus are synchronized to those data. We applied a fast Fourier transform to these results as well, and from that calculated a PSD. These PSDs can only study the $m = 0$ global oscillation mode since no spatial information is known. However, we can compare these results to those from the HSC analysis as a verification of the HSC technique. The HSC and the high-speed discharge current (HSDC) analysis should both capture the $m = 0$ oscillations of the thruster.

Figure 8.41 presents four example conditions comparing the HSC results with those from the HSDC analysis. The HSDC traces have been filtered using a Savitzky-Golay technique for clarity [183]. We find that the traces from both techniques exhibit the same shape and features. The peak locations, widths, and relative heights match between the two analysis techniques, though the absolute magnitudes of the traces differ due to the differences in both the raw data and in the analysis techniques. These example results are representative of the correlation between the two techniques across the test matrix. The HSDC analysis also captures the atypical results such as those for the Inner channel in the IMO configuration at 300 V, $0.63 j_{ref}$. We present a comparison of these results in Figure 8.42. The same series of peaks is clearly distinguishable on both traces, indicating that the result was not an artifact of the HSC analysis technique. No further data analysis is necessary with the HSDC traces, as we have extracted all relevant information from the HSC analysis above. However, we take the strong similarity between the two results, including the capturing of the atypical features of certain conditions, as a validation of the HSC technique.

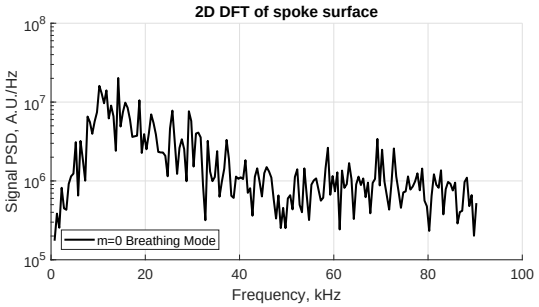

(a) HSC analysis for Inner in I configuration

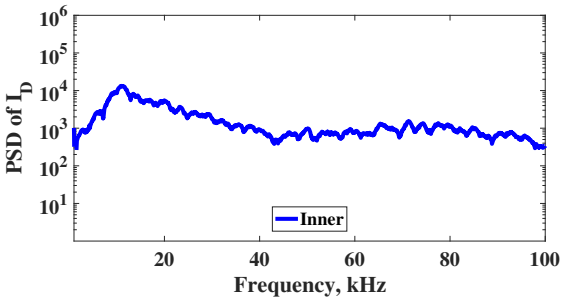

(b) HSDC analysis for Inner in I configuration

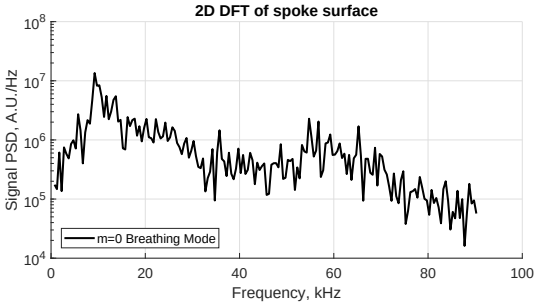

(c) HSC analysis for Middle in M configuration

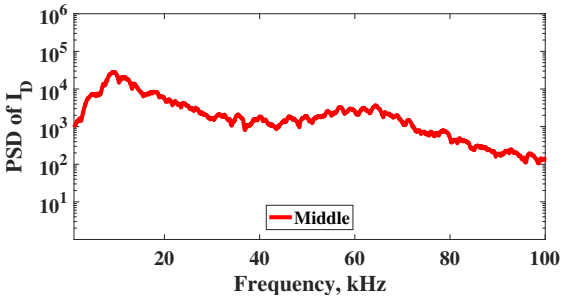

(d) HSDC analysis for Middle in M configuration

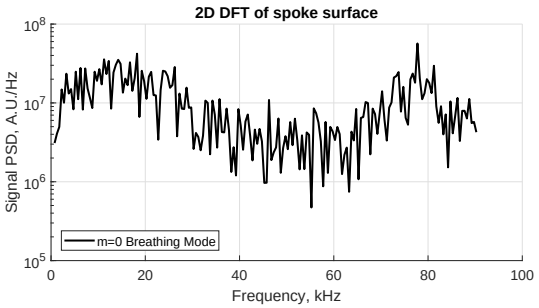

(e) HSC analysis for Inner in IM configuration

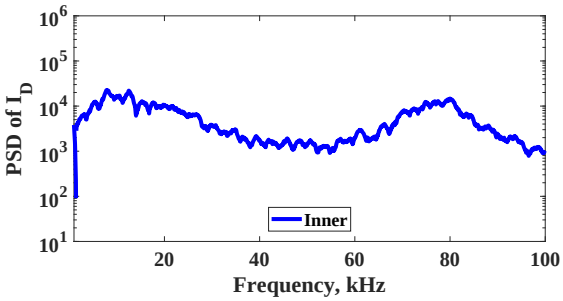

(f) HSDC analysis for Inner in IM configuration

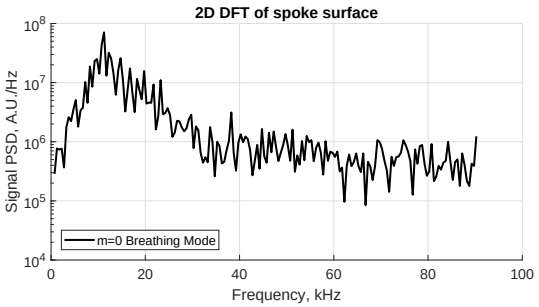

(g) HSC analysis for Middle in IM configuration

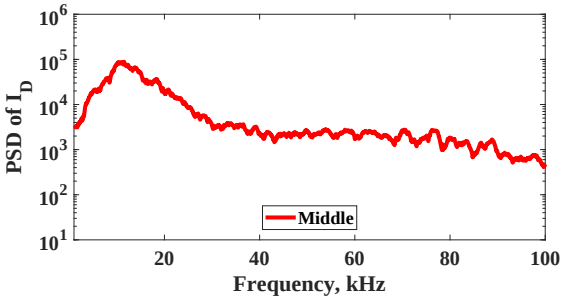

(h) HSDC analysis for Middle in IM configuration

Figure 8.41: Four examples of PSDs of the $m = 0$ oscillation mode captured with the HSC analysis (left) and the HSDC analysis (right) for the Inner and Middle channels operating at 400 V, $1.0 j_{ref}$. These results demonstrate that the HSC and HSDC analyses capture the same results in both single- and multi-channel operation.

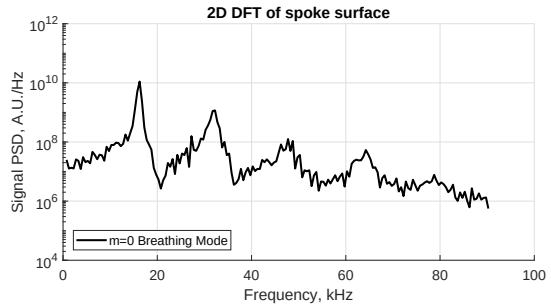

(a) HSC analysis for Inner in IMO configuration

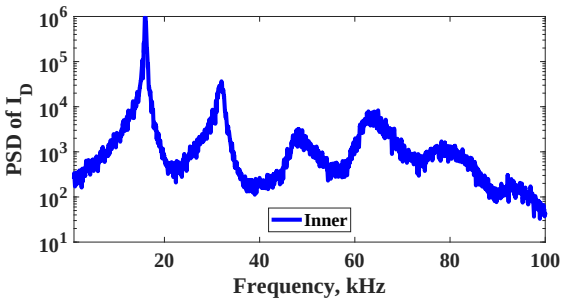

(b) HSDC analysis for Inner in IMO configuration

Figure 8.42: An example of PSDs of the $m = 0$ oscillation mode for an atypical condition (the Inner channel in the IMO configuration at 300 V, $1.0 j_{ref}$), demonstrating that the features were captured by both (a) HSC and (b) HSDC analysis.

8.6 Summary

We have shown that on average the X3 provides phenomenological efficiency values and trends that are very similar to the NASA-300M. This indicates that the X3 is operating as designed not only in terms of raw performance but also in terms of the various plasma mechanisms at work inside the discharge. These include:

- Anode efficiencies that increase with discharge voltage and range from 0.65–0.68 (±0.03).
- Charge utilization efficiencies that decrease slightly with discharge voltage from 0.98–0.97 (±0.005), which indicates an increasing population of multiply-charged ions with discharge voltage.
- Voltage utilization efficiencies that increase with discharge voltage from 0.93–0.96 (±0.01).
- Divergence-weighted current utilization efficiencies that increase with discharge voltage from 0.74–0.78 (±0.11).
- Divergence-weighted mass utilization efficiencies that decrease with discharge voltage. For the X3, these values ranged from 0.95–0.93 (±0.11), whereas for the NASA-300M

they ranged from 1.00–0.96 (±0.04). Discrepancies between thrust-derived and probe-derived anode efficiency for the NASA-300M, as well as results that were unphysically greater than unity for certain test conditions, made determination of the exact cause of the difference between the two thrusters difficult.

We then demonstrated that the various channel combinations of the X3 provided similar anode, charge, and voltage utilization efficiencies but that both the divergence-weighted mass and divergence-weighted current utilization efficiencies varied between combinations. The variation in η_{md} was explained once a calculation of the cross-channel ingestion effects was completed; with these effects accounted for, the values were much more similar across channel combinations. Preliminary calculations of these cross-channel effects, which will still be present in space and are not ground-test artifacts, indicate that they could account for average efficiency boosts on the order of 5%, although continued work is necessary to better characterize these effects and separate them from other facility-related effects.

High-speed camera analysis indicated that for nearly all conditions tested, the X3 operated in an oscillation mode characterized by a dominant breathing or $m = 0$ mode, a broad low-frequency breathing mode peak, and non-sinusoidal oscillations within the discharge that were de-coupled between channels operating simultaneously. This oscillation mode was present for both single- and multi-channel operation and a range of background pressures. This has important implications for future ground testing of the X3, as it indicates that the oscillation mode is not changing drastically across a large variation in facility background pressures. Performance can reasonably be compared between conditions if proper accommodations for facility neutral ingestion are made. The results from the high-speed camera analysis, including the small number of conditions that did not match the typical results, were also captured in high-speed discharge current analysis, a validation of the high-speed camera technique.

CHAPTER 9

Conclusions and Future Work

"There can be no thought of finishing [work on rockets], for "aiming at the stars," both literally and figuratively, is a problem to occupy generations, so that no matter how much progress one makes, there is always the thrill of just beginning."

– Robert Goddard in a fan letter to Orson Wells, 1932 [221]

"Predictions of when EP will be implemented have been repeatedly demonstrated to be speculative."

– Michael Patterson and James Sovey [136]

9.1 Conclusions

9.1.1 Summary of Work

The work performed here culminated in demonstrating that the X3 100-kW class nested Hall thruster provides state of the art performance across a throttling range of 300–500 V and 5–102 kW total power. In addition to that success, there are a number of technical achievements contained within this work.

By expanding an existing system mass and cost model to incorporate the unique benefits and features of NHTs, we demonstrated the system-level impact on mass and cost of implementing high-power NHTs on missions in excess of 500 kW total system power. NHTs are

capable of providing system mass savings on the order of 10%, which translate to similar cost savings by reducing the required number of thrusters to minimize system mass. Continued development work of NHTs will provide valuable additional information to refine this model further. Additionally, the model does not capture the wide throttling range of NHTs, which is another factor to consider when studying the full potential mission impact. The 10% savings can be viewed as a lower-bound estimate of the savings that NHTs are capable of providing.

The two low-power tests of the X3 that we performed at the University of Michigan identified a number of issues that may have precluded successful high-power operation. These included magnetic field optimization, thruster isolation, and propellant leak concerns. These issues were all addressed before high-power testing at NASA GRC, and the success of the GRC testing validated the repairs made. In addition to identifying these issues, the low-power tests demonstrated expected trends with swept magnetic field strength, which provided valuable information for subsequent high-power thruster performance mapping. We captured thermocouple measurements throughout the thruster operation, and we operated the thruster to a limited thermal steady-state condition at low power levels. These test results were later compared by Reilly to thermal modeling performed at JPL [93].

At GRC, we successfully characterized the X3 across a range of 5–102 kW total power, and in the process demonstrated new capabilities of Hall thrusters in total power, discharge current (a maximum of 247 A), and thrust (a maximum of 5.4 N). We measured the performance and studied the plasma plume of the thruster across the throttle table, which included test points for all seven available channel combinations. The performance results were compared to similar data from a series of high-power, single-channel Hall thrusters that were designed with similar techniques and principles as the X3: the NASA-300M, NASA-400M, NASA-457Mv1, and NASA-457Mv2. The results indicate that the X3 is operating with state of the art thrust-to-power ratio, specific impulse, and efficiency. The X3 also demonstrated similar trends with discharge voltage as the other high-power thrusters.

The plasma plume results indicated that the X3 operates with similar charge, voltage, current, and mass utilization as the NASA-300M single-channel Hall thruster, a 20-kW thruster developed with similar scaling and design techniques as the X3. These similarities are a validation of using these design techniques to scale to thruster powers in excess of 100 kW. Values measured include anode efficiencies between 0.65–0.68, charge utilization efficiencies of 0.97–0.98, and voltage utilization efficiencies of 0.93–0.96. Due to the lack of a swept Faraday probe to collect beam profile information, we calculated current and mass utilization efficiencies as "divergence-weighted" parameters because the divergence utilization efficiency could not be de-coupled from them. We performed similar calculations on Huang's published NASA-300M data [108]. Initial calculations indicated that the X3 matched the NASA-300M with divergence-weighted current utilization (with values ranging from 0.74–0.78) but that the X3 produced consistently lower values of divergence-weighted mass utilization (the X3's values ranged from 0.93–0.95 and were approximately 0.05 lower than those for the NASA-300M). Reasons for this difference were not clear from the data, and the uncertainty on both the X3 and NASA-300M data are large enough that the difference is not statistically significant.

In addition to the comparison of average behavior to that of the NASA-300M, we performed a comparison between the various channel combinations available in the X3. This comparison indicated that at 400 V (the largest set of data available in the experiment and the one selected for analysis here) the various combinations all exhibited comparable anode, charge utilization, voltage utilization, and divergence-weighted current utilization efficiencies, but that the divergence-weighted mass utilization efficiency was noticibly higher for multi-channel conditions over single-channel ones. However, a preliminary calculation of the effects of cross-channel neutral ingestion suggested that this effect may be responsible for this difference, and when this ingestion was separated from the divergence-weighted mass utilization efficiency, the values were comparable across all conditions. This effect was further shown to potentially be responsible for an increase in thruster efficiency of 0.01–0.09.

This effect is expected to occur on orbit, though further work is needed to de-convolve a number of effects captured together in the present analysis.

We made high-speed camera measurements to characterize the oscillatory behavior of the X3. This analysis demonstrated that each channel of the X3 operates in a similar oscillatory mode across the conditions tested. This mode is characterized by a dominant breathing or $m = 0$ mode, a broad low-frequency breathing mode peak, and non-sinusoidal oscillations within the discharge that were de-coupled between channels operating simultaneously. These measurements were the first of their kind for a 100-kW class Hall thruster and for discharge channels the size of the Middle and Outer channels of the X3, and the fact that these larger discharge channels demonstrated similar oscillatory behavior as the smaller Inner channel has important implications for future design of high-power Hall thrusters using existing scaling techniques. The power spectra of the high-speed camera analysis indicated that a high-frequency peak often existed during X3 operation, and results showed that this peak consistently occurred at values of 70–80 kHz for the Inner channel, 60–70 kHz for the Middle channel, and 40–60 kHz for the Outer channel, regardless of whether the channel was operating alone or with others. The source of this high-frequency content was not clear from the available data.

Discharge current oscillation measurements made with high-speed oscilloscopes and current measurement guns during thruster testing indicated that the strength of the oscillations changed dramatically for the Inner channel, and to a lesser extent the Middle channel, when operating with other channels as compared to operating alone, while the oscillation strength for the Outer channel was essentially unchanged between single- and multi-channel conditions. A limited exploration of this phenomenon with the Inner channel suggested that it was not related to the operation of the other channels or the increase in facility background pressure, though further exploration is certainly required. To study this phenomenon further, a correlation study was performed using various parameters of the high-speed camera analysis to see whether the difference in oscillation strength was correlated with any other parameters

of the discharge oscillations. This analysis studied the location of the aforementioned high-frequency peak, as well as the strength of both the low- and high-frequency peak, the width of those peaks, and the relative strength between the two, and found no correlation between any of these parameters and the difference in oscillation strength. Though this difference in oscillation strength requires further investigation, otherwise the oscillations across operating conditions of the X3 were very similar in character.

9.1.2 Implications of Work

This work has a number of important implications. The results from the X3 have demonstrated the capability of Hall thrusters to operate efficiently at 100 kW and 2300–2500 seconds specific impulse. A number of mission studies demonstrate that specific impulses of this magnitude are critical for near-term human exploration roles such as cargo and (at higher propulsion system power levels) crew transport. Previous 100-kW operation by the NASA-457Mv1 was at a lower total efficiency (0.58 versus 0.63 here) and at a much higher specific impulse (3460 seconds versus 2300–2500 seconds here), too high for most cargo and crew missions. Though the X3 is designed to be capable of specific impulses matching or exceeding the NASA-457Mv1, high-power, low-specific impulse operation has many mission applications. To illustrate this, we present in Figure 9.1 the same diagram that we showed in Figure 2.1, but with two markers added indicating the two 100-kW test points of the X3 (400 V/5.4 N/99 kW and 500 V/5.0 N/102 kW). As can be seen, the X3 has pushed the boundaries of Hall thruster operation.

These results also demonstrate the viability of using established Hall thruster scaling methodologies to design NHTs. The X3 demonstrated performance comparable to other thrusters designed with this technique, indicating both that efficient, state of the art Hall thruster discharge channels can be built with diameters of at least 80 cm, and that multiple discharge channels can be nested together and throttled to high power without issue. The X3 showed no indication of cathode coupling issues for the larger channels or during multi-

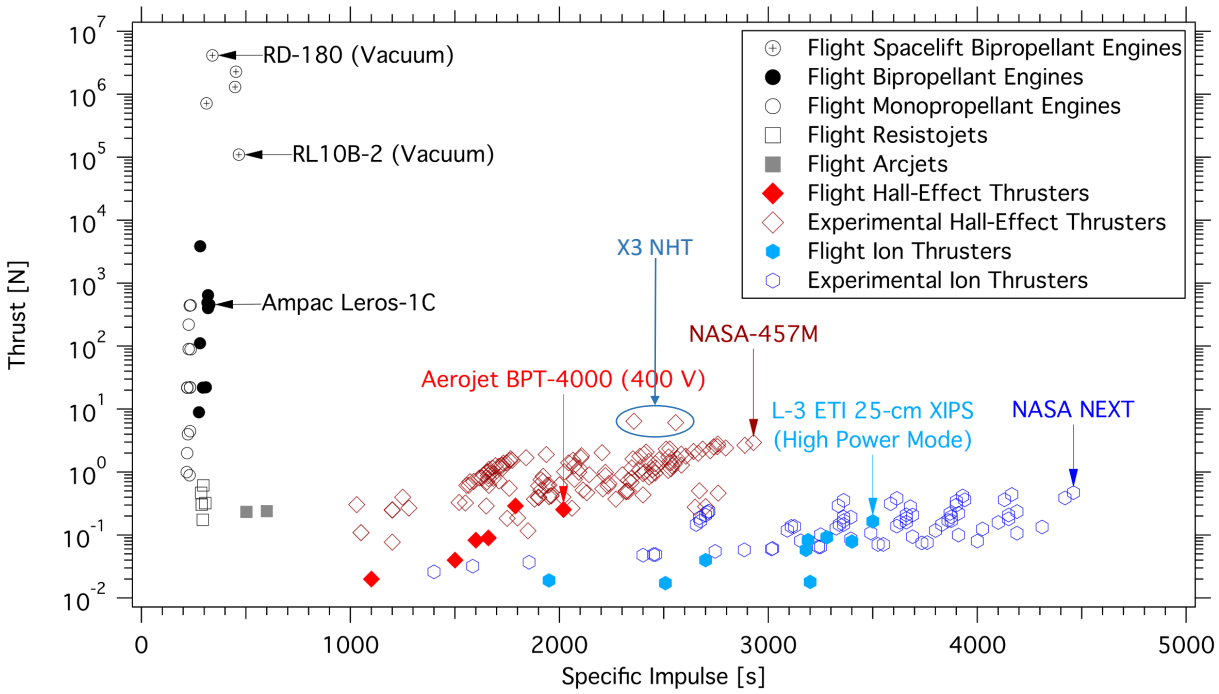

Figure 9.1: A plot of the thrust-versus-specific impulse trade space inhabited by various types of propulsion with two new markers indicating the two 100-kW operating points demonstrated with the X3 during this work.

channel operation, although work is needed to characterize this further.

This work also demonstrated the feasibility of ground-testing 100-kW class EP. We successfully operated at 240 mg/s xenon flow rate at 247 A discharge current while maintaining a facility background pressure of 4.2×10^{-5} Torr in VF5 at GRC. These pressures are low enough that significant thruster characterization is possible while maintaining a reasonable space-like environment. Additionally, the facility showed no indications of issues managing the thermal load from the thruster up to 102 kW. However, these results also demonstrate that as EP systems continue to grow in power level, new space simulation facilities will be required. Though reasonable characterization of the X3 is possible with existing facilities, thrusters larger than the X3 will begin to exceed facility pumping and heat-load capabilities. The need for improved testing facilities has become increasingly apparent with the continued increases in electric propulsion capability. NASA recently refurbished VF5's sister chamber Vacuum Facility 6 for EP testing and hope to add pumping speed capability to

that facility soon to match or perhaps exceed the capability of VF5 [222]. Looking farther ahead, Rodriguez recently presented a proposed design for a national space simulation facility that would offer ten times the pumping capability of VF5 [223]. Their work suggested an initial capital investment of $25M and an operating cost of $1.5M per year. As a metric of affordability, the authors noted that for the cost of a single launch at the current expected cost on NASA's Space Launch System (and thus one in-space demonstration of a high-power EP system), this facility could be built and then operated for 50 years. Both the need and the economic viability of developing new, advanced space simulation facilities for the next generation of EP technology are clear.

9.2 Recommendations for Future Work

The results from this work have suggested a number of future studies to be performed on the X3. We detail these briefly below. These studies are intended to highlight the vast experimental array available with a thruster as large and complex as the X3. One of the challenges in operating the X3 effectively is fixing enough variables to determine a meaningful result, and many of these suggestions for future work are intended to explore in more detail these parameters we did not vary or sweep during the work described above. Not all of these experiments specifically require the X3 (though some certainly do), so these suggestions can be viewed as future work toward the development of NHTs and high-power Hall thrusters in general.

9.2.1 Full-Power Operation

As has been noted, the X3 is designed to operate at powers up to 250 kW. Thus although the results here represent progress in increasing the demonstrated power, discharge current, and thrust capabilities of Hall thrusters, the X3 should be capable of even more. Brief testing at voltages in excess of 500 V at the end of the GRC campaign showed promising operation

from the thruster, and a small number of thruster improvements are being made based on lessons learned from that test. Future testing should be able to run the X3 at discharge voltages up to the 1000 V design point. A performance characterization similar to the one performed here for the higher-voltage, higher-power portion of the X3's throttle envelope is required to fully understand the device. These results may also provide more insight into some of the unexplained phenomena observed during these tests.

9.2.2 Magnetic Field Optimization for Multi-Channel Operation

As we noted throughout this work, magnetic field optimization during multi-channel operation of the X3 at high power was limited. We identified a number of performance trends in these data that indicate that a more detailed study of the magnetic field for these conditions is warranted. Further optimization of the field for these conditions may improve performance and/or identify performance losses due to multi-channel operation.

Another idea that was explored but ultimately not implemented was to perform fast I-B maps [178] on individual channels of the X3 while others are firing. No detailed study of the effect of sweeping one channel while another is firing has been performed yet. This kind of study could potentially provide valuable information on not only thruster operation and cathode coupling but also effects on the plasma impedance during multi-channel operation.

9.2.3 Low-Power Operation

Based on the evolved understanding of the X3 thruster and the optimal ways to operate it, we recommend returning to the low-power (30 kW) test conditions once again to evaluate whether the low performance observed during Florenz's work and the early work presented here can be improved upon. With the improved understanding of the magnetic field as well as the repaired anode welds, it is likely that the thruster now is capable of typical operation at these powers. The thruster operated as expected at a condition very close to that condition during the high-power tests (38 kW total power). One of the key aspects of the X3 is not

only its large total power but its wide throttling range, so verifying performance at the low end—and understanding why performance worsens, if it does—is an important part of fully characterizing the thruster.

9.2.4 Continued Magnetic Field Improvement

The X3 demonstrated state of the art efficiency during the high-power test at GRC. The thruster was typically operated at magnetic field current ratios optimized by modeling work done by Cusson [205]. However, for a small number of conditions, an additional magnetic field topology optimization was performed in which the current ratio was swept to minimize discharge current. The results from these two test conditions (for the Middle channel at 500 V) provided some of the best efficiency of the entire test matrix, indicating that more work may be necessary, including a verification that the model used to develop the settings matches the thruster as it was built.

9.2.5 Beam Profile Characterization

Although the probe suite used in the GRC testing provided critical information about the plasma plume of the X3 and thus the thruster's operation, a key diagnostic missing from the campaign was a swept Faraday probe to characterize the beam profile and divergence angle. In addition to providing a measurement of beam current I_b, a beam profile measurement is important for spacecraft integration. Measuring the beam profile across all channel combinations of the X3 will provide valuable information about how the beam profiles from each channel are merging and interacting in the plume region of the thruster. This analysis should be performed in both the near-field [107] and far-field plume [108] of the X3.

In addition to a swept-Faraday probe analysis, measurement of the ion velocity profile off-centerline will offer insight into the charge-exchange populations in the beam [219]. The role that charge-exchange processes play in the characterization and operation of the X3 in ground-test facilities is of special importance due to the relatively large background pressures

at which the thruster operates.

9.2.6 Continued Investigation of Oscillation Behavior

One of the biggest questions uncovered by this work concerns why the discharge current oscillations change so significantly in magnitude between the single-channel and multi-channel operation for the Inner and Middle channels. A number of different experiments are suggested to explore this phenomenon further. To determine whether the increase in oscillations is simply due to the strong wear bands already established on the discharge channel ceramics, new channel ceramics could be installed in the thruster and the test matrix repeated. Additional suggested tests include sweeping the magnetic ratio of each channel to study the effect on the oscillation strength and laser-induced fluorescence (LIF) measurements to study potential changes in the location of the acceleration zone within the channel between the single- and multi-channel magnet settings.

Another aspect to this study would be a more focused study of the cathode oscillations and the way they interact and coexist with the discharge channel oscillations. This should involve both high-speed camera analysis and high-speed discharge current measurements. A correlation study between the cathode and channel oscillations can be performed to study how oscillations—and perhaps thruster transients—transit between the discharge and the cathode in a complex multi-channel system.

9.2.7 Laser-Induced Fluorescence

In addition to the study above, a more thorough LIF campaign would provide important insight into the operation of the X3 and of high-power Hall thrusters in general. None of the NASA high power thrusters has been subject to a LIF study, and the highest power Hall thruster to be studied with LIF to date is the 12.5-kW NASA HERMeS thruster [224]. A LIF study of the X3 would provide information on how the acceleration region in each of the channels differs and varies between conditions. Because of how similar the channel

cross-sections of the X3 are, any differences in these measurements would be important from a thruster design perspective.

A study like this does present significant technical challenges to overcome. The larger beam of the X3 will make properly locating and shielding optics more difficult. Also, typical LIF is performed with the optics stationary in the vacuum facility (due to the difficulty involved in alignment) and the thruster on motion stages such that it can be moved to interrogate different points within the discharge. Safely and effectively mounting a 230-kg thruster onto motion stages will take significant design work.

9.2.8 Channel Coupling Study

To better understand potential channel coupling in the X3, the experiment performed by Georgin [187] and Cusson [204] could be repeated on the X3. In doing so, the X3 would be operated in a more similar manner between channel combinations to limit the differences between conditions (e.g., the same electromagnet settings for all conditions, which was not true for the present study). Results from this type of study would expand the NHT knowledge base as well as add important knowledge to Hall thruster scaling and design principles. If the results appear similar between the X2 and the X3, the X3 offers a testbed for further study and physical insight with the additional channel. In addition to testing the three channels together and configurations with two channels adjacent to each other, the X3 allows for operation with two channels separated by a larger distance (the IO configuration). Measuring the differences in effects between the various two-channel configurations, and how they sum in the IMO configuration, will offer further insight into the channel coupling mechanisms in NHTs. And if the results do not match the X2, as these preliminary results suggest may be the case, once again physical insight gained through the expanded available test matrix has the potential to offer great insight into the differences between the two thrusters. This type of knowledge is critical for the continued development toward flight readiness of NHTs.

9.2.9 Background Pressure Study

As has been discussed in detail in this work, the large throttling matrix of the X3 imposes a large range of test facility background pressures during thruster operation and characterization. Previous work has shown that a 6-kW magnetically shielded Hall thruster was significantly less susceptible to the effects of changing background pressure with a centrally-mounted cathode than with an externally mounted one [111]. This type of study should be repeated on the X3. The cost of xenon is such that maintaining a fixed background pressure during all X3 operation would be prohibitively expensive. However, performing this type of test a single time for a representative set of conditions would both eliminate concerns about changes in background pressure affecting thruster performance measurements and potentially offer further insight into how the thruster is sharing propellant during multi-channel operation. By sweeping the background pressure for each channel combination and studying how much propellant is ingested from the facility (by either operating the thruster with constant discharge current and monitoring the change in necessary flow rate, or by operating the thruster with constant anode flow rate and measuring the change in discharge current), an understanding of how much each channel ingests from the facility can be gathered. From these measurements, the remaining deficit can be more accurately accounted for and related to cross-channel neutral sharing.

9.2.10 Extended Total Cathode Flow Fraction Study

As we present in Appendix A, we operated the X3 stably down to a total cathode flow fraction (TCFF) of 3% at a single example operating condition. However, this test was by no means conclusive. A further study of low-TCFF operation of the X3 should be completed. This study should encompass multiple discharge voltages, current densities, and channel combinations to better map out the thruster response across the throttle table. Additionally, a stability study should be completed, including a detailed study of the effects of lighting an additional channel at a low TCFF. Thruster performance should also be thoroughly mapped

across the range of TCFF values.

9.2.11 Study of Cathode Downstream Neutral Injection

In addition to the study of low TCFF, a study should be completed that looks at the role that the cathode downstream neutral injection plays in the operation of the X3. No detailed study of X3 operation with varying flow split has been performed to date. By studying this with techniques including plasma plume diagnostics and thruster performance measurements, insight can be gained on the most optimal way to operate the JPL cathode with the X3 across the range of operating conditions. This test matrix should span a wide range of discharge currents and include different channel combinations and magnetic field settings to identify the effect that various parameters have on cathode operation and cathode/discharge coupling. At each test point, multiple TCFF values can be tested, and for each TCFF, the flow can be variably split between the injectors and the cathode center. As discussed in this work, we suspect that the downstream neutral injection from the cathode may play a role in the way the large Middle and Outer channels of the X3 couple with the cathode; this study would help to illuminate that mechanism better. This has broader implications for high-current cathodes, NHTs, and high-power Hall thrusters in general. For instance, the NASA-400M reported requiring much larger cathode flow rates than what is typical for Hall thrusters to maintain reasonable levels of thruster/cathode coupling [76]. Downstream neutral injection may show itself to be a means of surmounting these issues with TCFFs well below the 34% required by the NASA-400M.

9.2.12 Thruster Body Electrical Configuration Study

To date, the X3 has always been operated in a grounded configuration. The entire outer surface of the thruster is coated in a dielectric aluminum oxide coating, and results from the high-power test at GRC indicated that less than 2 A of current was collected by the body at discharge currents as high as 250 A, indicating that the coating is working in isolating the

thruster body from the plasma. However, a thorough study similar to that performed on the HERMeS thruster [179] should be performed just to verify that the electrical configuration of the body does not affect thruster performance. Plume measurements like those by Huang on the HERMeS thruster would shed more light into any effects from the body electrical configuration [219].

9.2.13 Other X3-based Experiments

The X3, with its large size and unique three-channel nested lens magnetic field topology, provides an opportunity for study of Hall thruster physics in greater detail than previously. One potential area of study is that of the effect of magnetic lenses between the cathode and the discharge. The X3 provides a means to study not only a single lens between the discharge and the cathode like the X2, but two lenses (for the Outer channel) and a lens between two firing channels (in the IO configuration with the M magnets on). A study like this should measure near-field plume properties, thruster performance, and perhaps even collect LIF measurements to study changes in the plasma within the discharge channel.

One other potential area of study is that of the role of dormant anodes in the electrical and plasma circuit of the X3. Normally, dormant anodes were left electrically floating, and typically floated to values on the order of 10 V above ground during operation. However, in some cases the anode of a dormant channel would float as high as 50 V above ground. Usually, anodes between the firing channel and the cathode were more susceptible to these effects than those outside of the firing channel. However, no thorough study was completed. An electrical study, similar to that performed on the thruster body with the HERMeS thruster [179], could be performed with the dormant anodes to study what the effect on thruster performance and operation is from leaving them floating, grounding them, or tying them to the cathode. This study could involve the thruster body electrical configuration as well.

A final thought for an experiment would be to use the dormant anodes as swept Langmuir

probes or use them for some other diagnostic means. Although they are large and these types of sweeps would likely be perturbing, they may shed light into the character of the plasma in these regions of the thruster. An alternative that would likely be more successful and less perturbative is outfitting the X3 with wall-mounted Langmuir probes in a style similar to those used by Shastry on various Hall thrusters [225–227]. This would be doubly beneficial because they could be used not only to characterize the near-wall plasma of a given channel when it was firing, but could also be used to characterize the plasma within that channel when it is dormant—and how that plasma changes with variations in operating condition and channel electrical configuration. This type of study, though technically challenging for a thruster the size of the X3 (and with BN segments like the X3), could provide very valuable characterization of the plasma between a firing discharge channel and the cathode.

APPENDIX A

Operation at Low Total Cathode Flow Fraction

Typically, the cathode flow in Hall thrusters is set as a fraction of the anode flow, often between 5–10% but in excess of 30% for some thrusters [76, 100]. As discussed, Hall thrusters are traditionally scaled to high power by maintaining discharge current densities and increasing the exit area and thus the total discharge current [11]. As such, anode flow rates typically scale linearly with thruster power level. For the X3, this translates to total anode flow rates of approximately 250 mg/s of xenon at 250 A operation. However, hollow cathodes do not require the same linear flow rate increase with current, and in fact, high flow rates (e.g., 25 mg/s for 10% TCFF at 250 A operation) can disrupt cathode operation. Hollow cathodes require internal pressures on the order of 1 Torr to operate effectively and to experience uniform insert heating [97].

The JPL-designed high-current cathode features a unique propellant design utilizing external gas injectors [174]. Recall that for this cathode, we use the total cathode flow fraction (TCFF), which includes both the cathode center and injector flow as described in Equation 4.1. These external injectors serve two purposes. First, they provide a secondary flow path for cathode flow in excess of what is necessary to maintain internal pressures of 1 Torr. Secondly, they reduce energetic ions, a result that has significant implications for cathode lifetimes. Work at JPL has shown that this cathode can operate at nominal xenon flow rates of 16 sccm through the cathode and 20 sccm through the injectors up to discharge currents of

250 A. That work, as well as work done with the H6 6-kW thruster [100], has suggested that it may be possible to operate the X3 with the JPL cathode at TCFFs much smaller than the traditional 7–10%, especially at high-current conditions. However, to date, all X3 testing has been performed with TCFFs of 7–10%. This TCFF was split with 16 sccm through the cathode and the remainder through the injectors.

During this test campaign we demonstrated for the first time X3 operation with the JPL cathode at reduced cathode flow fractions. For the IM condition at 300 V and 27.2 kW total discharge power (26.3 A on the Inner channel and 64.2 A on the Middle), we reduced the TCFF from 7% to 3% in 1% increments, leaving $\dot{m}_c$ at the JPL-recommended 16 sccm and reducing $\dot{m}_{inj}$. At each TCFF setting, we measured thruster performance with the thrust stand as well as the most probable voltage using the downstream RPA and the plasma potential using the downstream Langmuir probe described in Chapter 4.

Figure A.1 shows the trends of the cathode to ground voltage, coupling voltage, thrust, anode efficiency, and total efficiency as a function of TCFF. We found that cathode to ground voltage only dropped 1.5 V more negative in going from 7% TCFF to 3% TCFF. As the figure shows, over 1 V of this drop occurred between 4% and 3% TCFF. It is worth noting that 3% TCFF for this thruster operating condition resulted in an injector flow rate below the recommended 20 sccm from previous JPL work. Following work performed by Goebel on the H6 6-kW thruster [100], coupling voltage is defined as:

$$V_{coupling} = V_p + |V_{cg}| \tag{A.1}$$

where V_p is the downstream plasma potential as measured by the planar Langmuir probe and V_{cg} is the cathode to ground voltage as measured inside the thruster telemetry breakout box. Coupling voltage is an inefficiency or loss in Hall thrusters because it represents the portion of the applied discharge voltage unavailable for ion acceleration. It can be seen that coupling voltage increases with decreasing TCFF. This increase comes both from a decrease

in cathode to ground voltage and an increase in plasma potential.

Thrust dropped approximately 17 mN in going from 7% TCFF to 3% TCFF, which is approximately 1% of the thrust value. Half of this drop occurred in going from 7% to 4%, and half occurred in going from 4% to 3%, which corresponded to the 1 V drop more negative in cathode to ground voltage. Thrust uncertainty at this test point was approximately ±43 mN, corresponding to 2.6% of full scale. Thus, the recorded drop in thrust was within the uncertainty of the measurement.

We found that anode efficiency increased slightly between 7% and 5% TCFF before decreasing from 5% to 3% TCFF. Anode efficiency values fell within a range of approximately 0.01. Due to the thrust uncertainty, the anode efficiency value had an uncertainty of ± 0.03, so these changes fell within the uncertainty of the measurements. Finally, total efficiency increased 0.013 in going from 7% to 3% TCFF. This demonstrates that the efficiency gains from the lower cathode propellant flow rates overcame the inefficiencies imposed on the thruster discharge due to these lower flow rates.

Voltage utilization efficiency as a function of TCFF is presented in Figure A.2. We found that, much like the other parameters, η_V held nearly constant for TCFFs from 7% to 4% at approximately 0.93 before experiencing a drop of nearly 0.03 in going from 4% TCFF to 3%.

We collected high-speed measurements of the discharge current using the current guns and oscilloscopes described in Chapter 4. These measurements will indicate whether dropping the cathode flow fraction has induced a mode transition in the thruster. Figure A.3 plots the P2P value of the discharge current oscillations for each channel as a fraction of the mean discharge current of that channel against TCFF. As can be seen, the oscillation amplitudes of both channels decrease slightly with decreasing TCFF, but stay within 0.02 throughout the sweep. Work has indicated that mode transitions in Hall thrusters are typically accompanied by large changes in oscillation amplitude [48, 106]. Therefore, these results indicate that the mode of operation of both channels remained constant during the TCFF sweep.

This test demonstrated that the X3 NHT can be operated with the JPL cathode stably at

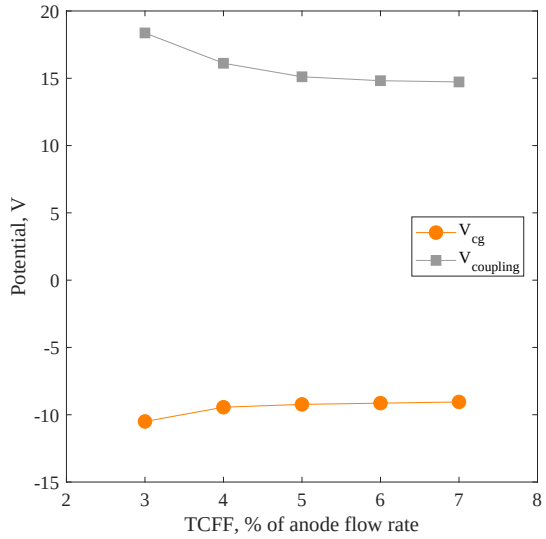

(a) Cathode to ground voltage

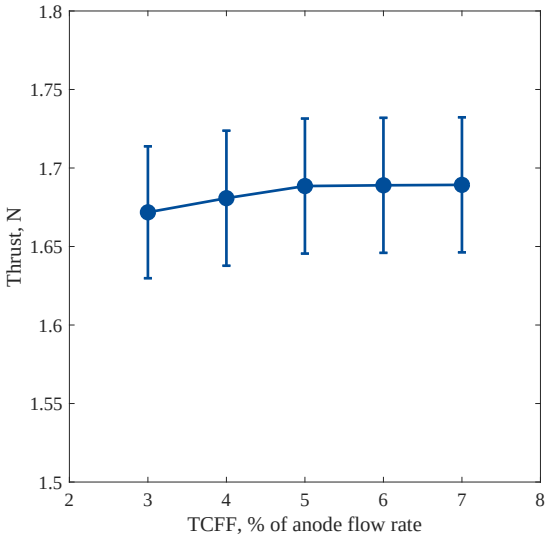

(b) Thrust

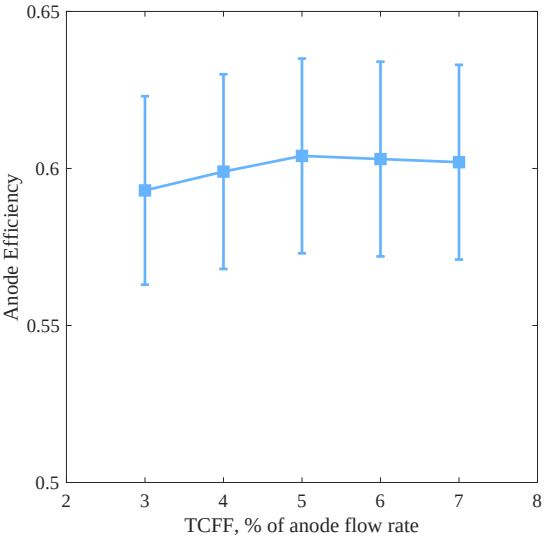

(c) Anode efficiency

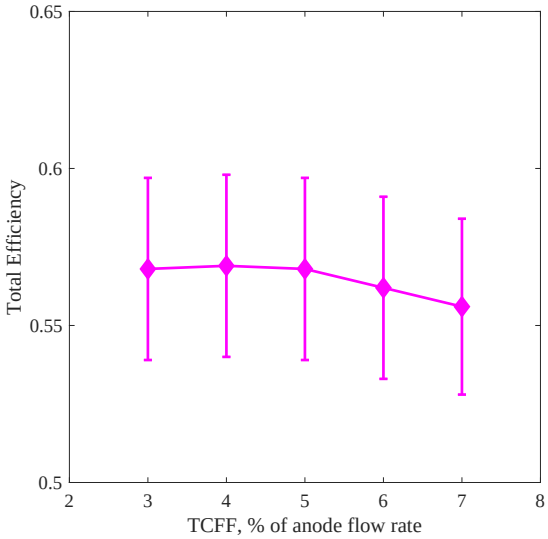

(d) Total efficiency

Figure A.1: Various thruster parameters plotted as functions of total cathode flow fraction: (a) cathode to ground voltage, (b) thrust, (c) anode efficiency, and (d) total efficiency.

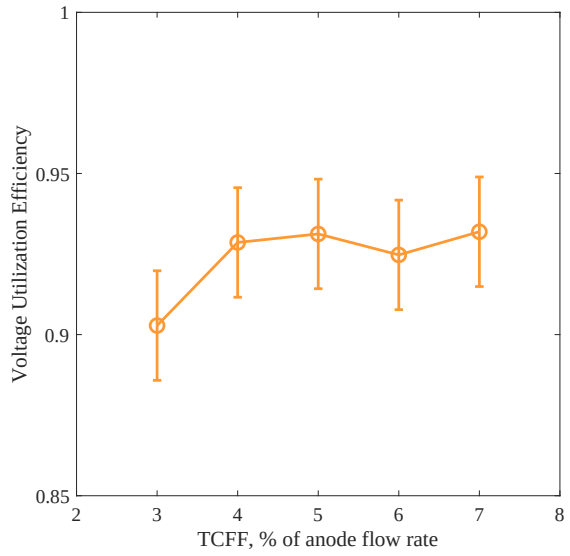

Figure A.2: Voltage utilization efficiency as a function of total cathode flow fraction.

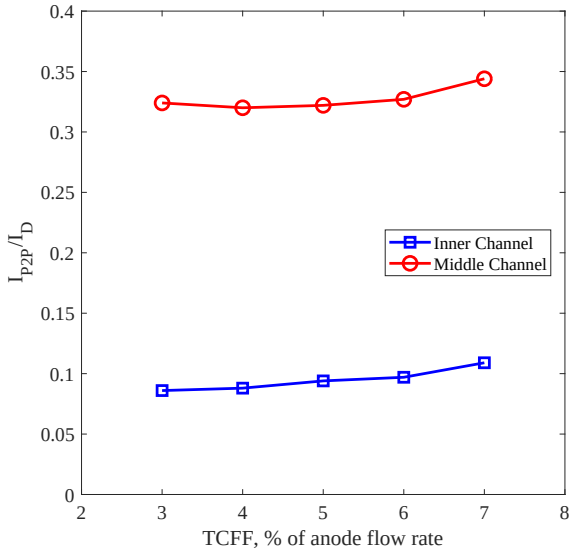

Figure A.3: Discharge current oscillations, characterized as the ratio of the peak to peak value to the mean value, as a function of total cathode flow fraction for the Inner and Middle channels of the X3.

TCFFs of as low as 3% without major impact to thruster operation. The slight inefficiencies encountered in the discharge due to these lower cathode flow rates are overcome by the total system efficiency gain from using less propellant. Low-TCFF operation can offer increased system efficiency as well as improved cathode lifetime, both of which are of interest to the XR-100 NextSTEP program, and can do so with little impact on thruster operation.

However promising these initial results are, significant further work is necessary to verify these trends with other thruster channel combinations and operating points. There are many questions left unanswered by this test. During another portion of the test, the thruster was operated in the MO condition at 5% TCFF when the Inner channel was started (to enter IMO operation). In doing so, the Inner channel experienced P2P oscillations in excess of 110% its mean discharge current and was visibly flickering through the vacuum chamber view port. The M and O channels appeared unaffected on the high-speed discharge current monitoring and through the view port. Adjusting the TCFF up to 7% settled these oscillations to values closer to nominal for that channel in multi-channel operation, near 80% its mean discharge current, and the visible flickering stopped.

This and other anecdotal experiences while operating the X3 during this campaign led us to the conclusion that a higher TCFF is likely necessary during changes in operation (at least changes in channel configuration if not large changes in voltage and/or current). Higher TCFF, especially on this cathode where this translates to more cold neutral gas being deposited at the exit of the cathode, helps to quench oscillations and helps the cathode couple to the anodes. The shock of lighting a new channel may cause issues with that anode coupling to the cathode which is already coupled to as many as two other channels. However, once the coupling is established between the new channel and the cathode, it is possible that the TCFF could again be dropped without issue.

However, this still leaves the question of thruster stability at low TCFF. The test we describe here consisted of operation at each TCFF value of less than 10 minutes, and during the entire test the thruster experienced no major discharge transients. It is possible that at

low TCFF, a spark event or other momentary change in thruster behavior could de-stabilize the thruster discharge for one or more channels. These questions, as well as those of the role the facility [228] and the electrical configuration of the thruster body [179] (grounded as it was here, floating, or cathode-tied) play in the phenomena observed, are a subset of broader questions being answered by the community at large. Understanding in more detail the source(s) of Hall thruster discharge oscillations and the role that ground-test facilities play in the phenomena observed during thruster operation will provide further insight into questions of cathode coupling and low-TCFF operation.

APPENDIX B

Low-Power Throttle Table

Table B.1: Low-power throttle table for X3 performance measurements at PEPL.

$V_{d,I}$ V	$I_{d,I}$ A	$V_{d,M}$ V	$I_{d,M}$ A	$V_{d,O}$ V	$I_{d,O}$ A	$P_{d,t}$ kW	T mN	$I_{sp,a}$ s	η_a –	V_{cg} V	p_b μTorr
300	13.6	0	0	0	0	4.2	278	1850	0.605	-12.8	4.4
0	0	300	31.8	0	0	9.5	385	1430	0.284	-11.8	9.6
0	0	0	0	300	54.7	16.4	611	1160	0.212	-16.7	29
300	13.3	301	31.9	0	0	13.6	809	2280	0.665	-13.0	13
299	13.6	0	0	298	54.9	20.4	1136	1800	0.490	-13.2	39
0	0	298	31.7	299	55.2	25.9	1162	1510	0.332	-12.5	66
300	13.7	300	32.6	300	55.2	30.4	1518	1840	0.450	-6.8	94

APPENDIX C

High-Power Throttle Table

Table C.1: 300 V throttle table for the X3 performance measurements.

$V_{d,I}$ V	$I_{d,I}$ A	$V_{d,M}$ V	$I_{d,M}$ A	$V_{d,O}$ V	$I_{d,O}$ A	$P_{d,t}$ kW	T N	$I_{sp,a}$ s	$I_{sp,t}$ s	η_a –	η_t –	V_{cg} V	p_b μTorr
300.2	16.3	0	0	0	0	4.9	0.35	1950	1820	0.69	0.64	-14.0	4.6
0	0	296.4	39.6	0	0	11.7	0.82	1950	1830	0.67	0.62	-11.7	10.2
0	0	0	0	300.3	68.6	20.6	1.35	2070	1950	0.67	0.62	-11.9	15.9
300.0	16.5	301.0	39.7	0	0	16.9	1.15	1930	1800	0.65	0.60	-9.6	13.8
300.0	16.0	0	0	300.4	68.9	25.5	1.63	2060	1940	0.65	0.60	-10.1	18.5
0	0	303.3	39.0	300.3	70.0	32.9	2.12	2100	1870	0.66	0.58	-10.3	23.0
299.8	16.5	303.1	39.4	300.3	70.1	37.9	2.38	2040	1920	0.63	0.58	-9.5	25.5
299.5	26.4	0	0	0	0	7.9	0.56	2050	1920	0.71	0.66	-10.9	6.6
0	0	292.6	63.8	0	0	18.7	1.28	2060	1930	0.70	0.64	-10.5	14.5
0	0	0	0	300.2	110.2	33.1	2.11	2150	2010	0.67	0.62	-9.9	22.9
299.3	27.5	300.4	66.1	0	0	28.1	1.85	2100	1960	0.68	0.63	-9.0	19.7
299.3	25.4	0	0	300.1	110.5	40.8	2.58	2110	2030	0.65	0.62	-12.7	26.0
0	0	305.1	63.1	300.0	110.6	52.4	3.30	2102	2002	0.65	0.61	-12.8	31.9
299.1	25.7	304.6	62.4	300.0	110.0	60.0	3.74	2070	1970	0.63	0.60	-14.4	36.2
298.9	32.9	0	0	0	0	9.84	0.66	2040	1910	0.67	0.62	-9.8	7.5
0	0	288.1	79.84	0	0	23.00	1.48	1950	1820	0.62	0.57	-9.6	17.0
0	0	0	0	299.3	137.5	41.2	2.55	2150	2010	0.65	0.61	-9.9	27.2
298.4	33.6	300.1	78.6	298.1	138.5	74.9	4.64	2130	2020	0.65	0.61	-11.4	42.2

Table C.2: 400 V throttle table for the X3 performance measurements.

$V_{d,I}$ V	$I_{d,I}$ A	$V_{d,M}$ V	$I_{d,M}$ A	$V_{d,O}$ V	$I_{d,O}$ A	$P_{d,t}$ kW	T N	$I_{sp,a}$ s	$I_{sp,t}$ s	η_a –	η_t –	V_{cg} V	p_b μTorr
400.4	16.2	0	0	0	0	6.5	0.39	2190	2050	0.64	0.59	-13.7	4.3
0	0	400.1	39.8	0	0	15.9	0.93	2210	2060	0.63	0.59	-12.6	9.7
0	0	0	0	400.0	68.9	27.6	1.58	2360	2200	0.66	0.61	-10.7	15.9
400.2	16.9	0	0	399.7	68.8	34.2	1.94	2380	2220	0.66	0.61	-10.1	18.5
0	0	394.8	38.9	399.6	69.7	43.2	2.45	2350	2200	0.65	0.61	-10.4	23.0
400.1	16.5	394.4	39.3	399.6	68.8	49.5	2.77	2310	2160	0.63	0.59	-9.8	25.6
399.77	26.4	0	0	0	0	10.6	0.61	2290	2140	0.65	0.60	-11.48	6.3
0	0	399.5	63.7	0	0	25.5	1.49	2360	2210	0.68	0.63	-10.6	14.1
0	0	0	0	399.1	106.9	42.7	2.40	2450	2330	0.68	0.64	-10.9	22.0
399.6	25.9	399.5	63.4	0	0	35.7	2.07	2420	2270	0.69	0.64	-8.8	19.0
399.7	25.0	0	0	399.8	109.0	53.6	2.98	2450	2290	0.67	0.62	-9.7	26.6
0	0	393.0	64.6	399.7	109.8	69.3	3.94	2460	2300	0.69	0.64	-10.8	32.6
399.5	25.7	401.0	64.1	399.6	109.9	79.9	4.49	2440	2240	0.67	0.61	-10.8	37.0
398.6	33.8	402.7	76.1	398.2	136.9	98.6	5.42	2470	2340	0.67	0.63	-11.3	41.8

Table C.3: 500 V throttle table for the X3 performance measurements.

$V_{d,I}$ V	$I_{d,I}$ A	$V_{d,M}$ V	$I_{d,M}$ A	$V_{d,O}$ V	$I_{d,O}$ A	$P_{d,t}$ kW	T N	$I_{sp,a}$ s	$I_{sp,t}$ s	η_a –	η_t –	V_{cg} V	p_b μTorr
500.7	16.5	0	0	0	0	8.2	0.43	2460	2300	0.64	0.59	-12.9	4.4
0	0	500.9	39.3	0	0	19.7	1.08	2590	2420	0.70	0.64	-12.9	9.9
0	0	0	0	500.4	67.4	33.7	1.77	2770	2580	0.71	0.66	-12.2	14.7
500.5	17.0	500.6	39.1	0	0	28.0	1.37	2450	2290	0.59	0.54	-9.1	12.6
500.4	16.8	0	0	500.3	67.5	42.2	2.19	2740	2560	0.70	0.64	-9.9	17.8
0	0	508.6	37.3	500.6	68.8	53.4	2.62	2790	2590	0.67	0.61	-12.1	20.7
500.2	18.0	499.2	38.1	500.6	68.6	62.4	3.28	2760	2580	0.71	0.66	-10.2	25.1
500.1	26.1	0	0	0	0	13.1	0.70	2680	2500	0.71	0.66	-10.9	6.2
0	0	500.4	62.3	0	0	31.2	1.69	2710	2530	0.72	0.67	-10.8	14.0
0	0	0	0	499.6	109.5	54.7	2.77	2840	2650	0.70	0.65	-10.4	21.9
499.9	26.3	499.8	63.7	0	0	45.0	2.34	2720	2540	0.69	0.64	-8.7	18.5
499.8	26.2	0	0	500.7	109.2	67.8	3.38	2800	2610	0.68	0.63	-9.6	25.3
0	0	504.3	63.8	499.4	110.0	87.1	4.34	2740	2560	0.67	0.62	-10.9	32.2
499.4	28.2	505.3	63.2	499.6	110.1	101.0	5.03	2750	2570	0.67	0.63	-10.3	35.9

BIBLIOGRAPHY

[1] Liang, R., *The Combination of Two Concentric Discharge Channels into a Nested Hall-Effect Thruster*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2013.

[2] Goebel, D. M. and Katz, I., *Fundamentals of Electric Propulsion: Ion and Hall Thrusters*, Vol. 1, John Wiley & Sons, 2008.

[3] Hofer, R. R., Jankovsky, R. S., and Gallimore, A. D., "High-specific impulse Hall thrusters, part 1: Influence of current density and magnetic field," *Journal of Propulsion and Power*, Vol. 22, No. 4, 2006, pp. 721–731.

[4] Florenz, R. E., *The X3 100-kW Class Nested-Channel Hall Thruster: Motivation, Implementation, and Initial Performance*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2013.

[5] Brophy, J. R., Gershman, R., Strange, N., Landau, D., Merrill, R., and Kerslake, T., "300-kW Solar Electric Propulsion System Configuration for Human Exploration of Near-Earth Asteroids," *47th Joint Propulsion Conference, AIAA-2011-5514*, American Institute of Aeronautics and Astronautics, San Diego, CA, Aug. 2011.

[6] Hofer, R. R. and Randolph, T. M., "Mass and cost model for selecting thruster size in electric propulsion systems," *Journal of Propulsion and Power*, Vol. 29, No. 1, 2012, pp. 166–177.

[7] Weir, A., *The Martian*, Broadway Books, New York, 2014.

[8] Morozov, A. I., "The conceptual development of stationary plasma thrusters," *Plasma Physics Reports*, Vol. 29, No. 3, 2003, pp. 235–250.

[9] De Grys, K., Fisher, J., Wilson, F., Beal, B., Dimicco, J., and Khayms, V., "4.5 KW Hall Thruster System Qualification Status," *40th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, AIAA-2004-3603*, American Institute of Aeronautics and Astronautics, Fort Lauderdale, FL, July 2004.

[10] Manzella, D., Jankovsky, R., and Hofer, R., "Laboratory Model 50 kW Hall Thruster," *38th Joint Propulsion Conference, AIAA-2002-3676*, American Institute of Aeronautics and Astronautics, Indianapolis, IN, July 2002.

[11] Manzella, D., *Scaling Hall Thrusters to High Power*, Ph.D. Dissertation, Stanford University, Stanford, CA, 2005.

[12] Sekerak, M. J., Longmier, B. W., Gallimore, A. D., Huang, W., Kamhawi, H., Hofer, R. R., Jorns, B. A., and Polk, J. E., "Mode Transitions in Magnetically Shielded Hall Effect Thrusters," *50th Joint Propulsion Conference, AIAA-2014-3511*, American Institute of Aeronautics and Astronautics, Cleveland, OH, July 2014.

[13] Clark, J. D., *Ignition!: An informal history of liquid rocket propellants*, Rutgers University Press, 1972.

[14] Stuhlinger, E., *Ion Propulsion for Space Flight*, McGraw-Hill New York, 1964.

[15] Wertz, J. R. and Larson, W. J., *Space Mission Analysis and Design*, 3rd ed., 1999.

[16] Bate, R. R., Mueller, D. D., and White, J. E., *Fundamentals of astrodynamics*, Courier Corporation, 1971.

[17] Sutton, G. P. and Biblarz, O., *Rocket Propulsion Elements*, John Wiley & Sons, 2010.

[18] Goddard, R. H., *Robert H. Goddard: An Autobiography*, Astronautics, 1959.

[19] Oberth, H., *Man Into Space*, Harper and Row Publishers, New York, 1957.

[20] Jahn, R. G., *Physics of Electric Propulsion*, Dover Publications, New York, 1st ed., 2006.

[21] Welander, B., Carpenter, C., de Grys, K., Hofer, R. R., Randolph, T. M., and Manzella, D. H., "Life and operating range extension of the BPT-4000 qualification model hall thruster," *42nd Joint Propulsion Conference, AIAA-2006-5263*, American Institute of Aeronautics and Astronautics, Sacramento, CA, July 2006.

[22] Garner, C. E., Rayman, M. D., and Brophy, J. R., "In-Flight Operation of the Dawn Ion Propulsion System-Arrival at Ceres," *34th International Electric Propulsion Conference, IEPC-2015-88*, Electric Rocket Propulsion Society, Kobe-Hyogo, Japan, July 2015.

[23] Shastry, R., Herman, D. A., Soulas, G. C., and Patterson, M. J., "Status of NASA's Evolutionary Xenon Thruster (NEXT) Long-Duration Test as of 50,000 h and 900 kg Throughput," *33rd International Electric Propulsion Conference, IEPC-2013-121*, Electric Rocket Propulsion Society, Washington D.C., Oct. 2013.

[24] Schmidt, T. D., Seboldt, W., and Auweter-Kurtz, M., "Flexible piloted Mars missions using continuous electric propulsion," *Journal of spacecraft and rockets*, Vol. 43, No. 6, 2006, pp. 1231–1238.

[25] Bussard, R. W. and De Lauer, R. D., *Nuclear rocket propulsion*, McGraw-Hill, 1958.

[26] Hohmann, W., *The Attainability of Celestial Bodies*, NASA Technical Translation, NASA (originally R. Oldenbourg Verlag, Munich, Germany), Washington D.C., 1960.

[27] Lovell, J. and Kluger, J., *Apollo 13*, Houghton Mifflin Harcourt, 2006.

[28] Polk, J. E., Brinza, D., Kakuda, R. Y., Brophy, J. R., Katz, I., Anderson, J. R., Rawlin, V. K., Patterson, M. J., Sovey, J., and Hamley, J., "Demonstration of the NSTAR ion propulsion system on the Deep Space One mission," *27th International Electric Propulsion Conference, IEPC-2001-075*, Electric Rocket Propulsion Society, Pasadena, CA, 2001.

[29] Williams, S. N. and Coverstone-Carroll, V., "Mars missions using solar electric propulsion," *Journal of Spacecraft and Rockets*, Vol. 37, No. 1, 2000, pp. 71–77.

[30] NASA, "Human Space Exploration Framework Summary," Tech. rep., National Aeronautics and Space Administration, Jan. 2010.

[31] Myers, R. and Carpenter, C., "High power solar electric propulsion for human space exploration architectures," *32nd International Electric Propulsion Conference, IEPC-2011-261*, Electric Rocket Propulsion Society, Wiesbaden, Germany, 2011.

[32] Seikel, G. and Reshotko, E., "Hall current ion accelerator," *Bull, Am. Phys. Soc*, 1962.

[33] Salz, F., Meyerand, R. G., and Lary, E. C., "Ion Acceleration in a Gyro-Dominated Neutral Plasma Experiment," *Bulletin of the American Physical Society, Sr. II*, Vol. 7, 1962, pp. 441.

[34] Lary, E. C., Meyerand, R. G., and Salz, F., "Ion Acceleration in a Gyro-Dominated Neutral PlasmaTheory," *Bulletin of the American Physical Society, Sr. II*, Vol. 7, 1962, pp. 441.

[35] Morozov, A. I., Kislov, A. Y., and Zubkov, I. P., "Strong-current Plasma Accelerator with Closed Electron Drift," *ZhETF Pisma Redaktsiiu*, Vol. 7, 1968, pp. 224.

[36] Choueiri, E. Y., "Fundamental difference between the two Hall thruster variants," *Physics of Plasmas (1994-present)*, Vol. 8, No. 11, 2001, pp. 5025–5033.

[37] Hofer, R. R. and Gallimore, A. D., "The role of magnetic field topography in improving the performance of high-voltage Hall thrusters," *38th Joint Propulsion Conference, AIAA-2002-4111*, American Institute of Aeronautics and Astronautics, Indianapolis, IN, 2002.

[38] Hofer, R. R. and Jankovsky, R. S., "The influence of current density and magnetic field topography in optimizing the performance, divergence, and plasma oscillations of high specific impulse Hall thrusters," Tech. Rep. NASA/TM–2003-212605, Cleveland, OH, 2003.

[39] Fisher, J., Wilson, A., King, D., Meyer, S., de Grys, K., and Werthman, L., "The development and qualification of a 4.5 kW Hall thruster propulsion system for GEO satellite applications," *27th International Electric Propulsion Conference, IEPC-2001-010*, Electric Rocket Propulsion Society, Pasadena, CA, Oct. 2001.

[40] Mikellides, I. G., Katz, I., and Hofer, R. R., "Design of a Laboratory Hall Thruster with Magnetically Shielded Channel Walls, Phase I: Numerical Simulations," *47th Joint Propulsion Conference, AIAA-2011-5809*, Vol. 2011-5809, American Institute of Aeronautics and Astronautics, San Diego, CA, 2011.

[41] Hofer, R., Goebel, D., Mikellides, I., and Katz, I., "Design of a Laboratory Hall Thruster with Magnetically Shielded Channel Walls, Phase II: Experiments," *48th Joint Propulsion Conference, AIAA-2012-3788*, American Institute of Aeronautics and Astronautics, Atlanta, GA, 2012.

[42] Goebel, D. M., Hofer, R. R., Mikellides, I. G., Katz, I., Polk, J. E., and Dotson, B. N., "Conducting Wall Hall Thrusters," *IEEE Transactions on Plasma Science*, Vol. 43, No. 1, 2013, pp. 118–126.

[43] Moraes, B. S., Ferreira, J. L., Ferreira, I. S., Winter, O. C., and Mouro, D. C., "A Permanent Magnet Hall Thruster for Pulsed Orbit Control of Lunar Polar Satellites," *Journal of Physics: Conference Series*, Vol. 511, IOP Publishing, 2014, p. 012074.

[44] Chen, F. F., *Introduction to Plasma Physics and Controlled Fusion*, Vol. 1, Plenum Press, New York, 2nd ed., 1984.

[45] Bellan, P. M., *Fundamentals of plasma physics*, Cambridge University Press, 2008.

[46] Lieberman, M. A. and Lichtenberg, A. J., *Principles of plasma discharges and materials processing*, John Wiley & Sons, 2005.

[47] Choueiri, E. Y., "Plasma oscillations in Hall thrusters," *Physics of Plasmas*, Vol. 8, No. 4, 2001, pp. 1411–1426.

[48] Sekerak, M. J., Gallimore, A. D., Brown, D. L., Hofer, R. R., and Polk, J. E., "Mode Transitions in Hall-Effect Thrusters Induced by Variable Magnetic Field Strength," *Journal of Propulsion and Power*, Vol. 32, No. 4, 2016, pp. 903–917.

[49] McDonald, M. S. and Gallimore, A. D., "Rotating spoke instabilities in Hall thrusters," *IEEE Transactions on Plasma Science*, Vol. 39, No. 11, 2011, pp. 2952–2953.

[50] Ellison, C. L., Raitses, Y., and Fisch, N. J., "Cross-field electron transport induced by a rotating spoke in a cylindrical Hall thruster," *Physics of Plasmas*, Vol. 19, No. 1, 2012, pp. 013503.

[51] Goebel, D. M., Jameson, K. K., Katz, I., and Mikellides, I. G., "Potential fluctuations and energetic ion production in hollow cathode discharges," *Physics of Plasmas*, Vol. 14, No. 10, 2007, pp. 103508.

[52] Fife, J. M., Martinez-Sanchez, M., and Szabo, J., "A numerical study of low-frequency discharge oscillations in Hall thrusters," *33rd Joint Propulsion Conference, AIAA-1997-3052*, American Institute of Aeronautics and Astronautics, Seattle, WA, 1997.

[53] Boeuf, J. P. and Garrigues, L., "Low frequency oscillations in a stationary plasma thruster," *Journal of Applied Physics*, Vol. 84, No. 7, 1998, pp. 3541–3554.

[54] Adam, J. C., Hron, A., and Laval, G., "Study of stationary plasma thrusters using two-dimensional fully kinetic simulations," *Physics of Plasmas*, Vol. 11, No. 1, 2004, pp. 295–305.

[55] Barral, S. and Ahedo, E., "Theoretical study of the breathing mode in Hall thrusters," *42nd AIAA Joint Propulsion Conference, AIAA-2006-5172*, 2006.

[56] Barral, S. and Ahedo, E., "Low-frequency model of breathing oscillations in Hall discharges," *Physical Review E*, Vol. 79, No. 4, 2009, pp. 046401.

[57] Hara, K., Boyd, I. D., and Kolobov, V. I., "One-dimensional hybrid-direct kinetic simulation of the discharge plasma in a Hall thruster," *Physics of Plasmas*, Vol. 19, No. 11, 2012, pp. 113508.

[58] Hara, K. and Boyd, I. D., "Low frequency oscillation analysis of a hall thruster using a one-dimensional hybrid-direct kinetic simulation," *33rd International Electric Propulsion Conference, IEPC-2013-266*, Electric Rocket Propulsion Society, Washington D.C., 2013.

[59] Szabo, J., Warner, N., Martinez-Sanchez, M., and Batishchev, O., "Full particle-in-cell simulation methodology for axisymmetric hall effect thrusters," *Journal of Propulsion and Power*, Vol. 30, No. 1, 2013, pp. 197–208.

[60] Sekerak, M. J., *Plasma oscillations and operational modes in Hall effect thrusters*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2014.

[61] McDonald, M. S., *Electron transport in Hall thrusters*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2012.

[62] Janes, G. S. and Lowder, R. S., "Anomalous Electron Diffusion and Ion Acceleration in a Low-Density Plasma," *The Physics of Fluids*, Vol. 9, No. 6, 1966, pp. 1115–1123.

[63] Huang, W., Kamhawi, H., and Haag, T., "Plasma Oscillation Characterization of NASAs HERMeS Hall Thruster via High Speed Imaging," *52nd Joint Propulsion Conference, AIAA-2016-4829*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, 2016.

[64] Sankovic, J. M., Hamley, J. A., and Haag, T. W., "Performance evaluation of the Russian SPT-100 thruster at NASA LeRC," *23rd International Electric Propulsion Conference, IEPC-1993-094*, Electric Rocket Propulsion Society, Seattle, WA, 1993.

[65] Garner, C. E., Polk, J. C., Goodfellow, K. M., Pless, L. C., and Brophy, J. R., "Performance evaluation and life testing of the SPT-100," *23rd International Electric Propulsion Conference, IEPC-1993-091*, Electric Rocket Propulsion Society, Seattle, WA, 1993.

[66] Herman, D. A., Santiago, W., Kamhawi, H., Polk, J. E., Snyder, J. S., Hofer, R. R., and Sekerak, M. J., "The Ion Propulsion System for the Asteroid Redirect Robotic Mission," *52nd Joint Propulsion Conference, AIAA-2016-4824*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, 2016.

[67] Kamhawi, H., Haag, T., Huang, W., Herman, D. A., Thomas, R., Shastry, R., Yim, J., Chang, L., Clayman, L., Verhey, T., and others, "Performance Characterization of the Solar Electric Propulsion Technology Demonstration Mission 12.5-kW Hall Thruster," *34th International Electric Propulsion Conference, IEPC-2015-007*, Electric Rocket Propulsion Society, Kobe-Hyogo, Japan, 2015.

[68] Kaufman, H. R., "Technology of closed-drift thrusters," *AIAA journal*, Vol. 23, No. 1, 1985, pp. 78–87.

[69] Brophy, J. R., "Stationary plasma thruster evaluation in Russia," Tech. Rep. NASA-CR-192823, Jet Propulsion Laboratory, March 1992.

[70] Haas, J. M., Gulczinski III, F. S., Gallimore, A. D., Spanjers, G. G., and Spores, R. A., "Performance characteristics of a 5 kW laboratory Hall thruster," *34th Joint Propulsion Conference, AIAA-98-3503*, American Institute of Aeronautics and Astronautics, Cleveland, OH, 1998.

[71] Hofer, R. R., *Development and Characterization of High-Efficiency, High-Specific Impulse Xenon Hall Thrusters*, Ph.D. Dissertation, University of Michigan, Ann Arbor, Michigan, 2004.

[72] Reid, B. M., *The Influence of Neutral Flow Rate in the Operation of Hall Thrusters*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2009.

[73] Jacobson, D. T. and Manzella, D. H., "50 kW Class Krypton Hall thruster Performance," *39th Joint Propulsion Conference, AIAA-2003-4550*, American Institute of Aeronautics and Astronautics, Huntsville, AL, 2003.

[74] Soulas, G., Haag, T., Herman, D., and Huang, W., "Performance Test Results of the NASA-457M v2 Hall Thruster," *48th Joint Propulsion Conference, AIAA-2012-3940*, American Institute of Aeronautics and Astronautics, Atlanta, GA, Aug. 2012.

[75] Kamhawi, H., Haag, T. W., Jacobson, D. T., and Manzella, D. H., "Performance Evaluation of the NASA-300M 20 kW Hall Effect Thruster," *47th Joint Propulsion Conference, AIAA-2011-5521*, American Institute of Aeronautics and Astronautics, San Diego, CA, Aug. 2011.

[76] Peterson, P. Y., Jacobson, D. T., Manzella, D. H., and John, J. W., "The performance and wear characterization of a high-power high-Isp NASA Hall thruster," *41st Joint Propulsion Conference, AIAA-2005-4243*, American Institute of Aeronautics and Astronautics, Tucson, AZ, 2005.

[77] Spores, R., Monheiser, J., Dempsey, B., Wade, D., Creel, K., Jacobson, D., and Drummond, G., "A Solar Electric Propulsion Cargo Vehicle to Support NASA Lunar Exploration Program," *29th International Electric Propulsion Conference, IEPC-2005-320*, Electric Rocket Propulsion Society, Princeton, NJ, 2005.

[78] Manzella, D., Jacobson, D., and Jankovsky, R., "High voltage SPT performance," *37th Joint Propulsion Conference, AIAA-2001-3774*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, 2001.

[79] Hofer, R. R. and Gallimore, A. D., "High-Specific Impulse Hall Thrusters, Part 2: Efficiency Analysis," *Journal of Propulsion and Power*, Vol. 22, No. 4, 2006, pp. 732–740.

[80] Haas, J. M., Hofer, R. R., Brown, D. L., Reid, B. M., and Gallimore, A. D., "Design of a 6-kW Hall Thruster for High Thrust/Power Investigation," *54th Joint Army Navy NASA Air Force Propulsion Meeting*, Denver, CO, May 2007.

[81] NASA, "Next Space Technologies for Exploration Partnerships Broad Agency Announcement," Oct. 2014.

[82] Florenz, R. E., Gallimore, A. D., and Peterson, P., "Development Status of a 100-kW Class Laboratory Nested Channel Hall Thruster," *32nd International Electric Propulsion Conference, IEPC-2011-246*, Electric Rocket Propulsion Society, Wiesbaden, Germany, Sept. 2011.

[83] Brown, D. L., Haas, J. M., Peterson, P. Y., and Kirtley, D. E., "Development of High-Power Electric Propulsion Technology for Near-Term and Mid-Term Space Power," *57th Joint Army Navy NASA Air Force Propulsion Meeting*, Colorado Springs, CO, May 2010.

[84] Hall, S. J., Jorns, B. A., Gallimore, A. D., and Hofer, R. R., "Expanded Thruster Mass Model Incorporating Nested Hall Thrusters," *53rd Joint Propulsion Conference, AIAA-2017-4729*, American Institute of Aeronautics and Astronautics, Atlanta, GA, 2017, p. 4729.

[85] Szabo, J., Pote, B., Hruby, V., Byrne, L., Tedrake, R., Kolencik, G., Kamhawi, H., and Haag, T. W., "A commercial one Newton Hall effect thruster for high power in-space missions," *47th Joint Propulsion Conference, AIAA-2011-6152*, American Institute of Aeronautics and Astronautics, San Diego, CA, 2011.

[86] McVey, J. B., Perrucci, A. S., and Britt, E. J., "Multichannel hall effect thruster," April 2006, US Patent 7,030,576.

[87] Grishin, S. D. and Leskov, L. V., *Electrical Rocket Engines of Space Vehicles*, Mashinostroenie, Moscow, 1989.

[88] Liang, R. and Gallimore, A. D., "Far-Field Plume Measurements of a Nested-Channel Hall-Effect Thruster," *49th AIAA Aerospace Sciences Meeting*, American Institute of Aeronautics and Astronautics, Orlando, FL, Jan. 2011.

[89] Liang, R. and Gallimore, A. D., "Constant-Power Performance and Plume Measurements of a Nested-Channel Hall-Effect Thruster," *32nd International Electric Propulsion Conference, IEPC-2011-049*, Electric Rocket Propulsion Society, Wiesbaden, Germany, Sept. 2011.

[90] Florenz, R., Liu, T., Gallimore, A., Kamhawi, H., Brown, D., Hofer, R., and Polk, J., "Electric Propulsion of a Different Class: The Challenges of Testing for MegaWatt Missions," *48th Joint Propulsion Conference, AIAA-2012-3942*, American Institute of Aeronautics and Astronautics, Atlanta, GA, July 2012.

[91] Hall, S. J., Florenz, R. E., Gallimore, A. D., Kamhawi, H., Peterson, P. Y., Brown, D. L., Hofer, R. R., and Polk, J. E., "Design Details of a 100-kW Class Nested-channel Hall Thruster," *62nd Joint Army Navy NASA Air Force Propulsion Meeting*, Nashville, TN, June 2015.

[92] Hall, S. J., Florenz, R. E., Gallimore, A. D., Kamhawi, H., Brown, D. L., Polk, J. E., Goebel, D., and Hofer, R. R., "Implementation and Initial Validation of a 100-kW Class Nested-channel Hall Thruster," *50th Joint Propulsion Conference, AIAA-2014-3815*, American Institute of Aeronautics and Astronautics, Cleveland, OH, July 2014.

[93] Reilly, S., Hofer, R., and Hall, S., "Thermal Analysis of the 100-kW class X3 Hall Thruster," 47th International Conference on Environmental Systems, 2017.

[94] Chang-Diaz, F. R., "The VASIMR rocket," *Scientific American*, Vol. 283, No. 5, 2000, pp. 90–97.

[95] Longmier, B. W., Cassady, L. D., Ballenger, M. G., Carter, M. D., Chang-Diaz, F. R., Glover, T. W., Ilin, A. V., McCaskill, G. E., Olsen, C. S., Squire, J. P., and others, "VX-200 magnetoplasma thruster performance results exceeding fifty-percent thruster efficiency," *Journal of Propulsion and Power*, Vol. 27, No. 4, 2011, pp. 915–920.

[96] Pancotti, A. P., Little, J. M., Neuhoff, J. S., Cornella, B. M., Kirtley, D. E., and Slough, J. T., "Electrodeless Lorentz Force (ELF) Thruster for ISRU and Sample Return Mission," *34th International Electric Propulsion Conference, IEPC-2015-67*, Electric Rocket Propulsion Society, Kobe-Hyogo, Japan, July 2015.

[97] Goebel, D. M. and Chu, E., "High-Current Lanthanum Hexaboride Hollow Cathode for High-Power Hall Thrusters," *Journal of Propulsion and Power*, Vol. 30, No. 1, 2014, pp. 35–40.

[98] Jackson, J., Allen, M., Myers, R., Hoskins, A., Soendker, E., Welander, B., Tolentino, A., Hablitzel, S., Hall, S. J., Jorns, B. A., Gallimore, A. D., Hofer, R. R., and Pencil, E., "100-kW Nested Hall Thruster System Development," *35th International Electric Propulsion Conference, IEPC-2017-219*, Electric Rocket Propulsion Society, Atlanta, GA, Oct. 2017.

[99] Polk, J. E., Pancotti, A., Haag, T., King, S., Walker, M., Blakely, J., and Ziemer, J., "Recommended Practice for Thrust Measurement in Electric Propulsion Testing," *Journal of Propulsion and Power*, Vol. 33, No. 3, 2017, pp. 539–555.

[100] Goebel, D., Jameson, K. K., and Hofer, R. R., "Hall thruster cathode flow impact on coupling voltage and cathode life," *Journal of Propulsion and Power*, Vol. 28, No. 2, 2012, pp. 355–363.

[101] Kim, V., "Main physical features and processes determining the performance of stationary plasma thrusters," *Journal of Propulsion and Power*, Vol. 14, No. 5, 1998, pp. 736–743.

[102] Hofer, R. and Gallimore, A., "Efficiency analysis of a high-specific impulse Hall thruster," *40th Joint Propulsion Conference, AIAA-2004-3602*, American Institute of Aeronautics and Astronautics, Ft. Lauderdale, FL, 2004.

[103] Brown, D. L., Larson, C. W., Beal, B. E., and Gallimore, A. D., "Methodology and historical perspective of a Hall thruster efficiency analysis," *Journal of Propulsion and Power*, Vol. 25, No. 6, 2009, pp. 1163–1177.

[104] Hofer, R., Katz, I., Goebel, D., Jameson, K., Sullivan, R., Johnson, L., and Mikellides, I., "Efficacy of electron mobility models in hybrid-PIC Hall thruster simulations," *44th Joint Propulsion Conference, AIAA-2008-4924*, American Institute of Aeronautics and Astronautics, Hartford, CT, 2008.

[105] Shastry, R., Hofer, R. R., Reid, B. M., and Gallimore, A. D., "Method for analyzing EB probe spectra from Hall thruster plumes," *Review of Scientific Instruments*, Vol. 80, No. 6, June 2009, pp. 063502.

[106] Brown, D. L., *Investigation of Low Discharge Voltage Hall Thruster Characteristics and Evaluation of Loss Mechanisms*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2009.

[107] Huang, W., Shastry, R., Herman, D. A., Soulas, G. C., and Kamhawi, H., "Ion Current Density Study of the NASA-300M and NASA-457Mv2 Hall Thrusters," *48th Joint Propulsion Conference, AIAA-2012-3870*, American Institute of Aeronautics and Astronautics, Atlanta, GA, 2012.

[108] Huang, W., Shastry, R., Soulas, G. C., and Kamhawi, H., "Farfield Plume Measurement and Analysis on the NASA-300M and NASA-300MS," *33rd International Electric Propulsion Conference, IEPC-2013-057*, Electric Rocket Propulsion Society, Washington D.C., 2013.

[109] Randolph, T., Kim, V., Kaufman, H., Kozubsky, K., Zhurin, V. V., and Day, M., "Facility Effects on Stationary Plasma Thruster Testing," *23rd International Electric Propulsion Conference, IEPC-1993-093*, Electric Rocket Propulsion Society, Seattle, WA, 1993, pp. 13–16.

[110] Kamhawi, H., Huang, W., Haag, T., and Spektor, R., "Investigation of the effects of facility background pressure on the performance and voltage-current characteristics of the high voltage hall accelerator," *33rd International Electric Propulsion Conference, IEPC-2013-446*, Electric Rocket Propulsion Society, Washington D.C., 2013.

[111] Hofer, R. R. and Anderson, J. R., "Finite Pressure Effects in Magnetically Shielded Hall thrusters," *50th Joint Propulsion Conference, AIAA-2014-3709*, American Institute of Aeronautics and Astronautics, Cleveland, OH, 2014.

[112] Frieman, J., *Characterization of Background Neutral Flows in Vacuum Test Facilities and Impacts on Hall Effect Thruster Operation*, Ph.D. Dissertation, Georgia Institute of Technology, Atlanta, GA, 2017.

[113] Frieman, J. D., Liu, T. M., and Walker, M. L., "Background Flow Model of Hall Thruster Neutral Ingestion," *Journal of Propulsion and Power*, Vol. 33, No. 5, 2017, pp. 1087–1101.

[114] Cai, C., *Theoretical and Numerical Studies of Plume Flows in Vacuum Chambers*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2005.

[115] Nakles, M. R. and Hargus Jr, W. A., "Background pressure effects on ion velocity distribution within a medium-power Hall thruster," *Journal of Propulsion and Power*, Vol. 27, No. 4, 2011, pp. 737.

[116] Kamhawi, H., Huang, W., Haag, T., Shastry, R., Soulas, G., Smith, T., Mikellides, I., and Hofer, R., "Performance and Thermal Characterization of the NASA-300MS 20 kW Hall Effect Thruster," *33rd International Electric Propulsion Conference, IEPC-2013-444*, Electric Rocket Propulsion Society, Washington D.C., Oct. 2013.

[117] Linnell, J. A. and Gallimore, A. D., "Krypton performance optimization in high-voltage Hall thrusters," *Journal of propulsion and power*, Vol. 22, No. 4, 2006, pp. 921–925.

[118] Linnell, J. A. and Gallimore, A. D., "Efficiency analysis of a hall thruster operating with krypton and xenon," *Journal of Propulsion and Power*, Vol. 22, No. 6, 2006, pp. 1402–1418.

[119] Linnel, J. A., *An Evaluation of Krypton Propellant in Hall Thrusters*, Ph.D. Dissertation, University of Michigan, Ann Arbor, Michigan, 2007.

[120] Florenz, R. E., Hall, S. J., Gallimore, A. D., Kamhawi, H., Griffiths, C. M., Brown, D. L., Hofer, R. R., and Polk, J. E., "First Firing of a 100-kW Nested-channel Hall Thruster," *33rd International Electric Propulsion Conference, IEPC-2013-394*, Electric Rocket Propulsion Society, Washington D.C., Oct. 2013.

[121] Walker, M. L. R., *Effects of facility backpressure on the performance and plume of a Hall thruster*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2005.

[122] Obama, B., "Remarks by the President on Space Exploration in the 21st Century," April 2010.

[123] Braun, R. D. and Blersch, D. J., "Propulsive options for a manned Mars transportation system," *Journal of Spacecraft and Rockets*, Vol. 28, No. 1, 1991, pp. 85–92.

[124] Chiravalle, V. P., "Nuclear electric ion propulsion for three deep space missions," *Acta Astronautica*, Vol. 62, No. 6, 2008, pp. 374–390.

[125] Johnson, R. A., Morgan, W. T., and Rocklin, S. R., "SNAP 10A, first reactor in space," Tech. rep., SAE Technical Paper, 1965.

[126] Oleson, S. and Katz, I., "Electric Propulsion for Project Prometheus," *39th Joint Propulsion Conference, AIAA-2003-5279*, American Institute of Aeronautics and Astronautics, Huntsville, AL, 2003.

[127] Palac, D., Gibson, M., Mason, L., Houts, M., McClure, P., and Robinson, R., "Nuclear Systems Kilopower Overview," NASA Technical Report 20160012354, 2016.

[128] Gibson, M. A., Oleson, S. R., Poston, D. I., and McClure, P., "NASA's Kilopower Reactor Development and the Path to Higher Power Missions," *IEEE Aerospace Conference*, IEEE, Big Sky, MT, 2017.

[129] Strange, N., Landau, D., Polk, J., Brophy, J., and Mueller, J., "Solar Electric Propulsion for a Flexible Path of Human Space Exploration," *61st International Astronautical Congress, IAC-10 A*, Vol. 5, 2010.

[130] Oh, D. Y., Goebel, D. M., Elkins-Tanton, L., Polanskey, C., Lord, P., Tilley, S., Snyder, J. S., Carr, G., Collins, S., Lantoine, G., and others, "Psyche: Journey to a Metal World," *52nd Joint Propulsion Conference, AIAA-2016-4541*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, 2016.

[131] Goebel, D. M., "Asteroid Retrieval and Mining with High Power Electric Propulsion," Feb. 2015.

[132] of US Human Spaceflight Plans Committee, R., *Seeking a human spaceflight program worthy of a great nation*, National Aeronautics and Space Administration, 2009.

[133] Strange, N., Merrill, R., Landau, D., Drake, B., Brophy, J., and Hofer, R., "Human missions to phobos and deimos using combined chemical and solar electric propulsion," *47th Joint Propulsion Conference, AIAA-2011-5663*, American Institute of Aeronautics and Astronautics, San Diego, CA, 2011.

[134] von Braun, W., *The Mars Project*, University of Illinois Press, Urbana, IL, 1953.

[135] Portree, D. S., *Humans to Mars: Fifty years of mission planning, 1950-2000*, National Aeronautics and Space Administration, 2001.

[136] Patterson, M. J. and Sovey, J. S., "History of electric propulsion at NASA Glenn Research Center: 1956 to present," *Journal of Aerospace Engineering*, Vol. 26, No. 2, 2012, pp. 300–316.

[137] Zubrin, R., *The Case for Mars*, The Free Press, New York, 1996.

[138] Dudzinski, L. A., Hack, K. J., Gefert, L. P., Kerslake, T. W., and Hewston, A. W., "Design of a solar electric propulsion transfer vehicle for a non-nuclear human mars exploration architecture," *26th International Electric Propulsion Conference, IEPC-1999-181*, Electric Rocket Propulsion Society, Kitakyushu, Japan, 1999.

[139] Price, H. W., Baker, J. D., Strange, N., and Woolley, R., "Human missions to Mars orbit, Phobos, and Mars surface using 100-kWe-class solar electric propulsion," *AIAA Space 2014 Conference*, American Institute of Aeronautics and Astronautics, San Diego, CA, 2014.

[140] Merrill, R. G., Strange, N., Qu, M., and Hatten, N., "Mars conjunction crewed missions with a reusable hybrid architecture," *2015 IEEE Aerospace Conference*, IEEE, Big Sky, MT, 2015, pp. 1–14.

[141] Sovey, J. S., Hamley, J. A., Haag, T. W., Patterson, M. J., Pencil, E. J., Peterson, T. T., Pinero, L. R., and others, "Development of an Ion Thruster and Power Processor for New Millenniums Deep Space 1 Mission," *33rd Joint Propulsion Conference, AIAA-97-2778*, American Institute of Aeronautics and Astronautics, Seattle, WA, 1997.

[142] Herman, D. A., Santiago, W., Kamhawi, H., Polk, J. E., Snyder, J. S., Hofer, R. R., and Parker, J. M., "The Ion Propulsion System for the Solar Electric Propulsion Technology Demonstration Mission," *34th International Electric Propulsion Conference, IEPC-2015-08*, Electric Rocket Propulsion Society, Kobe-Hyogo, Japan, 2015.

[143] Gefert, L. P., Hack, K. J., and Kerslake, T. W., "Options for the human exploration of mars using solar electric propulsion," *Space technology and applications international forum-1999*, Vol. 458, AIP Publishing, 1999, pp. 1275–1280.

[144] Hack, K. J., George, J. A., and Riehl, J. P., "Evolutionary Use of Nuclear Electric Propulsion," *AIAA Space Programs and Technologies Conference*, American Institute of Aeronautics and Astronautics, Huntsville, AL, Sept. 1990.

[145] Hack, K. J., George, J. A., and Dudzinski, L. A., "Nuclear Electric Propulsion Mission Performance for Fast Piloted Mars Missions," *AIAA/NASA/OAI Confernce on Advanced SEI Technologies*, American Institute of Aeronautics and Astronautics, Cleveland, OH, Sept. 1991.

[146] Clark, J. S., George, J. A., Gefert, L. P., Doherty, M. P., and Sefcik, R. J., "Nuclear electric propulsion: A better, safer, cheaper transportation system for human exploration of Mars," *11th Symposium on Space Nuclear Power and Propulsion Systems*, Jan. 1994.

[147] Chang-Diaz, F. R., Braden, E., Johnson, I., Hsu, M. M., and Yang, T. F., "Rapid Mars transits with exhaust-modulated plasma propulsion," Tech. Rep. NASA-TP-3539, National Aeronautics and Space Administration, March 1995.

[148] Keaton, P. W. and Tubb, D. J., "Nuclear electric propulsion," 1986.

[149] Rayman, M. D., Varghese, P., Lehman, D. H., and Livesay, L. L., "Results from the Deep Space 1 technology validation mission," *Acta Astronautica*, Vol. 47, No. 2, 2000, pp. 475–487.

[150] Brophy, J. R., Marcucci, M. G., Ganapathi, G. B., Garner, C. E., Henry, M. D., Nakazono, B., and Noon, D., "The ion propulsion system for Dawn," *39th Joint Propulsion Conference, AIAA-2003-4542*, American Institute of Aeronautics and Astronautics, Huntsville, AL, 2003.

[151] Gaidos, E. J. and Nimmo, F., "Planetary science: Tectonics and water on Europa," *Nature*, Vol. 405, No. 6787, 2000, pp. 637–637.

[152] Fortes, A. D., "Exobiological implications of a possible ammoniawater ocean inside Titan," *Icarus*, Vol. 146, No. 2, 2000, pp. 444–452.

[153] Waite Jr, J. H., Lewis, W. S., Magee, B. A., Lunine, J. I., McKinnon, W. B., Glein, C. R., Mousis, O., Young, D. T., Brockwell, T., Westlake, J., and others, "Liquid water on Enceladus from observations of ammonia and 40Ar in the plume," *Nature*, Vol. 460, No. 7254, 2009, pp. 487–490.

[154] Chang-Diaz, F. R., "An overview of the VASIMR engine: High power space propulsion with RF plasma generation and heating," *Radio Frequency Power in Plasmas: 14th Topical Conference*, Vol. 595, AIP Publishing, 2001, pp. 3–15.

[155] Kirtley, D., Brown, D. L., and Gallimore, A. D., "Details of an annular field reversed configuration plasma device for spacecraft propulsion," *29th International Electric Propulsion Conference, IEPC-2005-171*, Electric Rocket Propulsion Society, Princeton, NJ, 2005.

[156] Slough, J., Kirtley, D., and Weber, T., "Pulsed plasmoid propulsion: the ELF thruster," *31st International Electric Propulsion Conference, IEPC-2009-265*, Electric Rocket Propulsion Society, Ann Arbor, MI, 2009.

[157] Choueiri, E. Y. and Ziemer, J. K., "Quasi-steady magnetoplasmadynamic thruster performance database," *Journal of Propulsion and Power*, Vol. 17, No. 5, 2001, pp. 967–976.

[158] Burton, R. L., Clark, K. E., and Jahn, R. G., "Measured performance of a multimegawatt MPD thruster," *Journal of Spacecraft and Rockets*, Vol. 20, No. 3, 1983, pp. 299–304.

[159] Sovey, J. S. and Mantenieks, M. A., "Performance and Lifetime Assessment of MPD Arc Thruster Technology," *24th Joint Propulsion Conference, AIAA-88-3211*, American Institute of Aeronautics and Astronautics, Boston, MA, 1988.

[160] Sovey, J. S. and Mantenieks, M. A., "Performance and lifetime assessment of magnetoplasmadynamic arc thruster technology," *Journal of Propulsion and Power*, Vol. 7, No. 1, 1991, pp. 71–83.

[161] Lovberg, R. H. and Dailey, C. L., "Large inductive thruster performance measurement," *AIAA Journal*, Vol. 20, No. 7, 1982, pp. 971–977.

[162] Dailey, C. L., Lovberg, R. H., and Hieatt, J. L., "Multi-megawatt pulsed inductive thruster," May 1993, US Patent 5,207,760.

[163] Dailey, C. L. and Lovberg, R. H., "The PIT MkV pulsed inductive thruster," Tech. Rep. NASA-CR-191155, National Aeronautics and Space Administration, Cleveland, OH, July 1993.

[164] Polzin, K. A., "Comprehensive review of planar pulsed inductive plasma thruster research and technology," *Journal of Propulsion and Power*, Vol. 27, No. 3, 2011, pp. 513–531.

[165] Sengupta, A., Brophy, J. R., Anderson, J. R., Garner, C., de Groh, K., Karniotis, T., and Banks, B., "An Overview of the Results from the 30,000 Hr Life Test of Deep Space 1 Flight Spare Ion Engine," *40th Joint Propulsion Conference, AIAA-2004-3608*, American Institute of Aeronautics and Astronautics, Fort Lauderdale, FL, 2004.

[166] Nakanishi, S. and Pawlik, E. V., "Experimental investigation of a 1.5-m-diam Kaufman thruster." *Journal of Spacecraft and Rockets*, Vol. 5, No. 7, 1968, pp. 801–807.

[167] Patterson, M. J. and Benson, S. W., *NEXT ion propulsion system development status and performance*, National Aeronautics and Space Administration, Glenn Research Center, 2008.

[168] Patterson, M. J., Foster, J. E., Haag, T. W., Rawlin, V. K., Soulas, G. C., and Roman, R. F., "NEXT: NASAs Evolutionary Xenon Thruster," *38th Joint Propulsion Conference, AIAA-2002-3832*, American Institute of Aeronautics and Astronautics, Indianapolis, IN, July 2002.

[169] Goebel, D. M., Jorns, B. A., Hofer, R. R., Mikellides, I. G., and Katz, I., "Pole-piece interactions with the plasma in a magnetically shielded Hall thruster," *50th Joint Propulsion Conference, AIAA-2014-3899*, American Institute of Aeronautics and Astronautics, Cleveland, OH, 2014.

[170] Dannenmayer, K. and Mazouffre, S., "Elementary scaling relations for Hall effect thrusters," *Journal of Propulsion and Power*, Vol. 27, No. 1, 2011, pp. 236–245.

[171] Pynchon, T., *Gravity's Rainbow*, Penguin Books, New York, 2006.

[172] Kamhawi, H. and VanNoord, J., "Development and Testing of High Current Hollow Cathodes for High Power Hall Thrusters," *48th Joint Propulsion Conference, AIAA-2012-4080*, American Institute of Aeronautics and Astronautics, Atlanta, GA, 2012.

[173] Thomas, R. E., Kamhawi, H., and Williams, Jr., G. J., "High Current Hollow Cathode Plasma Plume Measurements," *33rd International Electric Propulsion Conference, IEPC-2013-76*, Electric Rocket Propulsion Society, Washington D.C., Oct. 2013.

[174] Chu, E., Goebel, D. M., and Wirz, R. E., "Reduction of Energetic Ion Production in Hollow Cathodes by External Gas Injection," *Journal of Propulsion and Power*, Vol. 29, No. 5, 2013, pp. 1155–1163.

[175] Dushman, S. and Lafferty, J. M., *Scientific foundations of vacuum technique*, Wiley, New York, 1962.

[176] Dankanich, J. W., Walker, M., Swiatek, M. W., and Yim, J. T., "Recommended Practice for Pressure Measurements and Calculation of Effective Pumping Speeds during Electric Propulsion Testing," *33rd International Electric Propulsion Conference, IEPC-2013-358*, Electric Rocket Propulsion Society, Washington D.C., 2013.

[177] Yim, J. and Burt, J. M., "Characterization of Vacuum Facility Background Gas Through Simulation and Considerations for Electric Propulsion Ground Testing," *51st Joint Propulsion Conference, AIAA-2015-3825*, American Institute of Aeronautics and Astronautics, Orlando, FL, July 2015.

[178] Kamhawi, H., Haag, T., Huang, W., Yim, J., Herman, D. A., Peterson, P. Y., Williams, G., Gilland, J., Hofer, R. R., and Mikellides, I. G., "Performance, Facility Pressure Effects, and Stability Characterization Tests of NASAs 12.5-kW Hall Effect Rocket with Magnetic Shielding Thruster," *52nd Joint Propulsion Conference, AIAA-2016-4826*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, 2016, DOI: 10.2514/6.2016-4826.

[179] Peterson, P. Y., Kamhawi, H., Huang, W., Williams, G., Gilland, J. H., Yim, J., Hofer, R. R., and Herman, D. A., "NASA HERMeS Hall Thruster Electrical Configuration Characterization," *52nd Joint Propulsion Conference, AIAA-2016-5027*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, 2016.

[180] Huang, W., Kamhawi, H., and Haag, T., "Facility Effect Characterization Test of NASAs HERMeS Hall Thruster," *52nd Joint Propulsion Conference, AIAA-2016-4828*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, 2016.

[181] Huang, W., Shastry, R., Soulas, G. C., and Kamhawi, H., "Farfield Plume Measurement and Analysis on the NASA-300M," NASA Technical Report 20130014362, NASA Glenn Research Center, 2013.

[182] Hutchinson, I. H., *Principles of plasma diagnostics*, Cambridge University Press, 2nd ed., 2002.

[183] Savitzky, A. and Golay, M. J., "Smoothing and differentiation of data by simplified least squares procedures." *Analytical chemistry*, Vol. 36, No. 8, 1964, pp. 1627–1639.

[184] King, L. B., *Transport-Property and Mass Spectral Measurements in the Plasma Exhaust Plume of a Hall-Effect Space Propulsion System*, Ph.D. Dissertation, University of Michigan, Ann Arbor, Michigan, 1998.

[185] Huang, W., Drenkow, B., and Gallimore, A. D., "Laser-Induced Fluorescence of Singly-Charged Xenon inside a 6-kW Hall Thruster," *45th Joint Propulsion Conference, AIAA-2009-5355*, American Institute of Aeronautics and Astronautics, Denver, CO, 2009.

[186] Durot, C. J., Georgin, M. P., and Gallimore, A. D., "Time-Resolved Laser-Induced Fluorescence Measurements in the Plume of a 6-kW Hall Thruster," *34th International Electric Propulsion Conference, IEPC-2015-399*, Electric Rocket Propulsion Society, Kobe, Japan, July 2015.

[187] Georgin, M. P., Dhaliwal, V., and Gallimore, A., "Investigation of Channel Interactions in a Nested Hall Thruster Part I: Acceleration Region Velocimetry," *52nd Joint Propulsion Conference, AIAA-2016-5030*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, July 2016.

[188] Durot, C. J., Gallimore, A. D., and Smith, T. B., "Validation and evaluation of a novel time-resolved laser-induced fluorescence technique," *Review of Scientific Instruments*, Vol. 85, No. 1, 2014, pp. 013508.

[189] Lobbia, R. B. and Beal, B. E., "Recommended Practice for Use of Langmuir Probes in Electric Propulsion Testing," *Journal of Propulsion and Power*, Vol. 33, No. 3, 2017, pp. 566–581.

[190] Segall, S. B. and Koopman, D. W., "Application of cylindrical Langmuir probes to streaming plasma diagnostics," *The Physics of Fluids*, Vol. 16, No. 7, 1973, pp. 1149–1156.

[191] Goel, S. K., Gupta, P. D., and Bhawalkar, D. D., "V-I characteristics of a plane disc probe in collisionless streaming plasmas," *Journal of Applied Physics*, Vol. 53, No. 4, 1982, pp. 2971–2974.

[192] Bransden, B. H. and McDowell, M. R. C., *Charge exchange and the theory of ion-atom collisions*, No. 82 in International series of monographs on physics, Clarendon Press, New York, 1992.

[193] Jorns, B. A. and Hofer, R. R., "Plasma oscillations in a 6-kW magnetically shielded Hall thruster," *Physics of Plasmas*, Vol. 21, No. 5, 2014, pp. 053512.

[194] Dale, E. T. and Gallimore, A. D., "High-Speed Image Analysis and Filtered Imaging of Nested Hall Thruster Oscillations," *34th International Electric Propulsion Conference, IEPC-2015-285*, Electric Rocket Propulsion Society, Hyogo-Kobe, Japan, July 2015.

[195] McDonald, M. S., Sekerak, M. J., Gallimore, A. D., and Hofer, R. R., "Plasma Oscillation Effects on Nested Hall Thruster Operation and Stability," *2013 IEEE Aerospace Conference*, IEEE, Big Sky, MT, 2013.

[196] Besnard, E., Chen, H. H., Mueller, T., and Garvey, J., "Design, manufacturing and test of a plug nozzle rocket engine," *38th Joint Propulsion Conference, AIAA-02-4038*, American Institute of Aeronautics and Astronautics, Indianapolis, IN, 2002.

[197] Haag, T., "Thrust Stand for High-Powered Electric Propulsion Devices," *Review of Scientific Instruments*, Vol. 62, No. 1186, 1991.

[198] Xu, K. G. and Walker, M. L., "High-power, null-type, inverted pendulum thrust stand," *Review of Scientific Instruments*, Vol. 80, No. 5, 2009, pp. 055103.

[199] Szabo, J., Pote, B., Paintal, S., Robin, M., Hillier, A., Branam, R. D., and Huffmann, R. E., "Performance evaluation of an iodine-vapor Hall thruster," *Journal of Propulsion and Power*, Vol. 28, No. 4, 2012, pp. 848–857.

[200] Hall, S. J., Gallimore, A. D., and Viges, E., "Thrust Stand for Very-High-Power Hall Thrusters," *63rd Joint Army Navy NASA Air Force Propulsion Meeting*, Phoenix, AZ, Dec. 2016.

[201] Reid, B. M., Gallimore, A. D., Hofer, R. R., Li, Y., and Haas, J. M., "Anode Design and Verification for the H6 Hall Thruster," *54th Joint Army Navy NASA Air Force Propulsion Meeting*, Denver, CO, May 2007.

[202] Huang, W., Yim, J. T., and Kamhawi, H., "Design and Emperical Assessment of the HERMeS Hall Thruster Propellant Manifold," Nashville, TN, June 2015.

[203] GeekWire, "Jeff Bezos discusses space flight and his vision for Blue Origin," Oct. 2016.

[204] Cusson, S. E., Dale, E. T., and Gallimore, A., "Investigation of Channel Interactions in a Nested Hall Thruster Part II: Probes and Performance," *52nd Joint Propulsion Conference, AIAA-2016-5029*, American Institute of Aeronautics and Astronautics, Salt Lake City, UT, July 2016.

[205] Cusson, S. E., Hall, S. J., Hofer, R. R., Jorns, B. A., and Gallimore, A. D., "The Impact of Magnetic Field Coupling Between Channels in a Nested Hall Thruster," *35th International Electric Propulsion Conference, IEPC-2017-507*, Electric Rocket Propulsion Society, Atlanta, GA, 2017.

[206] Sommerville, J. D. and King, L. B., "Hall-Effect ThrusterCathode Coupling, Part I: Efficiency Improvements from an Extended Outer Pole," *Journal of Propulsion and Power*, Vol. 27, No. 4, 2011, pp. 744–753.

[207] Jankovsky, R., Tverdokhlebov, S., and Manzella, D., "High Power Hall Thrusters," *35th Joint Propulsion Conference*, American Institute of Aeronautics and Astronautics, Los Angeles, CA, 1999.

[208] Hofer, R. R., Cusson, S. E., Lobbia, R. B., and Gallimore, A. D., "The H9 Magnetically Shielded Hall Thruster," *35th International Electric Propulsion Conference, IEPC-2017-232*, Atlanta, GA, 2017.

[209] Cusson, S. E., Hofer, R. R., Lobbia, R. B., Jorns, B. A., and Gallimore, A. D., "Performance of the H9 Magnetically Shielded Hall Thrusters," *35th International Electric Propulsion Conference, IEPC-2017-239*, Electric Rocket Propulsion Society, Atlanta, GA, 2017.

[210] Peterson, P. Y., *The Development and Characterization of a Two-Stage Hybrid Hall/Ion Thruster*, Ph.D. Dissertation, University of Michigan, Ann Arbor, MI, 2004.

[211] von Karman, T. and Edson, L., *The Wind and Beyond*, Little, Brown & Company, 1967.

[212] Hall, S. J., Cusson, S. E., and Gallimore, A. D., "30-kW Performance of a 100-kW Class Nested-channel Hall Thruster," *34th International Electric Propulsion Conference, IEPC-2015-125*, Electric Rocket Propulsion Society, Kobe-Hyogo, Japan, July 2015.

[213] Beal, B. E., Gallimore, A. D., and Hargus, W. A., "Effects of Cathode Configuration on Hall Thruster Cluster Plume Properties," *Journal of Propulsion and Power*, Vol. 23, No. 4, 2007, pp. 836–844.

[214] Sommerville, J. D. and King, L. B., "Effect of cathode position on Hall-effect thruster performance and cathode coupling voltage," *43rd Joint Propulsion Conference, AIAA-2007-5174*, American Institute of Aeronautics and Astronautics, Cincinnati, OH, 2007.

[215] Jameson, K. K., Goebel, D. M., Hofer, R. R., and Watkins, R. M., "Cathode coupling in Hall thrusters," *30th International Electric Propulsion Conference, IEPC-2007-278*, Electric Rocket Propulsion Society, Florence, Italy, 2007, pp. 2007–278.

[216] McDonald, M. S. and Gallimore, A. D., "Cathode position and orientation effects on cathode coupling in a 6-kW Hall thruster," *31st International Electric Propulsion Conference, IEPC-2009-113*, Electric Rocket Propulsion Society, Ann Arbor, MI, 2009.

[217] Dean, M. L., *Leaving Orbit: Notes from the Last Days of American Spaceflight*, Graywolf Press, Minneapolis, MN, 2015.

[218] Mikellides, I. G. and Ortega, A. L., "Numerical Simulations of a 100-kW Class Nested Hall Thruster with the 2-D Axisymmetric Code Hall2De," *35th International Electric Propulsion Conference, IEPC-2017-220*, Electric Rocket Propulsion Society, Atlanta, GA, 2017.

[219] Huang, W., Williams, G. J., Peterson, P. Y., Kamhawi, H., Gilland, J. H., and Herman, D. A., "Plasma Plume Characterization of the HERMeS during a 1722-hr Wear Test Campaign," *35th International Electric Propulsion Conference, IEPC-2017-307*, Electric Rocket Propulsion Society, Atlanta, GA, 2017.

[220] Dankanich, J. W., Walker, M., Swiatek, M. W., and Yim, J. T., "Recommended Practice for Pressure Measurement and Calculation of Effective Pumping Speed in Electric Propulsion Testing," *Journal of Propulsion and Power*, Vol. 33, No. 3, 2017, pp. 668–680.

[221] Burrows, W. E., *This New Ocean: The Story of the First Space Age*, Random House, New York, 1998.

[222] Peterson, P. Y., Kamhawi, H., Huang, W., Yim, J. T., Haag, T. W., Mackey, J. A., McVetta, M., Sorrelle, L., Tomsik, T., Gillagan, R., and Herman, D. A., "Reconfiguration of NASA GRC's Vacuum Facility 6 for Testing of Advanced Electric Propulsion System (AEPS) Hardware," *35th International Electric Propulsion Conference, IEPC-2017-028*, Electric Rocket Propulsion Society, Atlanta, GA, 2017.

[223] Rodriguez, W. A. and Dankanich, J., "National Space Simulation Facility Concept Design," *35th International Electric Propulsion Conference, IEPC-2017-16*, Electric Rocket Propulsion Society, Atlanta, GA, 2017.

[224] Chaplin, V. H., Conversano, R. W., Lobbia, R. B., Ortega, A. L., Mikellides, I. G., Hofer, R. R., and Jorns, B. A., "Laser-Induced Fluorescence Measurements of the Acceleration Zone in the 12.5 kW HERMeS Hall Thruster," *35th International Electric Propulsion Conference, IEPC-2017-229*, Electric Rocket Propulsion Society, Atlanta, GA, Oct. 2017.

[225] Shastry, R., Gallimore, A. D., and Hofer, R. R., "Near-wall plasma properties and EEDF measurements of a 6-kW Hall thruster," *45th Joint Propulsion Conference, AIAA-2009-5356*, American Institute of Aeronautics and Astronautics, Denver, CO, 2009.

[226] Shastry, R., Huang, W., Haag, T. W., and Kamhawi, H., "Langmuir Probe Measurements Within the Discharge Channel of the 20-kW NASA-300M and NASA-300MS Hall Thrusters," *33rd International Electric Propulsion Conference, IEPC-2013-122*, Electric Rocket Propulsion Society, Washington D.C., 2013.

[227] Shastry, R., Huang, W., and Kamhawi, H., "Near-Surface Plasma Characterization of the 12.5-kW NASA TDU1 Hall Thruster," *51st Joint Propulsion Conference, AIAA-2015-3919*, American Institute of Aeronautics and Astronautics, Orlando, FL, 2015.

[228] Frieman, J. D., Walker, J. A., Walker, M. L., Khayms, V., and King, D. Q., "Electrical Facility Effects on Hall Thruster Cathode Coupling: Performance and Plume Properties," *Journal of Propulsion and Power*, Vol. 32, No. 1, 2015, pp. 251–264.