Study of Hall Thruster Discharge Channel Wall Erosion via Optical Diagnostics

by

Wensheng Huang

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Aerospace Engineering)
in The University of Michigan
2011

Doctoral Committee:

Professor Alec D. Gallimore, Chair
Professor Iain D. Boyd
Associate Professor John E. Foster
Lecturer Timothy B. Smith
Hani Kamhawi, NASA Glenn Research Center

© Wensheng Huang
2011

Acknowledgements

My biggest thanks goes out to my advisor Alec Gallimore. Thank you for seeing my potential and giving me a chance to prove myself to the scientific community. Your unflagging support is what allowed me to finish this dissertation document. I do not think I could ever find a better advisor than you for me. You gave me just the right amount of push and leeway. I have always dreamt of working in NASA on advanced propulsion and now I will live out that dream thanks in large part to you.

I would like to thank my other committee members. Each of you is special in your own way to the development of my research. To Tim Smith, thank you for teaching me all you know about LIF and spectroscopy. I stand on the shoulders of giants and you have a giant shoulder. Thank you for all the advice you have given me over the years, I have greatly enjoyed all of our discussions together. To Hani Kamhawi, thank you for giving me my first opportunity to work in NASA. I know my GSRP experience will be invaluable to me in the future. I greatly appreciate that you were able to find time for me despite your busy schedule. To Ian Boyd, thank you for sparing the time and resources to aid me in my research. I learned a great many things from you and your students. I greatly enjoyed your unique sense of humor and the ease with which you can put smiles on people's faces. To John Foster, thank you for all of the advice, help, and knowledge you have given me these past years. You have a great knack for noticing things that I might otherwise miss. It was a lot of fun hanging out with you at GRC this past summer and I hope to enjoy your company again in the future.

Special thanks also go out to members of the PEPL gang. The biggest thanks goes to Bryan Reid, who taught me many of the basics of how to run a Hall thruster. Don't worry, the secrets of your embarrassing moments are safe with me. To the members of PEPL who preceded me, thanks goes to Robert Lobbia for your infinite patience and bottomless knowledge on electronics and fast diagnostics. And no, I do not plan on doing high-speed CRDS, yet. To Thomas Liu, thank you for aiding me with your unique

expertise on numerous occasions. I sometime wonder if there is anything you do not know or cannot do. By the way, what happened to “Jorge”? To Kristina Lemmer, thank you for looking out for me during my early years as a graduate student. Your RF shielding techniques have ever been a great inspiration. To Sonca Nguyen, thank you for helping me with my custom feedthrus. And no, I still have not switched over to LaTeX. To Bailo Ngom, I greatly enjoyed our partnership in developing an in-situ magnetic field sensor for Hall thruster. To Daniel Brown, thank you for giving me advice on various Hall-thruster-related subjects, your help with the cathode was greatly appreciated. To Ricky Tang, I will always remember our interesting Hartford experience together. Special thanks go to Peter Peterson and Richard Hofer. For although the two of you are many years my senior as PEPL alums you have each found time to aid my research, sometimes without you realizing.

To members of PEPL who are my year and younger, I like to start by thanking Rohit Shastry. Technically you are not my year but the number of experiments and amount of experience we have shared makes me feel otherwise. Thank you for always being there for me, I would not have finished all those experiments without your help. Let us show GRC what we can accomplish together. To Michael McDonald, it has been an interesting and fun experience sharing the lab with you. Thank you for sharing information with me for my dissertation. And hurry up and graduate. To Adam Shabshelowitz, thank you for all the help you have provided me during my dissertation work. I have greatly enjoyed your company, more than I usually let on. By the way, is Shab-E done yet? To Raymond Liang, thank you for all the help you have given me over the years. I have found your super-organized style of be a true inspiration, though I do not think I will try to mimic it. To Laura Spencer, thank you for keeping the lab an interesting and lively place for me to go to. And stop giving me that look. To Roland Florence, thank you for being one of my biggest supporters at the lab. Your antics always crack me up, even if I am not showing it on the outside. And last but not least, to Christopher Durot, thank you for your help during the early phase of CRDS. It has been fun working with you. Try not to get into too much trouble after I am gone.

Special thanks also go out to Azer Yalin and Lei Tao from the Colorado State University. Your contributions to our partnership on CRDS have been invaluable to me

and the University of Michigan team. On our behalf, I would like to thank you both for your excellent advice, speedy response, and generous aid. I hope to work with you again in future endeavors.

Next, I would like to thank the staff of the Aerospace Engineering department at the University of Michigan. Special thanks goes to Terry Larrow who did a very excellent job machining and welding the main rectangular frame for CRDS version two. Great thanks also go to Thomas Griffin, Christ Chartier, David McLean, and Eric Kirk for giving me great advice and performing various maintenance works at PEPL. Special thanks to Denise Phelps for handling all the official grad-student-related paperwork and helping me on numerous occasions. Many thanks to Suzanne Smith and Cynthia Enoch for helping me with various equipment purchases and paperwork. And last but not least, special thanks to Bonnie Bryant for looking out for me during IEPC events.

I would also like to thank personnel at NASA for rewarding me the GSRP fellowship and giving me interesting internship experiences. I greatly appreciated the help of Luis Pinero, Eric Pencil, Michael Meyer, Mark Kankam, Bernice Beznoska, and many others in making my NASA experience an enjoyable one.

Lastly, I would like to thank my family, particularly my father and mother. Though I never say it, I do love you both in my heart. I feel safe writing this here knowing that the two of you will probably never read this dissertation. God bless you both with long life and happiness.

Preface

This document describes how four years of research culminated in the placement of an uppercase “2” in one equation. The total equipment, consumables, and facility cost is ~$361,000. The total area taken up by the said “2” is ~2.3 mm^2. The cost density of the said “2” in physical area is 155 billion dollars/m^2. Assuming an ink surface coverage of ~500 m^2/g (common carbon black ink), the cost density of the said “2” in mass is 77 quadrillion dollars/kg. Thus, it is one the most expensive “2” ever produced. Luckily, the cost to create one copy of this “2” is on the order of much less than 1 cent. Therefore, by mass producing this “2”, its cost density can be greatly reduced. That or you waste a lot of paper.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

AEDF: Angular Energy Distribution Function

AOI: Angle-of-Incidence

AOM: Acousto-Optical Modulator

CFR: Carbon Fiber Rod

CRDS: Cavity Ring-Down Spectroscopy

CSU: Colorado State University

DAQ: Data Acquisition

FWHM: Full-Width-at-Half-Maximum

NIST: National Institute of Standards and Technology

HiVHAc: High Voltage Hall Accelerator

IR: Infrared

IW: Inner channel Wall

LIF: Laser-Induced Fluorescence

ND: Non-Dimensionalized

OES: Optical Emission Spectroscopy

OW: Outer channel Wall

PID: Proportional-Integral-Differential

PMT: Photomultiplier

PPT: Pulsed Plasma Thruster

QCM: Quartz Crystal Microbalance

SNR: Signal-to-Noise Ratio

SPT: Stationary Plasma Thruster

UM: University of Michigan

UV: Ultraviolet

VDF: Velocity Distribution Function

Chapter 1
Introduction

1.1 Problem Statement

The central question that this dissertation is trying to answer is "what is the mechanism by which Hall thruster channel walls are eroded?"

The Hall-effect thruster is a form of electric propulsion where an electric field is used as the primary acceleration mechanism to launch propellant out the back of a spacecraft. The performance of these devices has been well studied [1, 2], but the physical mechanism that limits their lifetime is still not well understood. In particular, the channel walls of Hall thrusters are vulnerable to erosion caused by the propellant particles accelerated by the thruster. Once the channel walls have been eroded away, the key components they protect will become exposed and the thruster's performance will degrade. This dissertation attempts to use novel optical diagnostics to study the physics of Hall thruster channel wall erosion.

Note that this dissertation focuses on magnetic-layer type Hall thrusters. Figure 1-1 shows one such magnetic-layer type Hall thruster in operation.

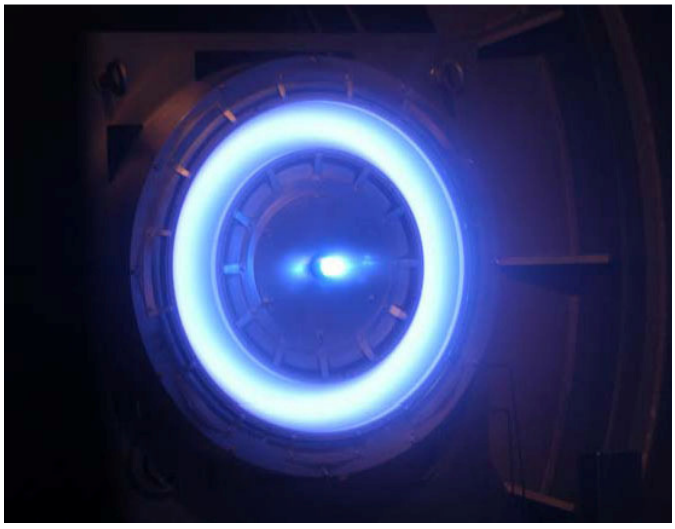

Figure 1-1. The NASA 457M in operation.
Alternate title: Motivation: it looks even prettier in person.

1.2 Motivation

There are all kinds of reasons why humanity must eventually head out to space, but the reason we want to explore space is because we are naturally curious. Unfortunately, our sense of curiosity has far outpaced our mobility. The root motivation for this dissertation is to aid in the physical accumulation of knowledge and the technological advancement that will enhance our mobility and grant us access to the farthest reaches of our Solar System.

The Hall-effect thruster is a unique space propulsion device that can vastly enhance our current reach in the Solar System. This device combines good efficiency, specific impulse, and a wide throttle range. The main drawback is that the thrust level is low compared to chemical propulsion system. These devices must operate for thousands to tens of thousands of hours in space in order to build up speed for the spacecraft [3, 4].

Unfortunately, even a low-temperature plasma like that found in Hall thrusters can greatly erode nearby components over a long period of time. In particular, due to the principles on which Hall thrusters operate, the bulk of the discharge plasma is located near vital magnetic components that sputter quickly when exposed to high-energy ions. The erosion of the discharge channel walls that protect these magnetic components is a key limiting factor on the life of the Hall thruster.

Great advances have been made in the study of Hall thruster channel wall erosion in both the experimental [5-7] and the theoretical [8-12] domains. However, the experimental studies have been limited to qualification tests. In these tests, a Hall thruster is operated at a set of throttle points for thousands of hours, and the channel wall profile is studied at various intervals. Each of these tests typically costs millions of dollars and takes years to complete, while yielding data for only a handful of thruster settings.

Theoretical development has also been greatly limited by the lack of parametric experimental data. If a thruster is ever qualified for a set of throttle points (after a great expenditure of resource and time) and it turns out a mission requires a different set of throttle points, it will generally be necessary to re-qualify the thruster from the beginning. Furthermore, limitations in our understanding of the wall erosion problem prevents us from accurately predicting the life of the thruster during the design phase. The Hall

thruster design process typically involves several test models and multiple wear tests before a design with a long life time is found. Once such a design is found, subsequent models can be created based on the prototype; however, it would be far better if we actually understood why the prototype design works and how to scale from the prototype while preserving a long operation life time.

With all that said, the main reason we currently lack physical understanding of the channel wall erosion problem is not due to a lack of effort but due to a lack of diagnostics. As mentioned before, our current best approach to studying the channel wall erosion problem is to simply measure the channel wall profile while running the Hall thruster for thousands of hours.

This dissertation is motivated by the possibility of using two novel optical diagnostics, laser-induced fluorescence (LIF) velocimetry and cavity ring-down spectroscopy (CRDS), to obtain both unique particle velocity data and erosion rate data in near-real-time. LIF velocimetry is used to study the propellant particles that bombard the channel walls, and CRDS is used to measure the channel wall erosion rates in near-real-time by studying the sputtered products.

1.3 Hypothesis

At the start of this dissertation, some four years ago, the hypothesis could be summed up in equation (1-1). For a given magnetic field topology, the rate at which channel wall material is sputtered off is described by,

$$\text{Sputter Rate} \sim \text{Discharge Voltage} * \text{Discharge Current} \tag{1-1}$$

A summary of the ideas that led to this scaling law can be found in Section 7.5 of Goebel and Katz's text [13]. Three important assumptions about the mechanism responsible for channel wall erosion are being made. One, the sputter yield of channel wall material is proportional to the energy of the bombarding particle. Two, the ion flux to the channel wall is proportional to the discharge current. Three, the acceleration zone does not change in structure and does not move. Equation (1-1) is also the starting point for this dissertation. The validity of the assumptions that make up this equation will be examined over the course of the dissertation. Note that for Hall thruster channel wall erosion, the

sputter rate refers to the same quantity as the erosion rate and the two are used interchangeably in this dissertation.

1.4 Detailed Description and Dissertation Goals

The original problem statement is very general and could take far more than five years of research by a graduate student to thoroughly answer. Instead, this dissertation is focused on just three aspects of the question: one, develop the tools needed to answer the question; two, attempt to narrow down the key parameters that drive channel wall erosion by using the said tools to perform parametric studies; three, try to quantify how the driving parameters relate to the erosion rate.

This dissertation contains two major experimental campaigns. In the first campaign, the main focus is on the development of an LIF-based diagnostic to study the particles that bombard the channel walls. The advantage of using LIF over other diagnostics is that LIF is non-intrusive, species-specific, and spatially resolved. LIF also provides excellent accuracy and is capable of mapping the velocity distribution functions (VDFs) of the particles in question.

There are several challenges to applying this technique in a Hall thruster environment. First, studying the bombarding particles inside of the discharge channel of the Hall thruster requires at least two laser axes. Mapping the near-wall VDFs of both singly-charged and neutral xenon (the bombarding particles) has never been tried before and has a number of associated geometric challenges. Second, VDFs are key to understanding the erosion problem because the sputtering of material by particle bombardment is strongly energy- and angle-dependent. Since we need to measure the VDFs, a number of spectral lineshape issues must be resolved in order to allow accurate measurements. The specifics of these issues are explained in Chapter 3 and Section 5.3.3. Third, the LIF diagnostics in question must be applied to a thruster that operates at up to 10 kW. There are a number of structural and thermal challenges that had to be worked out. The final solution was to use optical fiber beam delivery with protective shields that are transparent to the laser beam. Singly-charged and neutral LIF data were obtained for a 6-kW Hall thruster across seven operating conditions. These operating conditions

spanned discharge voltages from 150 V to 600 V, and anode mass flow rates from 10 mg/s to 30 mg/s.

In the second campaign, the main focus is on the development of a CRDS sensor to measure the density of ground state boron. This density is then correlated to the erosion rate of the thruster channel walls. To the best of the author's knowledge, this dissertation describes the first successful application of CRDS to erosion rate measurement on a magnetic-layer type Hall thruster. Since the CRDS diagnostic is very new to the field of Hall thruster research, there were a number of challenges that had to be overcome. These challenges included thermal and structural design issues, vibration isolation, mirror degradation, raw data analysis, and correlation of CRDS data to the erosion rate.

The specific developmental work and experimentation performed for this dissertation is a part of a much larger collaborative project between Prof. Alec Gallimore of the University of Michigan (UM) and Prof. Azer Yalin of the Colorado State University (CSU). CSU was responsible for the development of CRDS technique at the ultraviolet wavelength needed to study boron, and study many of the finer details of how to correlate boron density to erosion rate. UM was responsible for adapting the technique to study the aforementioned 6-kW Hall thruster. At the conclusion of the test, boron density data were obtained near the exit plane of the 6-kW Hall thruster at the same seven operating conditions used during the LIF tests.

In the final part of the thesis, the data obtained from the two major phases are combined to paint a picture of how the erosion of the 6-kW Hall thruster channel wall changed as a function of operating conditions. Detailed information on the acceleration zone length, the energy and angle of the singly-charged ion, and the erosion rate calculated from the boron density are studied in order to determine the main driving parameters in channel wall erosion physics.

In summary, the primary goals of this dissertation are,

1. Map the velocity distribution functions of the particles that are bombarding the channel wall,
2. Measure the erosion rate by studying the erosion product,
3. Determine the key parameters that drive the physics of channel wall erosion, and

4. Create an engineering model to relate the erosion rate from one operating condition to another operating condition for a given Hall thruster.

1.5 Prior and Current Research

In terms of magnetic-layer type Hall thrusters (also called stationary plasma thrusters, or SPTs), the oldest flight-like model on which erosion studies have been carried out is the SPT-100 designed in the former Soviet Union. In addition to the erosion profile measured by Absalamov, et al., [5] there is a vast amount of various types of data available on this and related designs, some of which are referenced in Goebel and Katz's text [13]. Some other SPTs on which extended wear tests and life tests have been performed include the T-100 [14], T-220 [15], PPS1350 [16], NASA 103M.XL [6], BHT-200 [12], BHT-600 [12], and BPT-4000 [7]. Unfortunately, most of these thrusters have only been tested for up to a few thousand hours, which is typically not enough to reach the end-of-life condition.

The oldest methods for studying Hall thruster channel wall erosion are weight loss and profilometry. As the name implies, the weight loss method involves measuring the weight of the eroded component before and after the erosion occurred [17, 18]. Profilometry involves measuring the eroded profile of the discharge channel walls at various stages of the Hall thruster's life. Currently, the most common and accurate approach in profilometry employs some form of laser scanning. The channel profile is typically illuminated with a laser sheet and the profile is measured with a camera [19, 20].

More advanced forms of erosion diagnostics include quartz crystal microbalance [21, 22], emission spectroscopy [22], and cavity ring-down spectroscopy [23, 24], which, until the experiment described in this dissertation, has only been implemented on anode-layer type Hall thrusters. Quartz crystal microbalance (QCM) has the potential to obtain differential sputtering profiles in tens of hours, but has many issues associated with thermal stabilization and etching of the sampling surface, and requires further development [22]. QCM is also limited to far-field interrogation and gives little information about what is happening near the sites of the channel erosion. Optical emission spectroscopy (OES) is another technique that can measure erosion data in hours to days, but is usually limited to relative density measurements. The accuracy of OES is

generally not very high because it requires accurate knowledge of the state of the plasma species in the region where the measurement is taken. Cavity ring-down spectroscopy will be described in Chapter 7.

Laser-induced fluorescence diagnostics were not originally used to study channel wall erosion when first applied to Hall thrusters. The first use of LIF in a Hall thruster environment was made by Manzella in 1994 to measure the velocity of the exhausted propellant [25]. The first use of LIF to measure propellant velocity inside the discharge channel of a Hall thruster was made by Hargus and Cappelli [26], using a small slot in the channel wall for optical access to the interrogation volume. The first LIF measurements made inside the discharge channel of an unmodified Hall thruster were made by Hargus and Nakles [27]. This dissertation will attempt to adapt LIF velocimetry to study the near-wall regions inside the discharge channel of a Hall thruster. The approach is described in more detail in Chapter 5.

In terms of theory and simulations, much is known about the process of material sputtering by bombarding particles. What has been generally lacking are the exact form of the sputtering function at low bombardment energies (<few hundred eV) and the ability to accurately determine the angular energy distribution function of the particles bombarding the channel walls. The most commonly used sputtering theories to describe the erosion of the channel wall are the Yamamura model [28] and the Zhang model [29], which is based on the Yamamura model. Yalin, et al., have also started using the modified Zhang model specifically to obtain better fits for experimental data obtained from low-energy xenon bombardment of Hall thruster wall material [21, 30]. Other models that are still often used include the original theory of sputtering by Sigmund [31], and a modification of the Sigmund theory to better fit low-energy sputtering by Bohdansky [32]. These models accept energy and angle (some models accept only energy) as inputs and outputs the sputter yield. They generally include constants that must be obtained by curve-fitting to experimental data. These models also all feature a threshold energy under which negligible sputtering occurs. The current state-of-the-art in Hall thruster channel wall erosion simulation is a 2-D hybrid code called HPHall2, which treats the ions as particles and the electrons as a fluid [33, 34]. The erosion rate is calculated using the Yamamura model fitted to Garnier's data [35, 36].

Aside from studying the problem of Hall thruster channel wall erosion, there is also active research into methods for reducing erosion. There are currently two approaches in use. The first approach is to design the magnetic field topology such that the bulk plasma is pushed as far down stream as possible. The idea is that the channel walls will eventually erode away enough that a terminal profile, which does not expose the magnetic circuit element, is reached. This is the approach used in the design of the BPT-4000 [11]. The second approach incorporates channel walls that creep forward with time so that eroded sections of the channel walls are being continuously replenished. This is the approached used in the design of the NASA 103M.XL [37]. Both approaches have shown great promise and are being actively studied.

1.6 Overview of Dissertation

The following is a brief overview of the remainder of this dissertation. Chapter 2 starts with some background information on why electric propulsion is promising for interplanetary exploration. It then transitions into a comparison of the Hall thruster with other common types of electric propulsion, and describes the unique role that only Hall thrusters can play. The chapter then goes over basic operational principles of the Hall thruster, and ends with a detailed description of the operational history of the 6-kW Hall thruster, which is the main test article for this dissertation.

Chapter 3 goes over the basics of xenon spectroscopy and various line broadening mechanisms that can interfere with the extraction of VDFs from LIF data. This chapter highlights some of the key pieces of information that are needed for processing LIF data. Several experiments that were performed to gather the said information are described in Chapter 4.

Chapter 5 begins with a description of how LIF velocimetry works, followed by the experimental setup for the LIF tests. This chapter then describes the details of how VDFs are extracted from the raw LIF data. The chapter ends with velocimetry results for both the singly-charged ion and the neutral propellant.

Chapter 6 goes over the basics of boron spectroscopy that pertain to the CRDS test. Chapter 7 begins the description of the CRDS test with an explanation of the basic operating principles. This is followed by descriptions of the experimental setup, data

reduction, and results. At the end of Chapter 7, additional ways to apply CRDS for Hall thruster research and development that are beyond the scope of this dissertation are briefly mentioned.

Chapter 8 begins by describing ways to combine LIF and CRDS data with erosion models and published wall probe data, to gain better insights into the physics of Hall thruster channel wall erosion. This chapter then identifies key parameters that drive the channel wall erosion phenomenon, and provides an assessment of how each parameter influences the erosion rate. The chapter ends with an attempt at creating a first-order engineering model for scaling the erosion rate as a function of the key parameters.

This dissertation concludes with Chapter 9, which contains a summary and a description of possible future work.

Chapter 2
Background

This chapter reviews the background information behind electric propulsion (EP) and the specific test article studied in this dissertation. The chapter begins with an overview of basic rocket principles and the physical reasons for the importance of developing electric propulsion devices. Next, it compares various common forms of electric propulsion and explains the advantages specific to Hall thrusters. After that, the chapter goes over the basics of how a Hall thruster operates. It ends with information and history of the specific 6-kW Hall thruster studied in this dissertation.

2.1 Basic Principles of Rocket Propulsion

The most defining engineering characteristic of a rocket is that it must carry its own propellant along. At our current level of technological development, a rocket is the most practical means of traveling through space. There is currently no technology in widespread use for space travel that does not require the use of propellant. Propellantless propulsion technologies like space tether [38], solar sail [39], and others are currently in active development.

The implications of the need for propellant can be seen in Tsiolkovsky's rocket equation,

$$\frac{M_f}{M_i} = \exp\left(-\frac{\Delta v}{u_e}\right) \tag{2-1}$$

where M_f is the final mass of the spacecraft after the rocket burn, M_i is the initial mass before the burn, Δv is the change in velocity of the spacecraft induced by the burn, and u_e is the average exhaust velocity of the propellant used during the burn. The exhaust velocity of the propellant can be related to the specific impulse, defined as

$$Isp = u_e / g_0 \tag{2-2}$$

where g_0 = 9.8 m/s^2 is the standard gravitational acceleration constant at the Earth's surface. The minimum Δv is typically fixed for a given mission. The Δv will vary slightly depending on the type of rocket engine chosen but the rough amount is a constant and is therefore a good representation of how much impulse is needed to deliver a spacecraft to its destination. For a given mission and a given payload mass, we can conclude from equations (2-1) and (2-2) that rocket engines with higher specific impulse require less propellant.

Figure 2-1 illustrates the importance of having a high specific impulse engine for interplanetary missions. This figure shows clearly that chemical rockets, being limited by physics to ~350 s specific impulse for storable propellants, are not a practical means of traveling beyond Mars. It is possible to extend the practical range of chemical rockets via gravity assists. However, incorporation of gravity assists puts two strong constraints on the mission: first, the launch time window becomes much narrower and specific in time; second, the trip time is greatly increased.

On the other hand, electric rockets are capable of thousands of seconds in specific

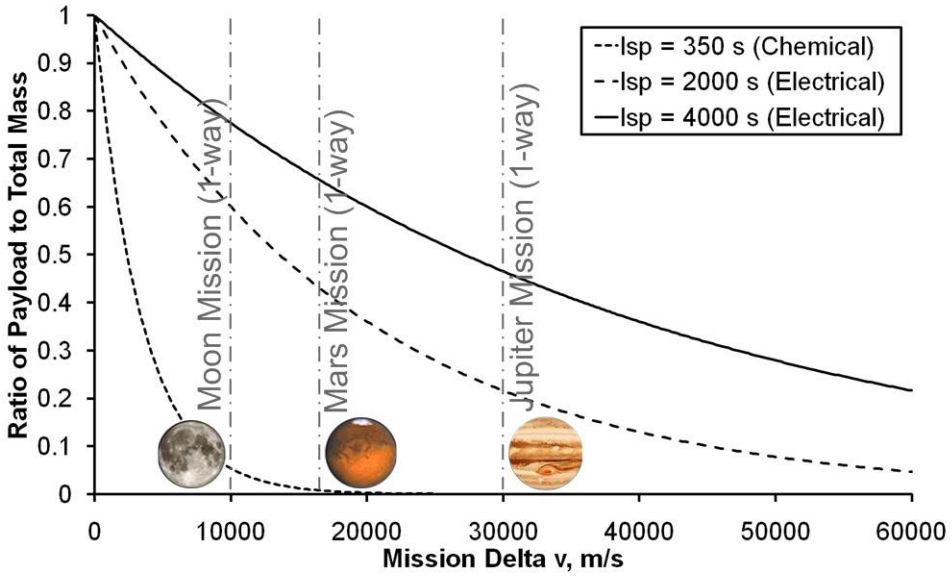

Figure 2-1. Ratio of payload to total mass versus mission delta v (no gravity assist). *Delta-v for Moon mission is based on the Earth-to-Moon transfer leg of the Lunar Reconnaissance Orbiter mission* [40]. *Delta-v for Mars mission is based on the Earth-to Mars transfer leg of a proposed Mars sample return mission* [41]. *Delta-v for Jupiter mission is based on the Earth-to-Callisto transfer leg from the preliminary design of a Jupiter Galilean moon orbiter mission* [42].

impulse, and make the delivery of large payloads (>10 ton) to planets beyond Mars practical. This high level of specific impulse can also extend the life of satellites by allowing electric rockets to use much less propellant than chemical rockets for a given orbit maintenance maneuver. The downside of electric propulsion devices is that they typically generate much lower thrust than chemical rockets, which can extend trip time.

The thrust provided by a rocket engine can be related to the average exhaust velocity as follows,

$$T = \dot{m} u_e \tag{2-3}$$

where $\dot{m}$ is the propellant mass flow rate. The propulsive power of a rocket engine can be calculated as follows.

$$P = \frac{1}{2}\frac{T^2}{\dot{m}} = \frac{1}{2} T * u_e = \frac{1}{2} T * Isp * g_0 \tag{2-4}$$

For an electric rocket, the available power is typically fixed. Thus, from equation (2-4), we can deduce that the rocket's specific impulse is in inverse proportion to the rocket's thrust. To achieve the high specific impulse needed to enable long distance missions, electric rockets typically, for the time being, output low levels of thrust. There are two important implications that follow from the preceding statement. First, spacecraft with an electric rocket tends to require a long time to build up speed. The trip time tends to be longer for an electrically-propelled craft than for a chemically-propelled craft when the required delta-v is low (e.g., missions to the Moon and Mars). Second, an electric rocket needs to operate almost continuously for the duration of the mission and must be robust enough to operate for a very long time (several thousand hours or more).

From the equations in this section, it can be concluded that the electric rocket has its own unique strengths and weaknesses when compared to the chemical rocket. While this technology is enabling for a host of long distance planetary missions, the thrust duration can be quite long, thus requiring the electric rocket to be robust.

2.2 Comparison of Electric Propulsion Devices

There are a large variety of electric propulsion devices currently in use or in development. This section will focus on only the devices that have been flown in space as of the writing of this document. Table 2-1, based on works by Larson and Wertz [43] and Goebel and Katz [13], compares key characteristics of flight-like electric propulsion technologies. Since this table shows only numbers from systems that have flown, it is not an accurate representation of the state of the art, especially in the case of ion and Hall thrusters. However, it illustrates the niche that Hall thrusters occupy in relation to other electric propulsion technologies. The specific impulse provided by a Hall thruster is high enough to enable long range planetary missions, though less than the specific impulse available with an ion thruster. In exchange, the Hall thruster can provide a greater thrust for a given amount of input power while maintaining good efficiency. Current research on the NASA HiVHAc and the BPT-4000 has also demonstrated wide throttle-ability and specific impulse up to 3000 s. These features make Hall thrusters ideal for station keeping, maneuvers within a planetary gravity well, as well as missions that require both high levels of thrust and specific impulse.

In terms of hardware complexity, the Hall thruster is simpler than the ion thruster but more complex than the resistojet, arcjet, and typical pulse plasma thruster (PPTs) [13]. The reduced complexity when compared to an ion thruster also means reduced power system mass, as fewer high-voltage supplies are needed to run a Hall thruster than an ion thruster.

The typical propellant for a Hall thruster system is the noble gas xenon. A mass

Table 2-1. Comparison of various flight-like electric propulsion technologies. *Information from* [13, 43].

	Specific Impulse, s	Input Power, kW	Efficiency, %	Thrust/Power, mN/kW
Resistojet	300	0.5-1	65-90	750-900
Arcjet	500-600	0.9-2.2	25-45	100-140
Ion thruster	2500-3600	0.4-4.3	40-80	25-40
Hall thruster	1500-2000	1.5-4.5	35-60	50-60
Pulse plasma thruster	850-1200	<0.2	7-13	15-20

flow system and gas storage tank is needed for a complete Hall thruster propulsion system. The need for a propellant management system is shared by the resistojet, arcjet, and ion thruster, but not PPT, which typically uses solid Teflon as propellant [13].

When all of the above information is accounted for, a Hall thruster propulsion system is not simple, but is versatile and can provide unique capabilities that other electric propulsion systems cannot.

2.3 Basic Principles of the Hall Thruster

A typical Hall thruster is made of three primary components: the chassis, the anode, and the cathode. Figure 2-2 shows a diagram of a Hall thruster that indicates the most common components. The chassis includes the body, magnetic coils, and the discharge channel. The body is typically made of iron to help shape the externally applied magnetic field and provide mechanical support for all other components. The most common body configuration is shaped roughly like a thick disk with a single annular discharge channel on the front face. The channel width is usually much smaller than the annulus diameter.

There will typically be two magnetic coils, an inner coil and an outer coil, though some configurations include trim coils. The inner coil, located radially inward from the channel, is typically a single spool. The outer coil, located radially outward from the channel, can be made of multiple spools. Some designs will incorporate a trim coil towards the back of the thruster (the channel opening being the front). The combination of the magnetic coils and the iron part of the body is called the magnetic circuitry.

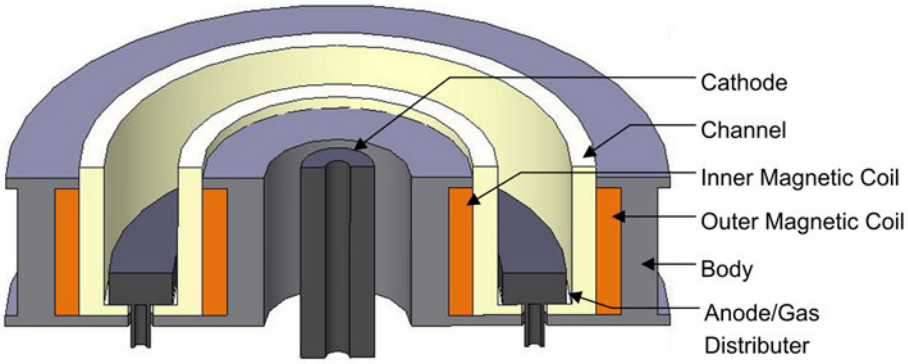

Figure 2-2. Diagram of a typical Hall thruster.

The main purpose of the magnetic circuitry is to form a roughly radial magnetic field in the annular discharge channel. This discharge channel will typically have walls that are separate from the main body to serve as protection for the magnetic iron from the energetic particles in the plasma. The channel wall material can be a ceramic, typically boron nitride, or a metal. Thrusters that incorporate ceramic walls are called magnetic-layer Hall thrusters while those that use metal walls are called anode-layer Hall thrusters. This dissertation is primarily focused on the magnetic-layer Hall thruster.

The anode is typically made of stainless steel to provide robustness and avoid interaction with the magnetic circuit. It usually also doubles as the gas distributor for the primary propellant gas, which is typically xenon. The anode is biased to a high voltage compared to the other components during operation. The cathode comes in a large variety of types and configurations. Its primary function is to act as an electron source and to complete the plasma discharge circuit. The most common type of cathode is a hollow cathode with either barium oxide or lanthanum hexaboride acting as the primary emitter. Both center-mounted (as shown in Figure 2-2) and side-mounted configurations exist. The main purpose of the anode and the cathode is to set up a roughly axial electric field in the discharge channel of the Hall thruster.

Figure 2-3 shows the basic principles of how the Hall thruster functions. As previously mentioned, there are strong axial electric and radial magnetic fields in the

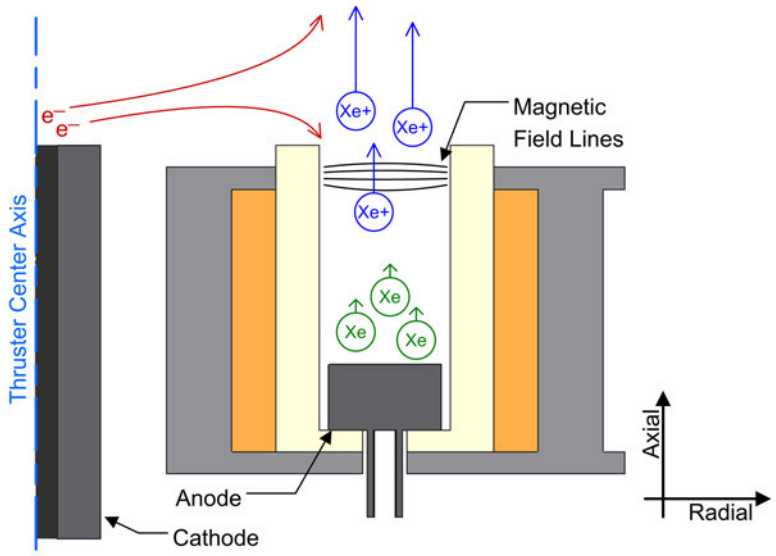

Figure 2-3. Diagram of basic Hall thruster operation principles.

channel. These fields are typically the strongest near the exit of the discharge channel. The cathode begins by providing a stream of electrons that are attracted by the positive bias on the anode. These electrons encounter the strong electric and magnetic fields near the channel exit and become trapped in an ExB drift, forming what is known as the Hall current from where the Hall thruster get its name. At the same time, neutral xenon gas is being released from the anode and drifts toward the channel exit. When these neutrals meet the Hall current, they are impact ionized by the kinetic energy that the electrons have acquired while traveling in the electric field. The newly born ions immediately see the large axial electric field present near the channel exit and rapidly accelerate out of the channel. This acceleration imparts a reaction force on the thruster propelling the thruster in the reverse direction. To keep the thruster charge neutral (so that the ions are not drawn back to the thruster and cancel out the thrust), the cathode also provides a steady stream of electrons that leaves with the outbound ions. Since both electrons and ions are present in great numbers in the acceleration zone at the same time, the plasma in the discharge channel can be kept quasi-neutral, thus avoiding current limits imposed by space charge effect.

As mentioned previously, the magnetic field is mostly radial and typically strongest near the channel exit plane. If the field is perfectly radial, the plasma is poorly confined and will bombard the channel walls causing rapid erosion, as was the case in early Hall thruster prototypes. Later designs use a magnetic field topology that is called a "lens", due to its ability to focus plasma away from the walls. Figure 2-4 shows the magnetic field topology of the NASA-173Mv1 [44], which

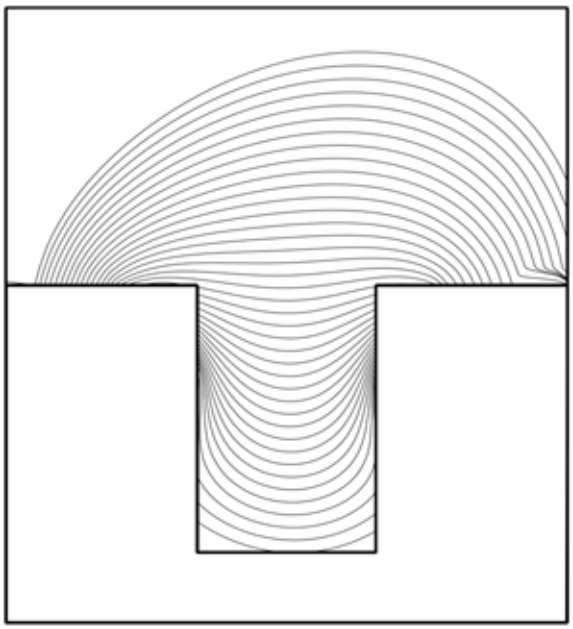

Figure 2-4. Magnetic field topology of the NASA-173Mv1 taken from [44].
The image is rotated to match the orientation used in the rest of this dissertation.

is representative of more recent Hall thruster designs. The figure shows that the magnetic field is only close to radial right outside the exit plane of the discharge channel, and is bowing away from the exit plane everywhere else. The basis for the "lens" topology comes from work done by Alexei Morozov that demonstrated the magnetic field lines are approximately electric equipotential field lines [45]. The underlying concept is that electrons have a much higher mobility traveling along the magnetic field lines than across them. Thus, the path along a magnetic field line is electrically shorted in comparison to the path perpendicular to it, and the electric potential is roughly the same along the magnetic field line.

In theory, the magnetic "lens" topology should focus ions toward the center of the channel as they travel downstream towards the exit. In reality, the equipotential field lines do not follow the magnetic field lines that closely. The biggest correction factor comes from the thermalizing potential associated with electron density gradient, which is shown in Morozov's calculations [45]. For a typical Hall thruster, the equipotential lines in the ion acceleration zone diverge most strongly from the magnetic field lines near the channel walls, where density gradients are large compared to the center of the channel. Thus, even with a magnetic "lens" it is possible for equipotential lines near the wall to be defocusing. The physics described in this section are important to keep in mind when looking at velocimetry data.

2.4 The 6-kW Hall Thruster

This section describes the specific Hall thruster studied in this dissertation. The 6-kW Hall thruster was originally conceived of as a state-of-the-art (SotA) research platform in a collaborative effort by the Air Force Research Laboratory (AFRL), the NASA Jet Propulsion Laboratory (JPL), and the University of Michigan (UM) Plasmadynamics and Electric Propulsion Laboratory (PEPL). The goal was to rapidly accumulate knowledge on a SotA design by simultaneously performing research on the exact same design at multiple institutions. All subsequent references to the 6-kW Hall thruster refer to the UM copy of the thruster.

The 6-kW Hall thruster is equipped with a single spool for the inner magnetic coil, eight spools evenly spaced in the azimuthal direction form the outer magnetic coil,

and a single spool for the trim coil. For the purpose of the experiments in this dissertation, the trim coil was not used, and the inner and outer coils were powered in a way that generated a magnetic lens topology that is symmetric about the channel centerline. The channel walls are made of boron nitride so this thruster is of the magnetic-layer type. The portion of the channel walls near the exit plane of the channel is removable so that replacement channel walls with different chamfer geometry can be inserted to simulate the effect of erosion. Unless otherwise specified, all tests related to the UM copy of the 6-kW Hall thruster use the original set of channel walls installed when the thruster was first assembled.

For the 6-kW Hall thruster, the cathode is mounted in the center of the thruster body. Two cathodes were used over the course of this dissertation. Both are hollow cathodes that use lanthanum hexaboride as the emitter material. Figure 2-5 shows a side-by-side comparison of the two cathodes. The two cathodes are nearly identical in design and the one used later is a slightly improved version of the one used earlier. Additionally, when mounted to the thruster, the cathode exit of the earlier version protrudes 2.5 mm beyond the exit plane of the discharge channel, while the cathode exit of the later version is flush with the same exit plane. The later version also has an extra orifice plate that the earlier version lacks. While the geometric difference was beyond the control of the author, performance testing showed that there is no measurable difference in the performance of the thruster when operating with the two cathodes. Nevertheless, it is noted here that the earlier version of the cathode was used in all laser-induced

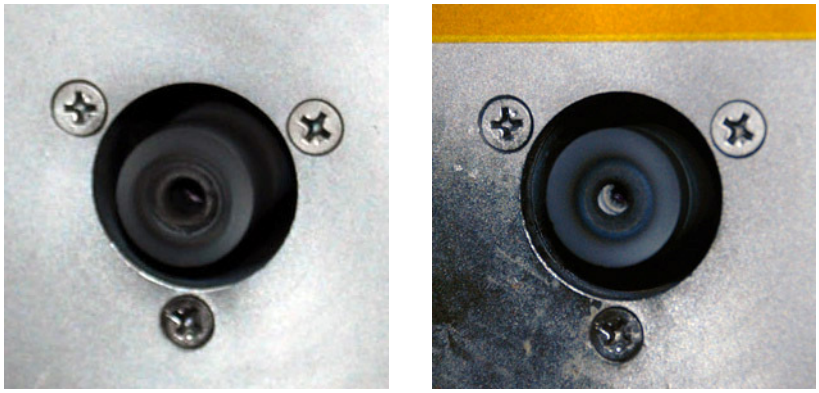

Figure 2-5. Side-by-side comparison of the two cathodes used for this dissertation mounted in the 6-kW Hall thruster.
Left image shows the version used for the LIF tests; right image shows the version used for the CRDS tests.

fluorescence (LIF) studies while the later version was used in all cavity ring-down spectroscopy (CRDS) studies.

2.4.1 Operational History

Many studies have been performed on the 6-kW Hall thruster since its inception. Figure 2-6 shows a timeline of various experiments that have been performed on the UM copy of the 6-kW Hall thruster. Each major block of operation time is marked with a letter for reference and occurs in sequence during the life of the 6-kW Hall thruster, unless otherwise specified. In particular, block H is the two-axis interior LIF study and block R is the CRDS study. These two studies are the most important dissertation-related studies and will be described in more detail in Chapter 5 and Chapter 7, respectively. Thus, the operation time starting at the beginning of block H and ending at the end of block R can be considered the part of the operation history that is the most important to this dissertation. For this reason, profilometry measurements were taken at 334 hours and 543 hours of operation, corresponding respectively to the beginning of block H and the end of block R. The profilometry measurement is described in Section 2.4.3. Next, the test and literature associated with each time block will be described.

Block A represents all of the experiments for which Bryan Reid was the principal investigator. He performed much of the initial performance and probe testing on the 6-kW Hall thruster. The majority of his work done during block A can be found in his doctoral dissertation [46]. The specific experiments performed in block A include performance testing and design verification [47, 48], high-speed axial reciprocating probing (HARP) of the plasma properties inside and in the near-field of the discharge channel [49, 50], and angularly-resolved ExB measurements, done in partnership with Rohit Shastry [51, 52].

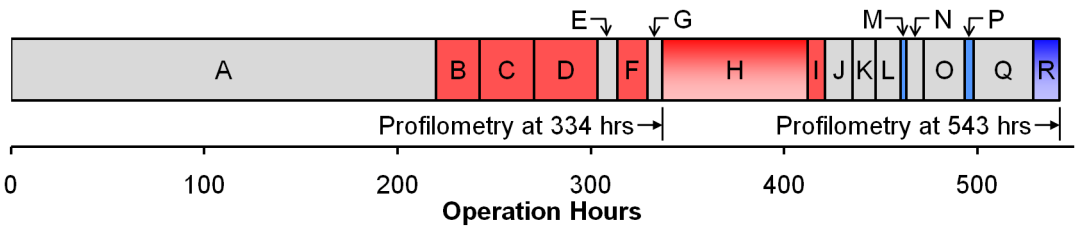

Figure 2-6. An operation timeline of the 6-kW Hall thruster.

Blocks B and C are pre-cursor studies performed prior to the formation of the dissertation problem statement to help the author gain experience in LIF velocimetry. These two experiments were performed using direct laser beam injection, as oppose to fiber-based beam delivery used in later LIF tests. Block B was an effort led by Tim Smith to learn more about neutral xenon flow in the near-field of the 6-kW Hall thruster [53]. Block C was the author's attempt to use the same setup to study xenon ions [54]. Although blocks B and C are shown in Figure 2-6 to be after block A the experiments actually took place concurrently with the end of block A. Tests represented by all other blocks take place sequentially as indicated by the associated alphabet letter.

Block D was a one-axis LIF test to map the ion behavior along the channel centerline in the acceleration zone of the 6-kW Hall thruster. Pertinent parts of this study will be described in Chapter 5, and other details can be found in [55].

Block E was an experiment by Michael McDonald to study the effects of using an external cathode while varying the cathode pointing angle and distance from the thruster [56].

Experiments in block F generated data that would eventually be used in a one-axis interior LIF study of neutral xenon flow [57] and an attempt to map the in-situ magnetic field of the 6-kW Hall thruster [58]. Parts of these studies relevant to the dissertation will also be described in Chapter 5.

Block G is a set of performance and equipment checkout tests to make sure the thruster-side of the experimental setup was working properly before the start of the main two-axis LIF test (block H). This was deemed necessary because the thruster had undergone 300+ hours of operation by this point and the performance characteristic might have changed along with the channel wall erosion pattern. Greater detail regarding this test and associated equipment can be found in Section 2.4.2.

As mentioned previously, block H is the main two-axis LIF test that will be described later in this dissertation. The primary goal of this test was to study the angular energy distribution functions of the singly-charged and neutral xenon particles that bombard the channel walls.

Blocks I to Q are various tests that took place between the two important dissertation-related tests (Blocks H and R). Due to technical challenges associated with

the development of CRDS, it was not possible to perform the two-axis interior LIF test and the CRDS test back-to-back. Block I is a set of follow-on experiments that took advantage of the setup already in place. The data gathered in these experiments, together with data from block F, produced a one-axis interior LIF study of the xenon neutral [57] and an in-situ magnetic field mapping study [58]. Additionally, Rohit Shastry performed an experiment where wall tabs were installed to gather Langmuir probe data from just downstream of the channel exit plane [59].

In block J, Mike McDonald performed a Langmuir probe experiment to map the near-field plume of the 6-kW Hall thruster in three-dimension using a high-speed (2-MHz) DAQ system [60].

In block K, Rohit Shastry performed a set of embedded wall probe experiments [61]. The operation times represented by block K were used for thruster validation. The actual test used a separate set of channel walls and the operation time associated with those special channel walls are not represented in Figure 2-6.

Block L is miscellaneous equipment validation.

Block M is the CRDS cavity thermal test. This test involves a barebone (no optics) CRDS setup sitting at the proper location with respect to the 6-kW thruster to study the head load experienced by the cavity structure during live operation. Four thermocouples were used to monitor the cavity temperatures at key locations during this test. The test results showed that thermal loading does not have a significant effect on the shape of the cavity. This knowledge drove some design decisions for version two of the CRDS cavity structure.

In block N, Rohit Shastry performed near-field Faraday probe studies on the 6-kW Hall thruster [61].

In block O and Q, Mike McDonald performed high-speed video, flush-mounted pole-face Langmuir probe, and HARP studies on the 6-kW Hall thruster. No citations are available for these tests at the time of writing.

Block P is the testing of the CRDS sensor version one. The information gained regarding the shortcomings of version one would eventually give rise to CRDS sensor version two. Greater detail on the version-one sensor can be found in Chapter 7.

As mentioned previously, block R is the main CRDS test that will be described later in this dissertation. The primary goal of this test is to study the erosion rate of the channel wall as a function of operating conditions. This test used the version-two CRDS sensor.

2.4.2 Performance of the 6-kW Hall Thruster

Performance and equipment checkout tests were performed prior to the start of the main two-axis LIF test to make sure the thruster still performed well after 300+ hours of operation. The performance data collected during this test are listed in Table 2-2. The listed anode efficiency does not account for magnet power, cathode power, and cathode mass flow. For this performance test, the currents for the magnet coils are set such that the magnetic field topology is symmetric about the channel centerline and the anode efficiency is maximized. The cathode mass flow rate is set to 7% of the anode mass flow rate. The discharge voltage and current are measured using high precision shunts and voltmeters and are accurate to ±0.1%. The mass flow rate is controlled using commercial mass flow controllers accurate to ±1%. These controllers are calibrated using the constant volume method taking into account the effect of compressibility. All experiments use 99.999% pure research-grade xenon as the propellant. Thrust is measured using an inverted-pendulum thrust stand based on the Glenn Research Center (GRC) design [62] and is accurate to ±2 mN. As the table shows, the thruster still exhibits state-of-the-art performance at 300+ hours of operation. Note that the nominal operating condition of the 6-kW Hall thruster is 300 V discharge voltage and 20 mg/s anode mass flow rate.

Table 2-2. Performance data for the 6-kW Hall thruster

Discharge voltage, V	Anode mass flow rate, mg/s	Discharge current, A	Thrust, mN	Anode efficiency
150	10	9.08	110	44%
150	20	21.55	250	48%
150	30	35.35	388	47%
300	10	8.94	180	60%
300	20	20.11	385	61%
300	30	33.25	612	63%
600	10	9.36	281	70%

The seven operating conditions listed in Table 2-2 are the throttling points for this dissertation. The discharge voltage and the anode mass flow rate are the main controls for these operating conditions and all other parameters are measured quantities. For the remainder of this dissertation, operating conditions will be referred to as YYY V, ZZ mg/s where YYY is the discharge voltage and ZZ is the anode mass flow rate.

2.4.3 Profilometry Results and Operation Time per Condition

As shown in Figure 2-6 (the thruster operation timeline) profilometry measurements were taken twice. The first set of measurements was taken prior to the main two-axis interior LIF experiment at 334 operational hours, and the second set of measurements was taken after the main CRDS experiment at 543 operational hours. These measurements were taken using a Brown & Sharpe MicroVal 343 coordinate measuring machine. The machine measures by sensing when its ruby-tip comes into contact with the sample. It has a resolution of 0.001 mm but due to limitation in the mounting scheme, the actual uncertainty is estimated to be ~0.05 mm. Figure 2-7 shows the profilometry results on a non-dimensionalized coordinate grid. The axial coordinate is non-dimensionalized by the channel length with Z = 0 at channel exit plane and Z = -1 at the anode exit plane. The radial coordinate is non-dimensionalized by the channel width with R = 0 at the nominal inner wall radius and R = 1 at the nominal outer wall radius.

Table 2-3 shows the breakdown of operation time by operating conditions. From this table, we can see that the 6-kW Hall thruster spent most of its operational hours in the 300 V, 10 mg/s and 20 mg/s conditions. All together, the 300 V operating conditions occupy roughly 70% of the total thruster run time. In this table, miscellaneous operating conditions include various conditions that are not one of the seven dissertation throttling points and spans discharge voltages from 105 V to 550 V and anode mass flow rates from 5 mg/s to 17.5 mg/s.

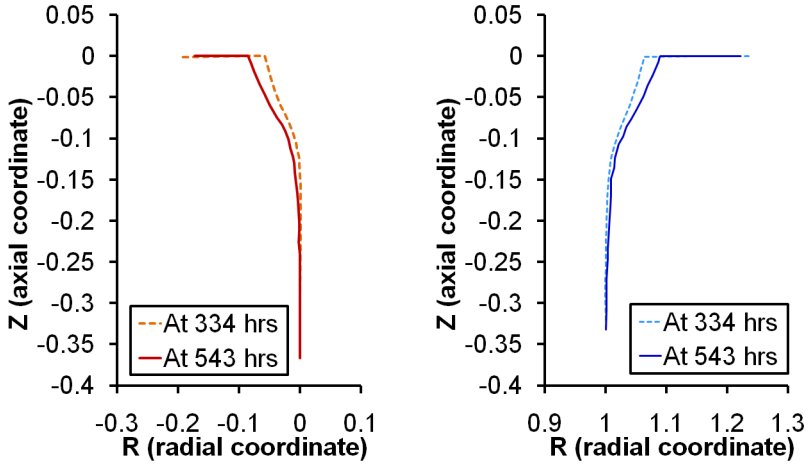

Figure 2-7. Profilometry measurements for the channel inner wall (left) and outer wall (right).

Table 2-3. Percentage breakdown of operation time by condition.

Discharge voltage, V	Anode mass flow rate, mg/s	Percentage time as of first profilometry measurement	Percentage time as of second profilometry measurement
150	10	7.5%	7.2%
150	20	5.7%	5.8%
150	30	5.8%	5.9%
300	10	30.6%	29.4%
300	20	29.5%	28.7%
300	30	10.7%	9.5%
600	10	6.4%	7.8%
Misc		3.8%	5.7%

Chapter 3
Physics of Xenon Spectroscopy

3.1 Basics of Xenon Transition Lines

The use of xenon spectroscopy, beyond simply observing the emitted spectra, is a relatively new diagnostic tool in the study of Hall thruster. The first work to use laser-based xenon spectroscopy to study Hall thrusters in the United States was a 1994 experiment by Manzella to study the velocity of the singly-charged ion [25]. Since then, significant advancement has been made in both the understanding of the spectroscopic parameters that govern xenon transitions and the development of laser-induced fluorescence (LIF) as a diagnostic for studying Hall thrusters. This chapter focuses on our current state of knowledge regarding useful xenon transitions for laser-based xenon spectroscopy and various line broadening mechanisms that may need to be accounted for when designing an LIF experiment.

Three basic types of atom-photon interactions are possible. The first is photon-induced excitation, where the incident photon energy matches the difference in energy between two energy states of an atom and the atom is excited to the higher state. The second is spontaneous emission, where an atom sitting at a high energy state spontaneously decays to a lower energy state, releasing a photon equal to the energy difference in the process. The third is stimulated emission, where a passing photon causes an atom of the right energy state to decay and release another photon of the exact same energy, direction, and phase. For the purpose of LIF velocimetry, we are primarily concerned with the first two interactions.

LIF velocimetry operates by first exciting an atom up to an excited state via photon absorption and then collecting the photon that is emitted during spontaneous decay. One advantage of using LIF as opposed to emission spectroscopy is that one can add a carrier frequency to the injected laser beam using some form of chopper. Then one

can use lock-in amplification on the extracted fluorescence signal to raise the signal-to-noise ratio (SNR) by several orders of magnitude over emission spectroscopy. Furthermore, the use of the laser allows the creation of spatially well defined interrogation zones.

The exact excitation and decay transition lines that are used for an LIF study are common called an LIF "scheme". A scheme that involves using the same transition for both excitation and decay is described as being resonant. A scheme that uses different excitation and decay transitions is described as being non-resonant. A non-resonant scheme is generally more advantageous than a resonant scheme because a resonant scheme cannot differentiate between LIF signal and surface reflection. Furthermore, certain energy states are considered forbidden in quantum mechanics. These states have a much lower decay rate to the ground state than regular excited states and are called metastable states. As a result, the number of atoms in these metastable states tends to be quite high compared to other excited states for a well-energized plasma. A scheme that uses a metastable lower state has the advantage of having a relatively large lower state population to excite and will generally give a higher SNR than other schemes. An LIF scheme can be any combination of resonant/non-resonant and metastable/not metastable.

Table 3-1 shows commonly used LIF schemes for xenon spectroscopy. Figure 3-1 shows the corresponding transition diagram. Although there are many other LIF schemes possible, many are inaccessible by commercially available lasers. The schemes listed here are all readily accessible by relatively cheap diode lasers with the exception of the Xe II 605.3 nm scheme, which is accessible by a dye laser. All wavelengths in this dissertation are vacuum wavelengths unless otherwise specified. Per convention, the

Table 3-1. Common LIF schemes for Hall thruster studies.

Source: NIST atomic spectra database, ver. 4, Sep. 2010.

Species	Start State	Excitation Wavelength (vac), nm	Excited State	Decay Wavelength (vac), nm	End State
Xe I	$6s'[1/2]^0_1$	834.912	$6p'[3/2]_2$	473.548	$6s[3/2]^0_1$
Xe I	$6s[3/2]^0_2$	823.390	$6p[3/2]_2$	823.390	$6s[3/2]^0_2$
Xe I	$6s[3/2]^0_1$	828.239	$6p[1/2]_0$	828.239	$6s[3/2]^0_1$
Xe II	$5d[4]_{7/2}$	834.953	$6p[3]^0_{5/2}$	542.066	$6s[2]_{3/2}$
Xe II	$5d[3]_{7/2}$	605.283	$6p[2]^0_{5/2}$	529.369	$6s[2]_{5/2}$

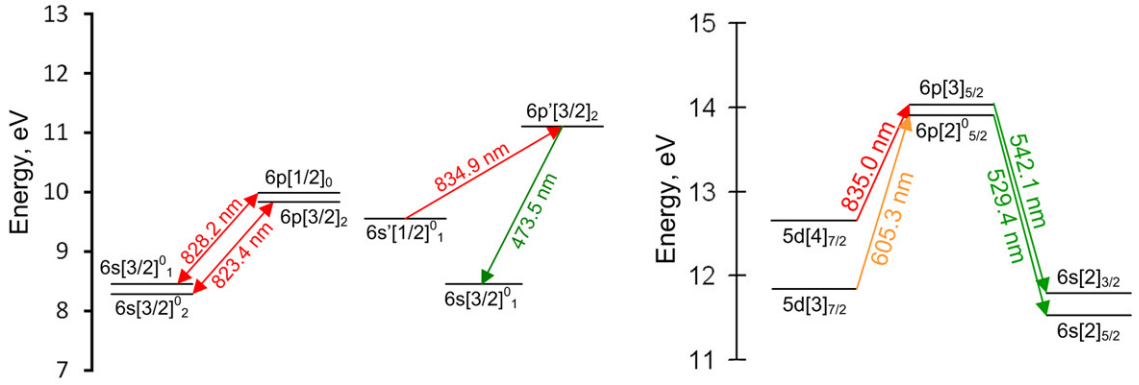

Figure 3-1. Transition diagrams for common neutral (left) and singly-charged (right) xenon LIF schemes.

species designation Xe I refers to neutral xenon and Xe II refers to singly-charged xenon. The most pertinent LIF schemes to this dissertation are the Xe I 834.9 nm, Xe I 823.4 nm, and the Xe II 835.0 nm schemes. Note the excitation wavelength of the Xe II 835.0 nm scheme is listed as 834.953 nm in Table 3-1. The Xe I 834.9 nm and Xe II 835.0 nm schemes are used to extract VDFs for their respective species. The Xe I 823.4 nm scheme is used mainly to troubleshoot the LIF setup and to study neutral density. We now briefly cover the advantages and disadvantages of each scheme and why certain schemes are better suited for certain types of measurements.

The Xe I 834.9 nm scheme has not been in use for as long as the other listed schemes because it is not as easy to excite as the other neutral transitions. However, it is also not as easy to saturate with laser light. The saturation effect will be described in greater detail in Section 3.2.4. For the time being it is enough to know that the saturation effect distorts LIF lineshapes making the interpretation of velocity distribution functions (VDFs) difficult to impossible. Additionally, the Xe I 834.9 nm scheme listed in Table 3-1 is non-resonant and is ideal for measurements near reflective surfaces. For experiments where the neutral xenon VDF is the desired measurement this is a great scheme.

The Xe I 823.4 and 828.2 nm schemes were first used to probe neutral xenon by Cedolin in 1996 [63, 64]. The early adoption of these schemes in the history of LIF studies on Hall thrusters is due to the relative ease of laser absorption by the lower state. In particular, the lower state of the Xe I 823.4 nm transition is metastable. This transition is strongly responsive to even the weakest of laser beams and is a good way to

troubleshoot an LIF setup. However, the strength of this line also means that it is highly susceptible to saturation broadening and is not useful for VDF measurements. It is better suited for bulk velocity, peak velocity, and, under the right conditions, density measurements. There is a non-resonant LIF scheme for the Xe I 823.4 excitation transition. It involves collecting fluorescence at 895.5 nm with the end state being $6s[3/2]^0_1$ (the initial lower state of the Xe I 828.2 nm scheme). However, as part of this dissertation, LIF traces were taken using an optogalvanic cell that showed the collected signal for this non-resonant scheme is three orders of magnitude lower than for the resonant scheme.

The Xe I 828.2 nm scheme differs from the 823.4 nm scheme in that the lower state is not metastable and the hyperfine structures are half as narrow (~2 GHz wide for 828.2 nm). Nevertheless, it still responds fairly strongly to a laser and is susceptible to saturation broadening. Hyperfine structures will be discussed in greater detail in Section 3.2.2. For now it is sufficient to know that hyperfine structures distort lineshape and generally must be accounted for when extracting VDFs.

The Xe II 835.0 nm scheme is one of two commonly used LIF velocimetry schemes in the study of Hall thrusters. This scheme has gained popularity in recent years due to the ease with which it can be accessed using a relatively inexpensive diode laser. The disadvantage (and, under certain specific conditions, the advantage) of using this transition is that the hyperfine structures are so tightly packed together that no accurate determination of the hyperfine structure constants has been made to date. The width of the hyperfine structures, in the absence of strong external electric and magnetic fields, has been narrowed down to ~450 MHz during the course of this dissertation. Xenon ion VDFs regularly reach several GHz in width in the frequency domain. The implication, shown in [65], is that for large regions inside and in the plume of the Hall thruster, the hyperfine structure can be modeled as a simple Gaussian or ignored all together. The key to getting the correct VDFs is to know where the hyperfine structures can be ignored and where they can be approximated. This is described in greater detail in Chapter 4. The lower state of the Xe II 835.0 nm is metastable and the scheme is non-resonant. The response of this line to a laser is not as strong as the 823.4 nm and 828.2 nm neutral lines but is also generally not easy to saturate, making it ideal for extracting VDFs.

The Xe II 605.3 nm scheme is the other commonly-used LIF scheme. The associated lower state is metastable, the scheme is non-resonant, the hyperfine structure constants are known, and the transition responds well to a laser. The only downside to using this scheme is that the wavelength is not easily accessible by commercial diode lasers. In the past, researchers have used a dye laser to access this line and had to deal with all the accompanying difficulties. The other potential downside is that the associated hyperfine structure is ~4 GHz wide, so there are no conditions under which the hyperfine structure can be safely ignored.

Next, we will discuss various line broadening mechanisms that may need to be accounted for when extracting VDFs from raw LIF lineshapes.

3.2 Line Broadening Mechanisms

Although energy states are generally assigned a single average energy on a macroscopic scale, their spectra are typically not sharp peaks when you zoom in to the microscopic scale. There are a large number of physical phenomena that can give rise to lineshape broadening and distortion. For example, Figure 3-2 shows the Xe I 834.912 nm line as it appears on a macroscopic scale (left) and on a microscopic scale (right). To extract VDFs from LIF raw traces, one must properly account for all broadening and distorting factors. Note that although laser linewidth is a broadening mechanism, the

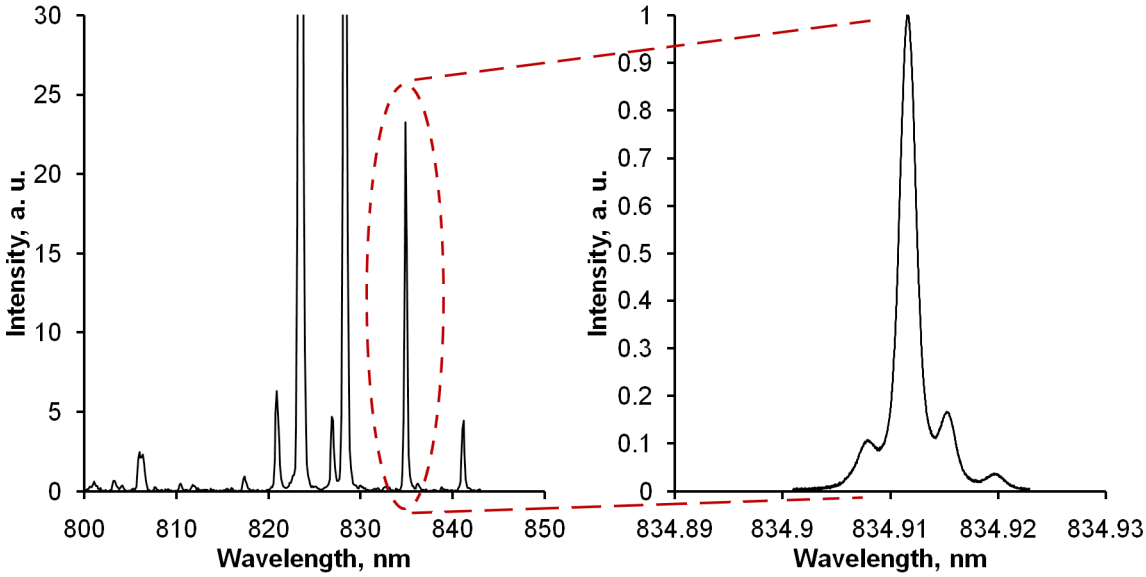

Figure 3-2. General example of how line broadening influence spectral lineshape.

linewidth of the lasers used in this dissertation is negligibly narrow compare to other broadening factor and will not be discussed in detail.

3.2.1 Doppler Shift and Broadening

When an atom is moving in the direction of the photon's wave vector, the atom absorbs the photon at a shifted frequency. This is known as the Doppler effect. The governing equation for the Doppler effect is,

$$\frac{\Delta\upsilon}{\upsilon_0} = -\frac{\mathbf{v}}{c}\cdot\frac{\mathbf{k}}{|\mathbf{k}|} \tag{3-1}$$

where υ_0 is the photon frequency, $\Delta\upsilon$ is the shift in photon frequency from the perspective of the particle, **v** is the particle velocity, c is the speed of light, and **k** is the photon wave vector. In fact, this is the principle on which LIF velocimetry is based. Both random and directed velocity will show up in the LIF spectra as having shifted from the stationary wavelength. By varying the frequency of the injected laser light and measuring the intensity of the collected fluorescence, one can construct the VDF of the probed species.

Since the Doppler shift equation, along with a number of convenient spectroscopic equations, is simpler in the frequency domain, subsequent discussion will mostly be in frequency space. A particularly common convention is to plot the intensity of a lineshape against what is called the "detuning". The term stems from tuning of sound frequencies in music and is basically the frequency difference between the current frequency and an arbitrary reference marker. The purpose of detuning is usually to show the relative frequency spacing between various spectral features, and the choice of reference marker is not important. Most of the plots in this dissertation use either the intensity-averaged frequency or the peak frequency of a spectrum as the reference marker when detuning is used.

To describe Doppler shift and broadening (as well as a number of other broadening phenomena) in mathematical terms, one can use the mathematical operation called convolution. The equation for the convolution operation is shown in equation (3-2).

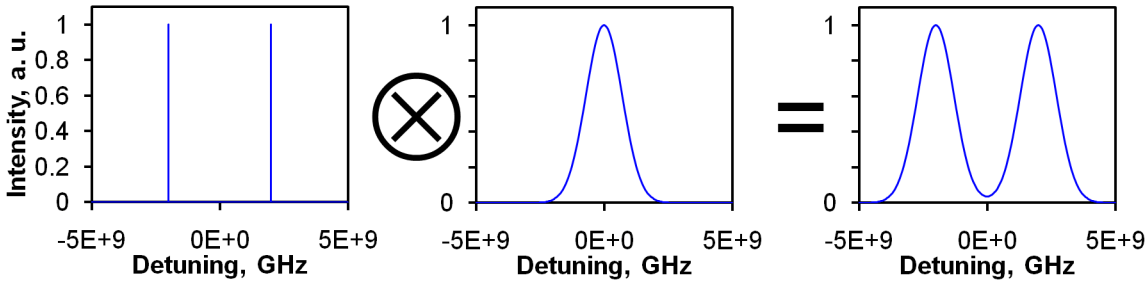

Figure 3-3. Simple graphical example of how convolution works.

$$C = A \otimes B \quad \Leftrightarrow \quad C(x) = \int_{-\infty}^{\infty} A(x-y)B(y)dy \tag{3-2}$$

The symbol $\otimes$ represents the convolution operation. Figure 3-3 shows, graphically, what happens when a function is convolved with another function. Intuitively, it is almost like multiplying two graphs together to form a third graph. In this example, the leftmost plot shows two spectral lines spaced 4 GHz apart. The middle plot shows a typical Maxwellian distribution found in room-temperature plasma. The rightmost plot shows what one would see if one used LIF to scan this line, assuming no other broadening factor is important.

The inverse of the convolution operator is called the deconvolution operator, represented by the $\varnothing$ symbol. Just as nature convolves multiple broadening factors together to form the spectral lineshape, we can use the deconvolution operation to remove the effect of these broadening factors and extract specific pieces of information. For example, in Figure 3-3, if we wanted to extract the VDF (middle plot) from the raw LIF trace (right plot), we could do so using deconvolution techniques as long as we knew a priori what the broadening factor (left plot) looks like. Conversely, since convolution is commutative, the VDF (middle plot), if known, can be treated as the broadening factor and the base line structure (left plot) can be extracted. More detail about deconvolution methods can be found in Section 5.3.3.

3.2.2 Hyperfine Structure and Isotopic Shifts

Hyperfine structure is another major broadening factor in laser-based xenon spectroscopy. The exact amount and nature of the splitting varies quite a bit from transition line to transition line, but is generally wide enough to interfere with the extraction of VDFs from raw LIF lineshapes.

Hyperfine structure arises primarily out of the coupling between the intrinsic electric and magnetic fields of the electrons and the nucleus in the atom. Paired nucleons have zero net spin because the spin vectors of nucleons in a pair point in opposite directions. The atomic number of xenon is 54 so the number of protons is even. Of the nine stable isotopes of Xenon, only two (Xe-129 and Xe-131) have an odd number of neutrons. The energies of the otherwise degenerate states in these two isotopes will shift slightly, leading to hyperfine structure. The shift in energy is described by equation (3-3),

$$\begin{aligned} E_{hfs} &= A\frac{C}{2} + BD \\ C &= F(F+1) - I(I+1) - J(J+1) \\ D &= \frac{(3C/4)(C+1) - I(I+1)J(J+1)}{2I(2I-1)J(2J-1)} \end{aligned} \tag{3-3}$$

where A is the magnetic dipole interaction constant, B is the electric quadrupole interaction constant, I is the nuclear spin angular momentum quantum number, J is the total electron angular momentum quantum number, and F is the total resultant angular momentum quantum number and obeys equation (3-4). The selection rule for allowed hyperfine structure transitions is in equation (3-5).

$$F = I + J, I + J - 1, ..., |I - J| \tag{3-4}$$

$$\Delta F \equiv F_U - F_L = [0, \pm 1],\ F_L{=}0 \nrightarrow F_U{=}0 \tag{3-5}$$

In equation (3-5) F_L and F_U are the total angular momentum quantum numbers for the lower and upper states, respectively. The zero to zero transition is forbidden. For a given xenon transition and a given xenon isotope, the values of I, J_L, and J_U are fixed. The symbols J_L and J_U denote the total electron angular momentum of the lower and upper states, respectively. However, the value of F for the lower and upper states can still independently take on the range of values shown in equation (3-4) because F represents different ways in which I and J couple vectorially. This is the mathematical basis for how hyperfine structures can arise for a nucleus that has an odd number of nucleons. A nucleus with an even number of nucleons has I = 0.

The calculations for the relative intensity of the individual split transitions in the hyperfine structure differ depending on the value of J_L and J_U. The relative intensity for a J→J transition is given by equations (3-6), (3-7), (3-8), and (3-12) [66]. The relative

intensity for a J-1→J transition is given by equations (3-9), (3-10), (3-11), and (3-12) [67].

$$\mathrm{Int}(F \to F-1) \sim P(F)Q(F-1)/F \tag{3-6}$$

$$\mathrm{Int}(F \to F) \sim \frac{2F+1}{F(F+1)}[R(F)]^2 \tag{3-7}$$

$$\mathrm{Int}(F\text{-}1 \to F) \sim P(F)Q(F-1)/F \tag{3-8}$$

$$\mathrm{Int}(F \to F-1) \sim Q(F)Q(F-1)/F \tag{3-9}$$

$$\mathrm{Int}(F \to F) \sim \frac{2F+1}{F(F+1)}P(F)Q(F) \tag{3-10}$$

$$\mathrm{Int}(F\text{-}1 \to F) \sim P(F)P(F-1)/F \tag{3-11}$$

$$\begin{cases} P(F) = (F+J)(F+J+1) - I(I+1) \\ Q(F) = I(I+1) - (F-J)(F-J+1) \\ R(F) = F(F+1) + J(J+1) - I(I+1) \end{cases} \tag{3-12}$$

In addition to the nucleus-electron electromagnetic interactions, variations in nuclear mass and volume cause slight variations in the energy of the electron orbital states. Each stable isotope of xenon has an associated shift in transitional energy described by the parameter called the isotopic shift. The relative intensity of each transition is proportional to the abundance of the associated isotope. There are quantum

Table 3-2. Magnetic dipole, electric quadrupole, and total electron momentum quantum number for the Xe I 823.4 nm and Xe I 834.9 nm transitions.
The values of these constants are from the references [68-70].

	Lower state constants, MHz				Upper state constants, MHz			
Transition	A_{129}	A_{131}	B_{131}	J_L	A_{129}	A_{131}	B_{131}	J_U
Xe I 823.4 nm	-2384.4	705.4	255.7	2	-889.6	262.7	21.3	2
Xe I 834.9 nm	-5808.0	1709.3	30.3	1	-2894.6	858.9	-14	2

Table 3-3. Isotopic shift for the Xe I 823.4 nm and Xe I 834.9 nm transitions.
The values of isotopic shift constants are from the reference [69].

	Isotopic Shift, MHz						
Transition	Xe-128	Xe-129	Xe-130	Xe-131	Xe-132	Xe-134	Xe-136
Xe I 823.4 nm	117.2	123.7	52.3	94.7	0	-52.5	-150.3
Xe I 834.9 nm	100.5	88.5	43.5	74.3	0	-30.0	-115.4

models for isotopic shift, but the orbit calculations for an atom as complex as xenon are computationally prohibitively expensive so isotopic shifts for xenon are typically found by empirical methods.

The hyperfine structure and isotopic shifts typically cause broadening of similar magnitude in the lineshape. Thus, the two effects must often be accounted for at the same time. The literature sometimes refers to both effects together as the hyperfine structures. For simplicity, this dissertation will also do the same.

For the three LIF schemes that are pertinent to this dissertation, only the hyperfine structure constants of the Xe I 823.4 nm and the Xe I 834.9 nm lines are known. Table 3-2 lists the nucleus-electron interaction constants. Note that for Xe-129 and Xe-131, I equals 1/2 and 3/2, respectively. Table 3-3 lists isotopic shifts for seven of the nine stable xenon isotopes. Isotopic shift information is not available for Xe-124 and Xe-126, but they each make up less than 0.1% of the typical xenon composition and can be safely neglected [71].

Since the hyperfine structure constants for the Xe II 835.0 nm transition are not known, a series of experiments were completed as a part of this dissertation to obtain a good approximation for the hyperfine structure. These experiments are described in detail in Chapter 4.

3.2.3 Lifetime Broadening

Lifetime (or natural) broadening can be described by the Heisenberg uncertainty principle. Since energy and temporal measurements are complementary, one cannot know with absolute certainty the exact energy of a photon emitted from a particle whose decay time is uncertain. Since photon energy scales with photon frequency, this effect broadens an absorption line, modeled in frequency space as a Dirac delta function, into a Lorentzian function [72]. The modeling Lorentzian function is shown in equations (3-13) and (3-14).

$$\mathrm{nat}(\upsilon) = \frac{\Delta\upsilon_n}{2\pi} \frac{1}{(\upsilon - \upsilon_0)^2 + (\Delta\upsilon_n/2)^2} \tag{3-13}$$

$$\Delta\upsilon_n = \frac{1}{2\pi\tau_n} \tag{3-14}$$

where $\Delta\upsilon_n$ is the characteristic width of the Lorentzian function in frequency domain, υ_0 is the frequency of the transition, and τ_n is the uncertainty in the decay time. This uncertainty turns out to be the same as the lifetime of the upper state of the transition and the derivation is shown in [72]. Table 3-4 shows the lifetime of the upper state for the Xe I 823.4 nm, Xe I 834.9 nm, and Xe II 835.0 nm transitions. The values for the neutral upper states are from [73] and the values for the singly-charged upper state are from [74]. In practice, the amount of broadening due to the lifetime effect is one to two orders of magnitude smaller than the amount of broadening due to the Doppler effect, even in room-temperature plasma, and can usually be neglected.

Table 3-4. Upper state lifetime of several xenon transitions.

Transition	Lifetime of the upper state, ns
Xe I 823.4 nm	48
Xe I 834.9 nm	33
Xe II 835.0 nm	9.3

3.2.4 Saturation Broadening

Physically, the saturation effect occurs when the laser beam intensity is high enough such that the size of the upper state population is comparable to the size of the lower state population. At low input beam intensity, the intensity of the emitted fluorescence is proportional to the beam intensity. However, as the input beam intensity rises, the rate of increase in the emission intensity begins to drop until the emission intensity asymptotically reaches a maximum. At a high enough input beam intensity, the recorded emission signal becomes independent of the input beam intensity. For a two-level system, saturation can be modeled by equation (3-15) [72].

$$I_R \alpha \frac{S(\upsilon)}{1+S(\upsilon)}, \text{ where } S(\upsilon) \equiv \frac{I_\upsilon}{I_S(\upsilon)} \tag{3-15}$$

where I_R is the response (LIF or optogalvanic) of the sampled population, S is called the saturation parameter, I_υ is the injected beam intensity, and I_S is the saturation intensity. It can be shown that I_S is the intensity of the injected beam necessary to equally populate the upper and lower states of a transition. Note that $I_S(\upsilon)$ is frequency dependent and is inversely proportional to the magnitude of the lineshape $i(\upsilon)$. For this reason, $S(\upsilon)$ is proportional to $i(\upsilon)$ and is higher at the peaks and lower at the troughs in the lineshape. It follows that the peaks in the LIF lineshape will experience a greater reduction in response

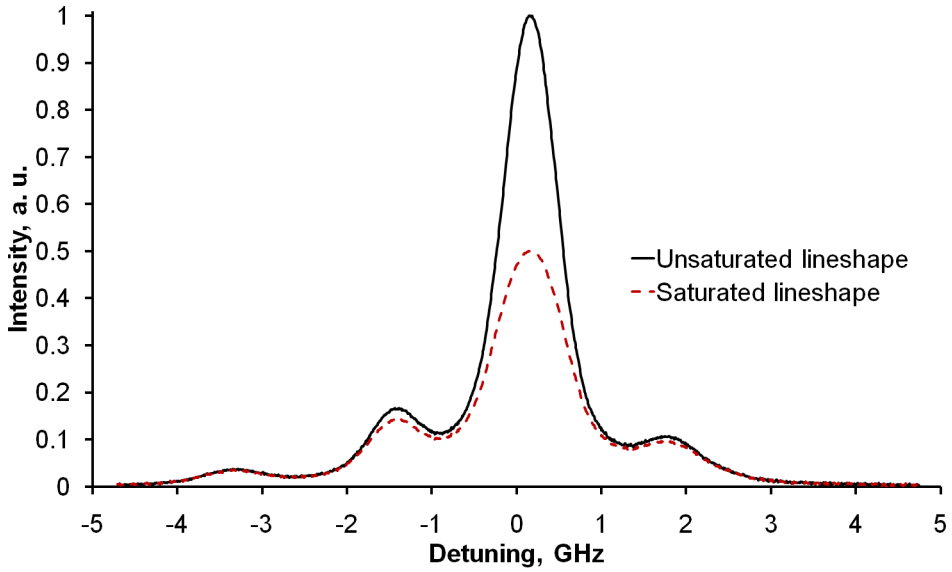

Figure 3-4. Comparison of saturated versus unsaturated spectral lineshape.

intensity than the troughs. This effect causes the peaks to appear shorter than they would be for an unsaturated lineshape. As a result, the lineshape will appear to have been broadened. Figure 3-4 shows a simulated example of how the saturation effect shortens the main peak of the Xe I 834.9 nm lineshape while making that peak look broader.

In general, it is not possible to remove the distortion brought on by saturation from raw LIF traces in post processing. This is because modeling the effects of saturation requires knowledge of the laser intensity profile at the interrogation zone. It is possible to make an approximation or to measure this profile, but the preferable approach is to simply avoid saturation when trying to extract VDFs. This is accomplished by performing saturation studies to map the response of the sample as a function of input laser intensity. Then the input laser intensity is adjusted, e.g. using a neutral density filter, such that the response is linear.

3.2.5 Influence of External Electric Field

It is also possible for the spectral lineshape to undergo further splitting under the influence of an external electric field. A strong enough external electric field causes the centroid of the electron cloud to no longer coincide with the center of the nucleus, inducing a dipole as a result. The total resultant angular momentum vector formed by the nucleus and the electron can couple with a dipole vector in a quantized manner and give rise to the splitting of otherwise degenerate energy states. This is called the Stark effect.

The earliest literature with regards to the Stark effect for the specific LIF schemes used for this dissertation dates back to 1933 in an experiment by Harkness and Heard [75]. The applied field strength in this experiment was 67,000 V/cm or $6.7*10^6$ V/m. Harkness and Heard found no measurable splitting for the Xe I 834.9 nm transition with a converted uncertainty of ~3 GHz. The highest electric field in the Hall thruster is around 30 kV/m (based on LIF measurements to be described in Chapter 5). Assuming a roughly linear Stark effect, the maximum Stark shift possible within the uncertainty bounds is ~14 MHz. Thus, the Stark effect for the Xe I 834.9 nm transition is negligible when compared to other broadening factors. Harkness and Heard also commented that they found no measurable shift in the singly-charged xenon line, but did not specifically state which line they studied. Since no subsequent literature has surfaced that suggests the Stark effect for the xenon transitions used in this dissertation is important in a Hall thruster environment, the Stark effect is assumed to be negligible.

Associated with the Stark effect is a phenomenon called Stark broadening. In essence, it is the Stark effect observed when particles collide with ions and electrons. Due to variation in distance from the charged particle during collision, there is a spread in Stark splitting and a broadening of the lineshape when collision frequency is high. To estimate the amount of broadening, data from work by Manola, et al., will be used [76]. For the Xe II 542.1 nm de-excitation transition (from the Xe II 835.0 nm scheme), at an electron density of 10^{23} m^{-3}, the FWHM of Stark broadening is ~0.2 nm. Assuming linear scaling, for Hall thruster plasma where electron density reaches only ~$1*10^{19}$ m^{-3}, the Stark broadening is ~20 MHz, which is negligible compared to other broadening factors.

3.2.6 Influence of External Magnetic Field

Much as an external electric field can cause otherwise degenerate energy states to split, an external magnetic field can also do the same. It turns out that the strength of the magnetic field typically found in the Hall thruster (a few hundred Gauss) is high enough to cause very noticeable distortions to the neutral xenon lineshapes. This effect is less noticeable for the singly-charged ion lines because the Doppler broadening tends to be around over an order of magnitude higher for the singly-charged ion lines than the neutral lines.

For a weak magnetic field (less than a few hundred Gauss), this splitting is called the Zeeman effect. For a strong magnetic field (more than a few thousand Gauss), the effect is called the Paschen-Back effect. There is a great difference in the way the degenerate energy states behave in the two regimes. For typical field strengths found in Hall thrusters, a linear Zeeman effect model is sufficient if one only wants to approximate the lineshape. On the other hand, a nonlinear Zeeman model is necessary if one wants to use the lineshape to measure the in-situ magnetic field. The reason for this is that the lineshape is not very sensitive to small changes in the magnetic field strength. A high fidelity model is needed if one wishes to accurately calculate the magnetic field strength from Zeeman splitting. For in-depth descriptions of both the linear and nonlinear Zeeman models, refer to Bailo Ngom's dissertation [77].

The Zeeman effect can also be categorized as σ- or π-polarized, based on the orientation of the laser beam polarization with respect to the magnetic field direction. The atom is σ-polarized if the polarization of the laser beam is perpendicular to the external magnetic field ($\mathbf{E} \perp \mathbf{B}$). The atom is π-polarized if the polarization of the beam is parallel to the external magnetic field ($\mathbf{E} \parallel \mathbf{B}$). As a part of this dissertation, experiments performed for the Xe I 834.9 nm and Xe II 835.0 nm transitions showed the π-polarized splittings are at least an order of magnitude narrower than σ-polarized splittings. These experiments are described in Chapter 4.

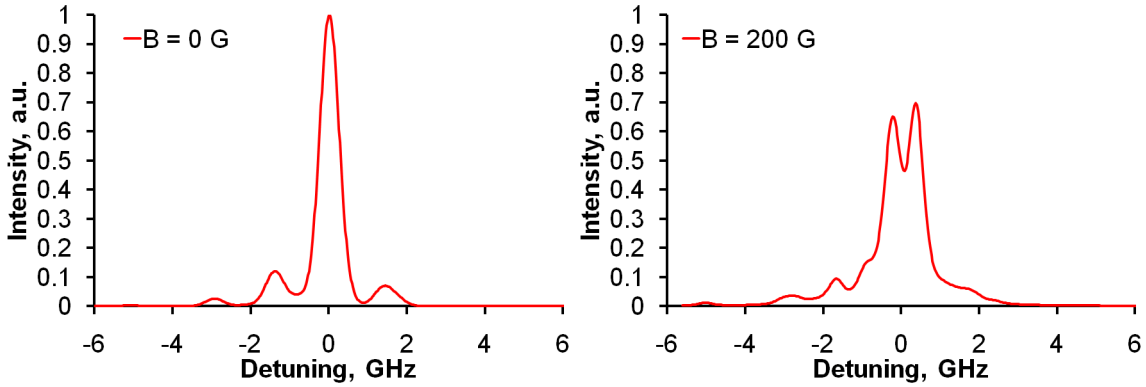

Figure 3-5. An example of how external magnetic field affects the Xe I 834.9 nm lineshape.

Figure 3-5 shows an example of how the Xe I 834.9 nm lineshape can become distorted when under the influence of an externally applied magnetic field. As the figure shows, the distortion can be quite severe and must either be removed by taking proper care when setting up the LIF diagnostics or by using some form of deconvolution in post processing if one is interested in extracting VDFs. However, the Zeeman effect produces no net displacement of the intensity average of the spectra. If one was only interested in the mean velocity of the atoms, the Zeeman effect can be ignored.

Having described the major broadening factors that are important to the application of xenon LIF velocimetry in a Hall thruster environment, it should be clear that a number of preparatory experiments must be performed in order to interpret xenon LIF velocimetry data. In particular, there are two objectives that must be accomplished. The first objective is to develop a hyperfine structure model for the Xe II 835.0 nm transition, for which no established hyperfine structure constants are available. The second objective is to develop Zeeman effect models for the Xe I 834.9 nm and Xe II 835.0 nm transitions should it prove physically impossible to acquire data that are free from this effect. These experiments are described in Chapter 4.

Chapter 4
Intermodulated Optogalvanic and Zeeman Effect Studies

4.1 Motivation for the IMOG and Zeeman Experiments

As previous mentioned in Chapter 3, extraction of the velocity distribution functions from raw LIF traces requires proper accounting of all applicable broadening effects. In particular, the hyperfine structure of the Xe II 835.0 nm transition and the Zeeman effect of the Xe I 834.9 nm and Xe II 835.0 nm transitions need to be characterized before those two transitions can be used to extract VDFs.

To accomplish the first objective of hyperfine structure characterization, a technique known as intermodulated optogalvanic (IMOG) spectroscopy is used. This technique is designed to take advantage of the nonlinearity associated with the saturation effect to extract lineshapes that are free of Doppler broadening. A lineshape freed from the influence of external electric and magnetic field is broadened primarily by the Doppler effect, the hyperfine structure, lifetime broadening, and saturation broadening. By using saturation broadening in a clever fashion and removing the Doppler effect, only the hyperfine structure and lifetime broadening remains. This resulting lineshape is called a kernel in this dissertation. Since lifetime broadening is much narrower than the hyperfine structure, the hyperfine structure can usually be extracted.

To carry out IMOG spectroscopy, two laser beams of approximately equal intensity are sent from the opposite directions to converge in the interrogation zone. Normally, one beam interacts with the population travelling one way while the other beam interacts with the population travelling in the opposite direction. When the photon frequency of the beams matches the zero-velocity frequency, both beams interact with the same population. If the interaction is linear, no change occurs in the response of the signal. However, if the beams are intense enough to produce a strong saturation effect, the total response at the zero-velocity frequency will be different than the response at

other frequencies. By modulating (mechanically chopping) the input beams at frequencies f_1 and f_2, respectively, a Doppler-free response is obtained at frequencies f_1+f_2 and f_1-f_2. Refer to Svanberg's text for an overview of intermodulated techniques [72]. Note that the modulation frequency refers to the rate at which light is physically blocked and unblocked, and is not to be confused with the wave frequency of the light.

In practice, some Doppler broadening shows up in IMOG spectra due to velocity-changing collisions that re-direct particles to go in the exact opposite direction. In addition, the two beams cannot be perfectly aligned. While the components of signal corresponding to these broadening effects are low in intensity, so is the Doppler-free component of the signal. Defining these broadening effects as noise, it can be shown with the saturation model, equation (3-15), that the SNR for an IMOG spectrum asymptotically reaches a maximum as the degree of saturation is increased. Furthermore, for a given total power, the signal theoretically maximizes when that power is equally split between the two beams of equal cross-sectional area.

Realistically, the degree of saturation is dependent on the power of the laser and the ability to focus two beams onto the same interrogation volume. Furthermore, an IMOG spectrum is not immune to the same undesirable broadening that can be found in LIF and regular optogalvanic traces. This is part of the reason why normally IMOG is only used to extract the hyperfine constants, which are then used to construct the kernel. For the purpose of the completed experiment, saturation broadening puts an additional lower limit on the observed kernel line width. The thinner the kernel line width is, the closer it is to being Doppler-free.

It should be noted that although IMOG specifically uses an optogalvanic cell as the response collection device, the associated principle can be applied to another collection method, such as photomultiplier collection of fluorescence. For this dissertation, such an approach is termed intermodulated LIF (IM-LIF). However, it will be shown that IM-LIF results have a lower SNR than IMOG results, so ultimately the IMOG results were used. A description of the IMOG and IM-LIF experiments can be found in Sections 4.2.1 and 4.2.2.

The second objective, which is to study the Zeeman effect, is accomplished by a set of carefully designed experiments where the magnetic field and the laser beam

polarization directions are controlled. By taking optogalvanic traces of the spectral lineshape while the interrogated atoms experience varying magnetic field strength, it is possible to construct and validate models for the Zeeman effect. During these experiments, we also discovered that the lineshape is insensitive to an externally applied magnetic field when the atoms are π-polarized. These experiments are described in Sections 4.2.3 and 4.2.4.

4.2 Experimental Setup for the IMOG and Zeeman Experiments

This section begins by describing two variations of the intermodulated optogalvanic experimental setups used to extract the hyperfine structure of the Xe II 835.0 nm transition. It then describes the experimental setup for the Zeeman effect experiment on the Xe I 834.9 nm transition. The section ends with a description of the experimental setup for the Zeeman effect experiment on the Xe II 835.0 transition.

4.2.1 IMOG Spectroscopy Experimental Setup I

Figure 4-1 shows the laser and optics setup for the first version of the IMOG spectroscopy experiments. Note that diagrams in this dissertation are generally not drawn to scale. A small rectangular slab in the figure denotes either a mirror or a beamsplitter. A

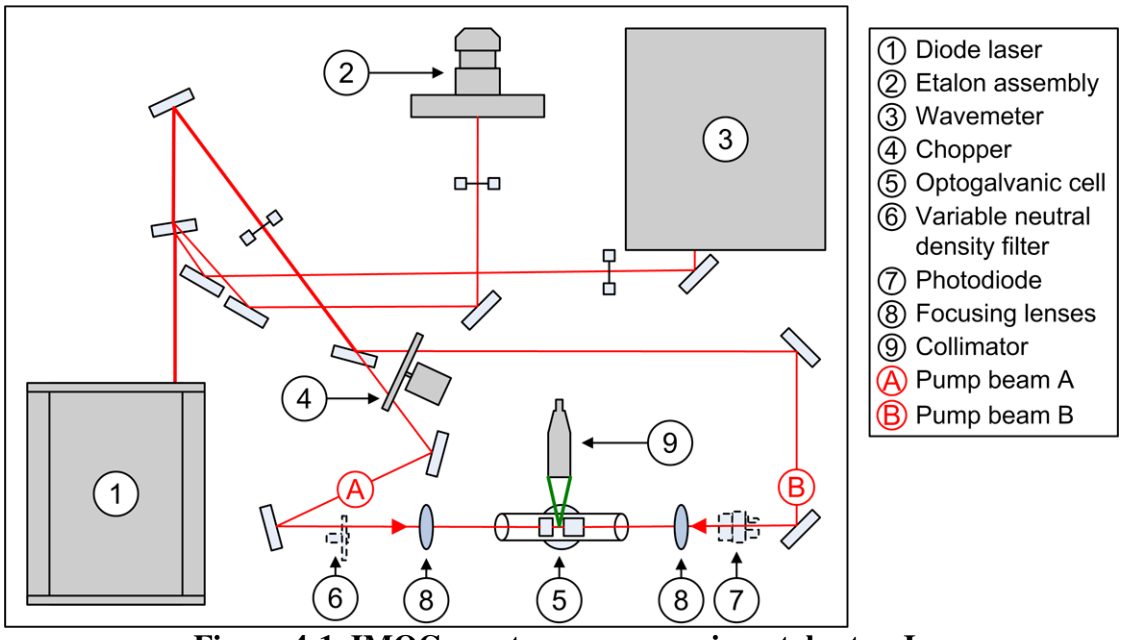

Figure 4-1. IMOG spectroscopy experimental setup I.

line connecting two small squares denotes an iris. All other components are labeled.

The laser system used is a TOPTICA TA-100/830 tapered-amplifier diode laser. This system has a nominal linewidth of ~10 MHz and a mode-hop-free range of ~12 GHz. The total output power is ~250 mW. The setup uses a Burleigh SA-91 etalon assembly (2-GHz free spectral range, finesse >250) and a WA-1000 wavemeter (±1 pm) for reference information. The main laser beam is split roughly equally into pump beam A and pump beam B.

Pump beam A, chopped at ~1160 Hz, was typically ~10% weaker than pump beam B due to limitations in the optics. During the saturation study, a variable density filter and a Thorlabs DET-110 photodiode (shown in dashed lines) are added while pump beam B is blocked. Both pump beams have a physical cross-sectional diameter of ~1 mm.

Pump beam B is chopped at ~970 Hz using a different set of slots on the same chopper that chops pump beam A. This beam is aligned to overlap with pump beam A in the middle of the optogalvanic cell with the beam axes being within 0.1° of each other.
A pair of plano-concave focusing lenses designed to focus the beams into a much smaller volume can be added or removed. A collimator focused on the interrogation zone is added to allow fluorescence signal to be collected during the intermodulated tests for IM-LIF. An optical fiber brings the fluorescence signal to a SPEX-500M monochromator, the output of which is amplified by a Hamamatsu R928 photomultiplier. The amplified signal is sent to an SR810 lock-in amplifier.

The optogalvanic cell used in this study is a Hamamatsu L2783-42 XeNe-Mo galvatron. The core of the cell is a pair of φ6.25 mm cylindrical Mo tube electrodes centered in a φ25 mm x 120 mm glass cylinder. The cylinder is filled with approximately 3 Torr of xenon and 4 Torr of neon. The ends of the cylinder are angled approximately 10 degrees from being perpendicular to the electrode axis to eliminate retro-reflection. This design permits unobstructed passage of a laser beam along the axis shared by the cylinder and electrodes.

The optogalvanic cell is operated at 250 V to maintain a dense, stationary plasma rich with ion species. A ballast resistor prevents runaway current growth after the discharge is struck, and provides a simple method of measuring the discharge current. The voltage drop across the ballast resistor is connected to another SR-810 lock-in

amplifier through an RC filter that passes only the AC component of the signal. Data are collected at both the sum and the difference of the two chopping frequencies.

IMOG setup I was designed to investigate various parameters that can be varied to improve the SNR associated with the Doppler-free traces. Two methods of data collection were investigated including the use of the optogalvanic cell and the collection of emitted fluorescence. Different degrees of saturation are also investigated by controlling the use of focusing lenses.

4.2.2 IMOG Spectroscopy Experimental Setup II

Data gained from IMOG setup I (Section 4.3.1) showed that the limitations imposed by the various noise and saturation effects may give rise to a particular saturation level that yields the best IMOG results. Setup II, shown in Figure 4-2, was designed to allow better control of the saturation level. The variable neutral density filter simultaneously changes the power of both pump beams by the same percentage. This setup has no fluorescence collection capability because setup I showed that optogalvanic signal collection yielded a much higher SNR.

A few mirrors were removed going from setup I to setup II to increase the maximum power delivered. Better optics management also brought the power of the two

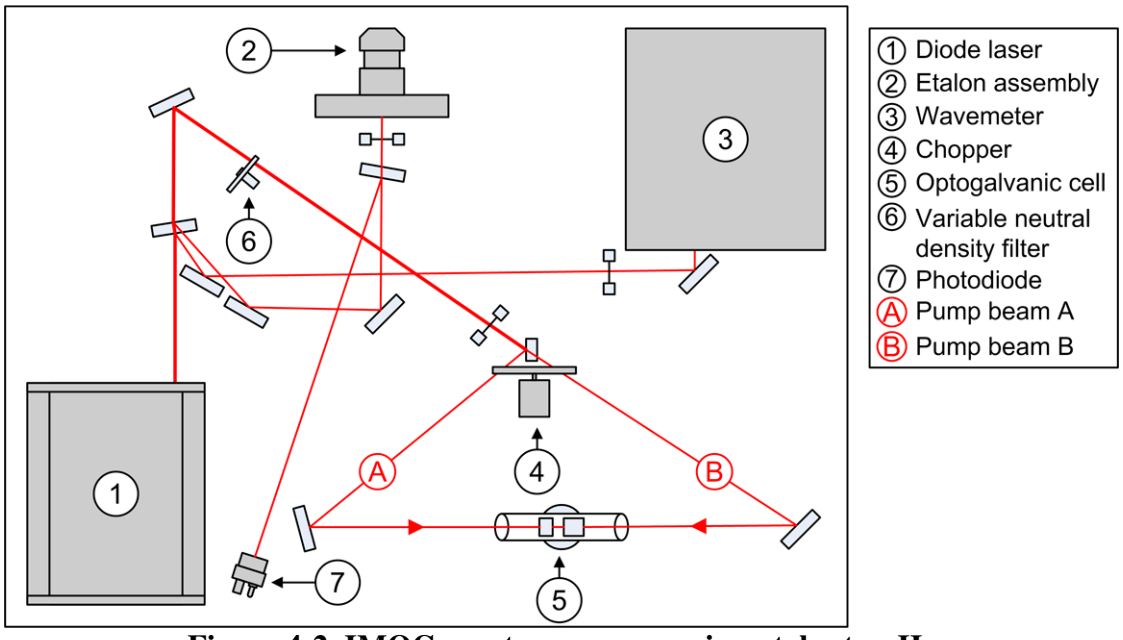

Figure 4-2. IMOG spectroscopy experimental setup II.

pump beams to within ~6% of each other. Chopping frequencies were slightly lower, at ~1140 and ~950 Hz for pump beams A and B, respectively. A saturation study was performed in much the same manner as with setup I, though the associated elements are not shown. A Thorlabs PDA-36A photodiode, different from the one that is used for the saturation study, was added to monitor laser power drift.

4.2.3 Zeeman Effect Study for Neutral Xenon

Figure 4-3 shows the laser and optics setup for the first version of the IMOG spectroscopy experiments. As before, the laser system used is a TUI Optics TA-100/830 tapered-amplifier diode laser operating at ~250 mW. The setup uses a Burleigh SA-91 etalon and a WA-1000 wavemeter for reference information. Unlike the previous IMOG experimental setup, our goal is to obtain a good SNR while avoiding saturation. The majority of the main beam is ejected to a beam dump while a weak beam split off from the main beam is used to probe the optogalvanic cell.

The probe beam goes through, one iris, two mirrors, one half-wave plate, one SR540 chopper (at ~1.6kHz), one polarizing beamsplitter cube, the optogalvanic cell (surrounded by the solenoid electromagnet), and into a Thorlabs DET-110 photodiode. The photodiode is linked to a SR810 lock-in amplifier. By controlling the orientation of the polarizing cube, the electric field direction of the inbound photons can be set to be

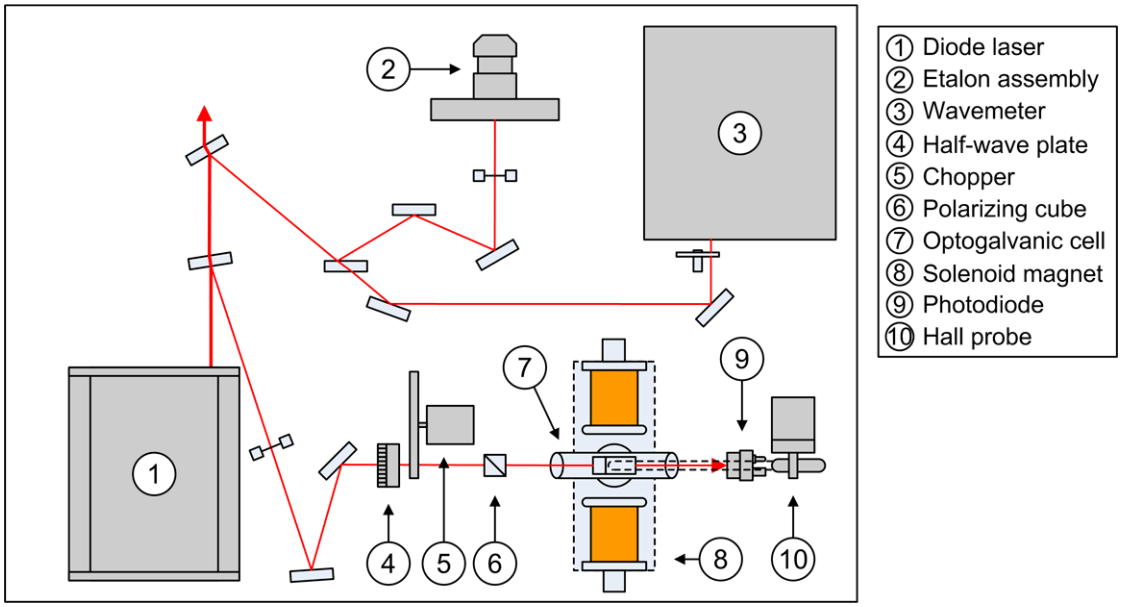

Figure 4-3. Experimental setup for the neutral xenon Zeeman effect study.

parallel or perpendicular to the magnetic field direction at the center of the optogalvanic cell. The addition of the half-wave plate turns the polarization of the laser beam to control the laser power passing through the polarizing cube. During the saturation study, an additional neutral density filter post is added between the cube and the optogalvanic cell for better control of the laser power. The DET-110 photodiode, combined with an internal photodiode that reports the power of the master oscillator, helps monitor the laser power over the course of a given scan. This power level is then used to normalize the optogalvanic cell output to remove any influence that variation in the laser power may have on the lineshape. This technique only works if the lineshape is unsaturated. The optogalvanic cell is operated at 250 V. The solenoid magnet is designed to generate a uniform magnetic field that points from one coil to the other in a region roughly 1 cm in size right at the center of the two coils. It is powered by a Kikusui PAD 16-10L power supply. The magnet is cooled by external air flow and can generate up to a 300-Gauss field strength at the center without becoming too hot to touch. A MMZ-2508-UH 3-axis Hall probe connected to a LakeShore 460 gaussmeter is used to monitor the magnetic field strength with a measurement uncertainty of ±0.25%.

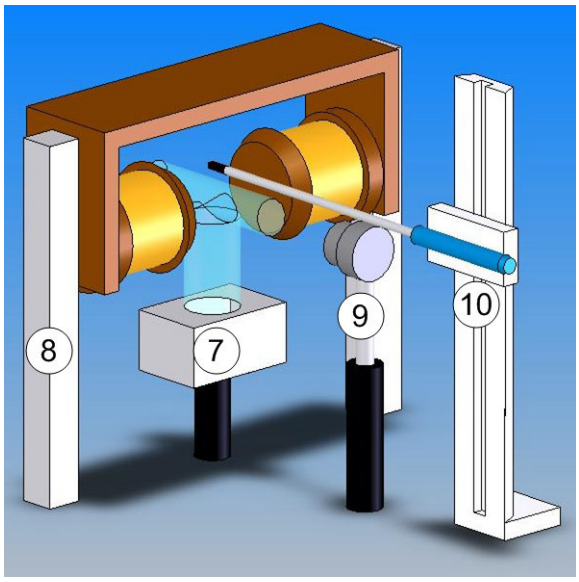

Figure 4-4. 3D picture of the experimental setup for the neutral xenon Zeeman effect study.

Figure 4-4 shows a diagram of part of the experimental setup. The figure shows how the Hall probe tip sits 2.1 cm above the center of the optogalvanic cell. Prior to testing, the optogalvanic cell and the photodiode are removed so that the Hall probe can be moved vertically on its motion stage. The magnetic field direction is checked to make sure it runs parallel to the table surface. The magnetic field strength at the center of where the optogalvanic cell will reside and where the Hall probe will sit during live experimentation are recorded. The field strength values at these two locations are compared across a wide range of solenoid currents in order to form a calibration curve.

Since the optogalvanic cell is made of non-magnetic materials, we can monitor the magnetic field strength at the center of the optogalvanic cell using this calibration curve and the off-center Hall probe measurement.

Figure 4-5 shows the excellent linearity of the calibration curve.

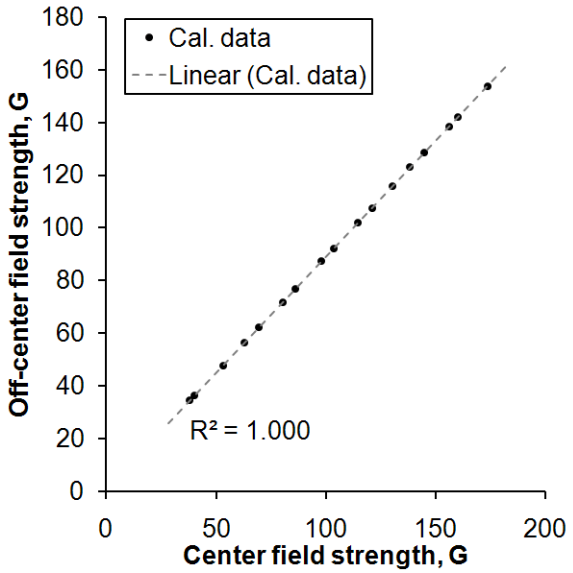

Figure 4-5. Calibration curve for using an off-center Hall probe to monitor the magnetic field at the center of the optogalvanic cell.

4.2.4 Zeeman Effect Study for Singly-Charged Xenon

The Zeeman effect study for the singly-charged xenon used a setup that is very similar to that used for the neutral xenon. The singly-charged xenon study took place a few years after the neutral xenon study, so the major difference is in the availability of newer and better equipment. In particular, this study used a TOPTICA TA-pro 850 tapered-amplified diode laser with a linewidth of <1 MHz and a 50+ GHz mode-hop-free range, operating at ~200 mW of power. This study also used a High Finesse

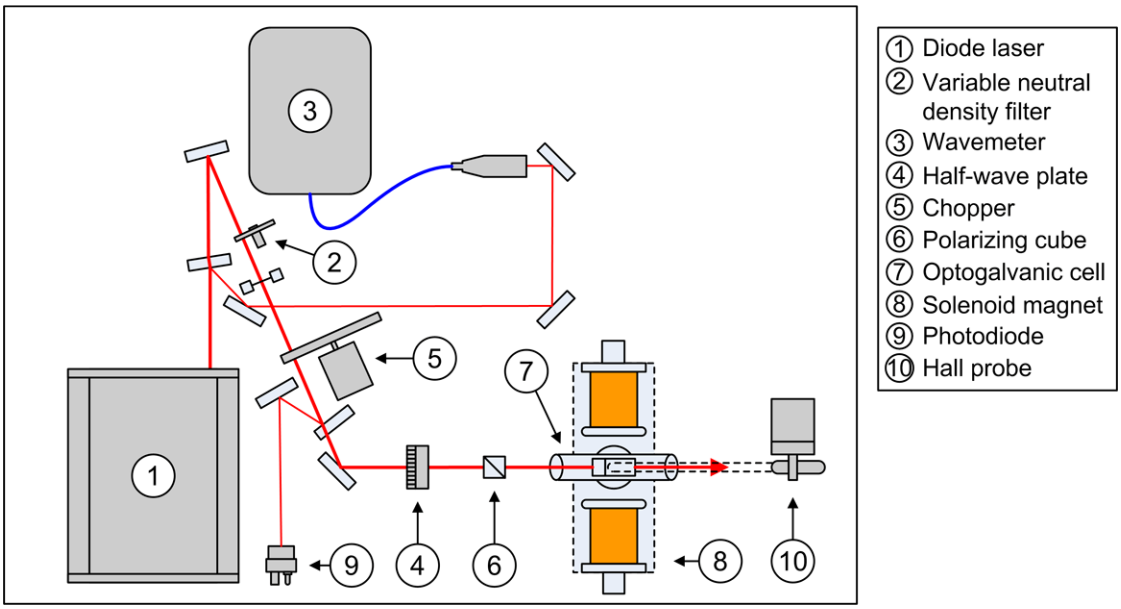

Figure 4-6. Experimental setup for the singly-charged xenon Zeeman effect study.

WS7 Fizeau-type wavemeter with a resolution of 10 MHz (0.023 pm at 835 nm) and absolute accuracy of 60 MHz (0.14 pm at 835 nm). The excellent resolution of the WS7 made the use of the SA-91 etalon unnecessary. The same solenoid magnet as in the neutral xenon Zeeman effect study was used, but the power supply was replaced by a Kepco ATE 36-30M instead. For this experiment, the SR540 chopper operated at ~1.8 kHz, and the photodiode was a Thorlabs PDA-36A. Figure 4-6 shows a diagram of the experimental setup for the singly-charged xenon Zeeman effect study. Note that the optogalvanic cell, operating at 250 V, contains a much lower density of singly-charged xenon than of neutral xenon. As such, more laser power is needed to get an adequate signal for this study than during the neutral xenon study. The variable neutral density filter is used to control the exact amount of laser power being fed into the optogalvanic cell. In order to get a statistically significant signal, the laser power had to be set to a point that slightly saturated the transition. The degree of saturation can be seen in the corresponding saturation study plot in Section 4.3.1.

4.3 Data Reduction for the IMOG and Zeeman Experiments

All uncertainties reported in this dissertation have associated confidence intervals of 95% (or roughly two standard deviations). Signal-to-noise ratio is defined as peak signal intensity divided by root-mean-square of noise.

4.3.1 IMOG Spectroscopy Preparatory Studies

Since IMOG spectroscopy typically has a low SNR, the first step of the experiment was to try different signal collection methods. Setup I allowed a comparison of intermodulated optogalvanic (IMOG) against intermodulated-laser-induced fluorescence (IM-LIF) spectroscopy. Also, the focusing lenses made it possible to investigate the effect of changing saturation level without changing the total pump beam power.

Figures 4-7 and 4-8 show the saturation study plots for the optogalvanic and the fluorescence collection schemes, respectively. The uncertainty on the peak intensity is ±3% and uncertainty on the normalized laser power is ±1%. The associated error bars are not drawn since they are around the same size as the data dots. The uncertainty on the

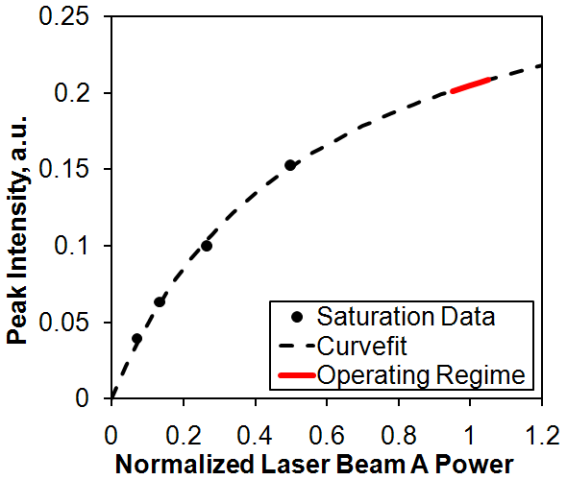

Figure 4-7. Saturation study plot for IMOG spectroscopy setup I.

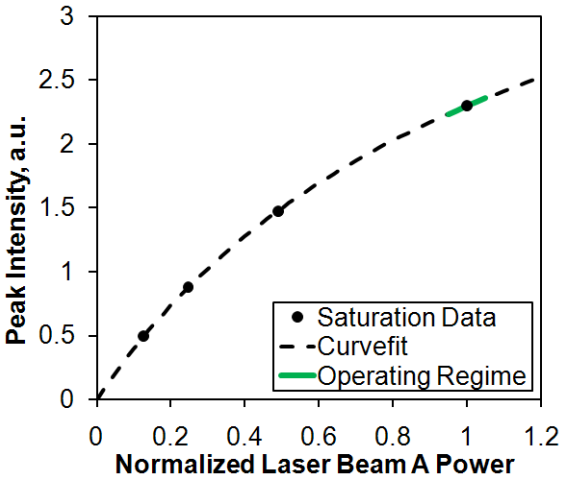

Figure 4-8. Saturation study plot for IM-LIF spectroscopy setup I

normalized laser power comes from photodiode readings. The actual laser intensity is ~60 mW/mm^2 when the normalized laser power is 1 for the setup with no focusing lenses. This value is calculated from thermopile and beam waist measurements. Each figure contains dashed lines representing fits to the data based on equation (3-15), the saturation model equation. Each figure is also overlaid with solid colored lines indicating the approximate laser output at which the IMOG and IM-LIF traces are taken. From these two figures, we can see the interrogated gas sample is well saturated by the input laser beams.

Figure 4-9 compares representative IMOG and IM-LIF traces on the same plot. Aside from the collection method, the two raw traces were taken under identical conditions. The IMOG trace was taken immediately after the IM-LIF trace. Each trace took about 8 minutes to complete. The results shown in Figure 4-9 indicate that the optogalvanic signal collection scheme produced a far higher SNR than the fluorescence signal collection scheme. The optogalvanic signal was more sensitive to saturation than the fluorescence signal. Furthermore, similar

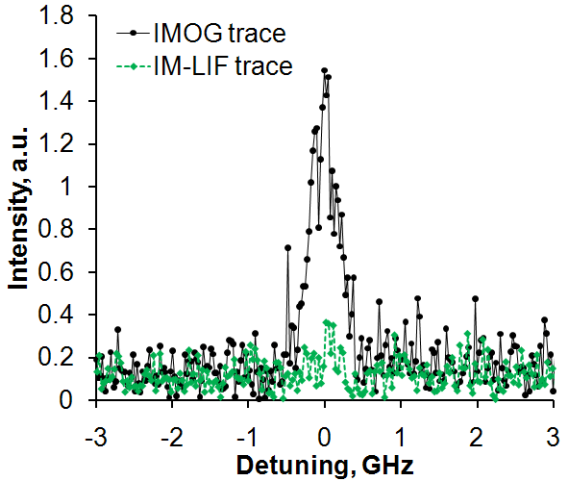

Figure 4-9. Comparison of sample data taken with IMOG versus IM-LIF.

comparisons were made with the focusing lenses in place, and the signal became undetectable for IM-LIF. Thus, the optogalvanic cell was chosen for data collection.

Studies were also done comparing data taken at the sum of the two beams' chopping frequencies to data taken at the difference of the frequencies; data taken at the sum of frequencies consistently show roughly an order of magnitude better SNR. This may have been due to the fact that the difference of the frequencies was ~190 Hz, where environmental electrical interference could have been an issue. All data shown here are collected at the sum of the chopping frequencies.

The last preparatory study compares the effect of beam focusing on the degree of saturation and the SNR. The focusing lenses are plano-convex lenses with a focal length of 100 mm. Given the laser beam has a nominal beam waist of 1 mm, the calculated focused beam waist was ~100 μm at the center of the optogalvanic cell. Figure 4-10 is the saturation study plot associated with having the focusing lenses in place.

Figure 4-11 compares an IMOG trace taken with the focusing lenses in place with an IMOG trace taken without the lenses. This figure shows that the SNR decreased greatly when the focusing lenses were in place. The most likely explanation for this decrease is that placing the focusing lenses also changed the size of the interrogation zone, greatly decreasing the number of particles being interrogated. Additionally, it was difficult to align two narrow beams with respect to each other. Any misalignment could have greatly decreased the effective interrogation volume and the signal strength. The

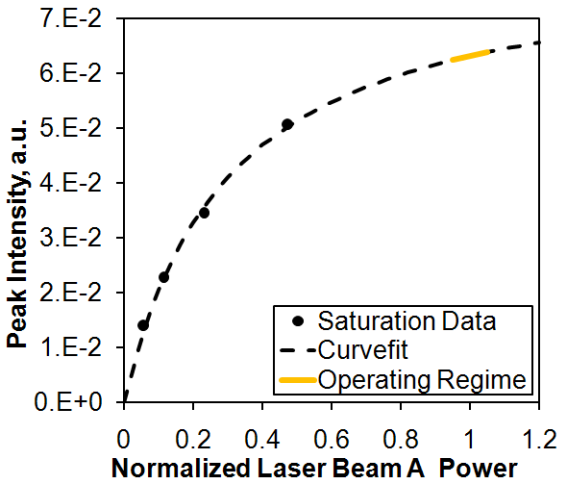

Figure 4-10. Saturation study plot for IMOG with focusing lenses in place.

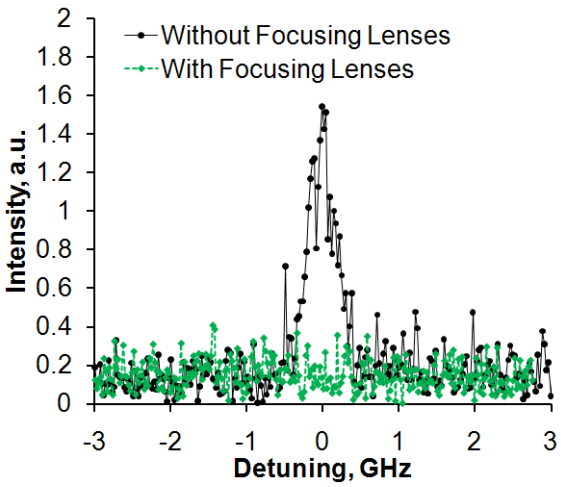

Figure 4-11. Raw IMOG traces taken with and without focusing lenses in place.

results of all these preparatory tests inspired IMOG experimental setup II, where the saturation level is easily controlled by a variable neutral density filter. The size of the interrogation region for setup II remains the same regardless of the saturation level, though the saturation level is now tied to the laser power.

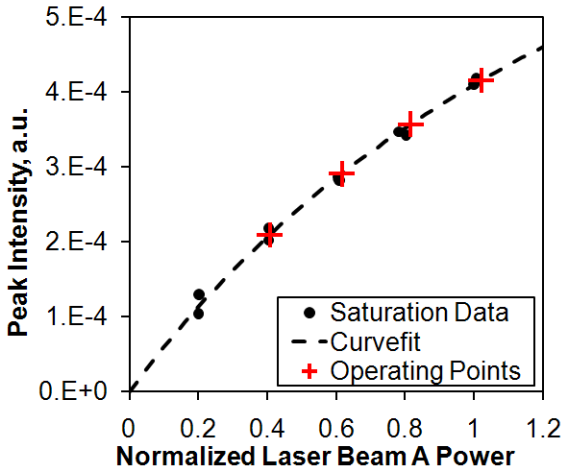

Figure 4-12. Saturation study plot for the IMOG data taken with setup II.

4.3.2 IMOG Spectroscopy Results and Uncertainty

Figure 4-12 shows the saturation study plot for data taken with setup II. Unlike setup I, IMOG data from setup II are taken at different saturation levels along the fitted curve. The laser power levels at which traces were taken are marked as crosses, which roughly correspond to 40, 60, 80, and 100 percent of the available laser power.

Four IMOG traces were taken at each power level. Each set of four traces were then averaged to form one trace for the associated power level. Each trace used a lock-in time constant of 1 second and data was collected at the sum of the two chopper

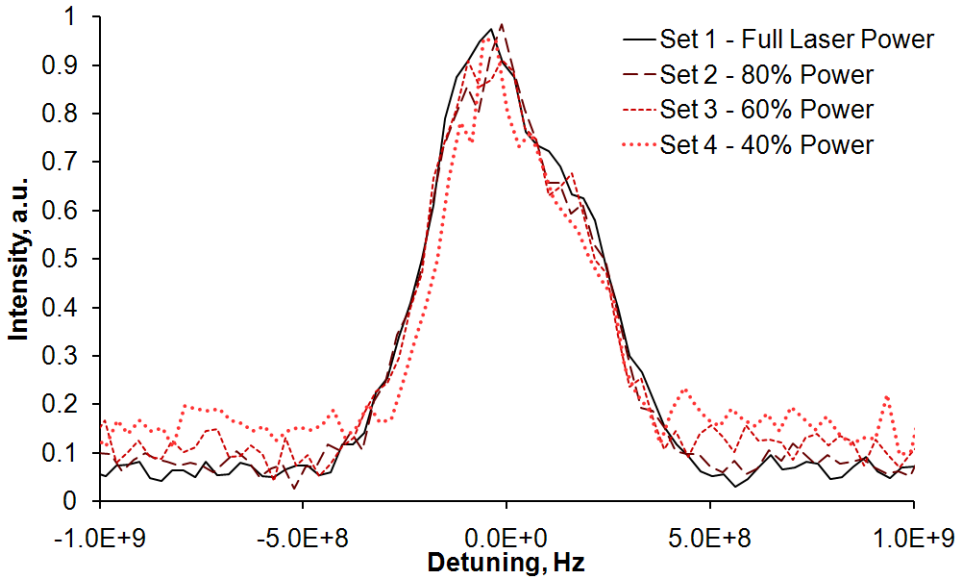

Figure 4-13. IMOG results by input laser power level.

Full laser power corresponds to a laser intensity of ~60 mW/mm^2.

frequencies. Data was also taken at 20 percent of full power but the SNR was too low for those traces to be useful. The traces were averaged by first finding the mean frequency of the trace in order to line them all up with respect to each other. Then, the traces were merged and every four data points averaged. Figure 4-13 shows the averaged results of the four data sets plotted together. From this figure, we see that while the full-width-at-half-maximum (FWHM) frequency of the Doppler-free scans goes down as the sample becomes less saturated, the SNR also worsens. Here SNR is defined as the peak signal divided by the root-mean-square of noise, which is calculated using the parts of the traces that do not appear to contain signal. Table 4-1 summarizes the FWHM frequency and SNR value for each of the four data sets. The 40-percent-laser-power trace (dotted line in Figure 4-13) is ~15% thinner than the full-power trace but the scan is so much noisier that it is difficult to determine whether the features are noise or real hyperfine structures.

Table 4-1. Summary of IMOG results as a function of input laser power.

Data Set	Percentage of full power	FWHM frequency, MHz	SNR
1	100	449	15.0
2	80	447	11.6
3	60	421	7.5
4	40	373	5.8

As a compromise between increasing SNR and decreasing saturation broadening, a set of 7 scans were taken with pump beam A laser power at ~55% of full power. The

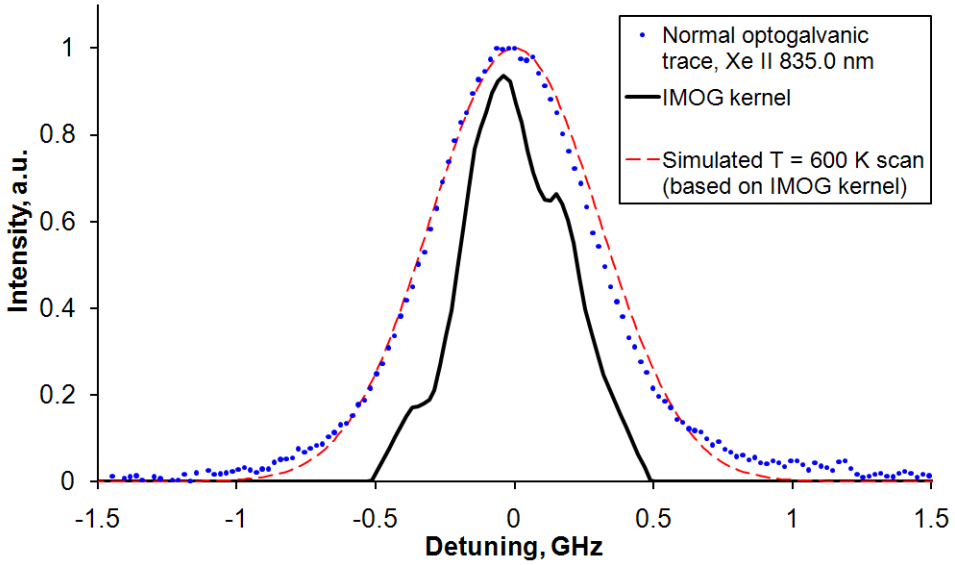

Figure 4-14. IMOG spectroscopy results and validation.

averaged result was then smoothed with a seven-point Savitzky-Golay filter [78]. Lastly, the noisy floor away from the peak was removed by linearly extrapolating for where the two slopes of the peak meet the horizontal zero axis. The final result is shown in solid black line in Figure 4-14 and will subsequently be referred to as the IMOG kernel. Also shown in this figure in dots is a normal (unsaturated) optogalvanic trace taken under the same optogalvanic cell operating condition as was used to obtain the IMOG spectra. From previous measurements of xenon lines with known hyperfine structure constants, it is determined that the gas temperature is 600±50 K at this operating condition. A dashed curve showing a simulated 600 K predicted spectrum is shown in Figure 4-14; this spectrum is obtained by convolving the IMOG kernel with a 600 K Maxwellian velocity distribution. The normal optogalvanic trace and the simulated 600 K trace based on the obtained kernel show excellent agreement. The simulated spectrum is slightly broader than the normal optogalvanic trace most likely due to some Doppler broadening that still remains in the IMOG kernel. Even a small amount of Doppler broadening can smooth out sharp hyperfine features that are packed closely together. This is why the IMOG kernel does not look like a series of spikes as would be expected from normal hyperfine structure. The IMOG kernel has a FWHM frequency of ~450 MHz; most of the strong hyperfine lines should lie within that confine. This kernel does not capture the wings of the hyperfine structure lineshape perfectly since the wings are composed of weak, relatively unsaturated hyperfine lines. Nevertheless, it is usable for deconvolution purposes.

A detailed comparison between using and not using the IMOG kernel to extract VDFs from raw LIF data was carried using data obtained along the channel centerline in the acceleration zone of the 6-kW Hall thruster. For a more detailed description of this analysis, please refer to [65]. The basic results of the analysis showed that there is no statistically significant difference in the mean velocity between using and not using the IMOG kernel. However, the FWHM velocity of the final VDFs can be up to 6% lower when the IMOG kernel is accounted for in the region upstream of the acceleration zone, where the ions are freshly born and have not begun accelerating. In this region inside the Hall thruster, the Doppler broadening is observed to typically be ~1 GHz wide (i.e. FWHM frequency is ~1 GHz). Downstream of this region, the Doppler broadening is

typically a few GHz wide and deconvolving the IMOG kernel out of the raw LIF traces only reduces the FWHM velocity by 1-2%. Thus, if 1-2% error is considered acceptable, the hyperfine structure of the 835.0 nm transition can be ignored when analyzing data from the acceleration zone.

4.3.3 Neutral Xenon Zeeman Effect Study Results and Validation

The Zeeman effect study was performed for the Xe I 834.9 nm and the Xe II 835.0 nm transitions up to 300 Gauss for both π- and σ-polarized laser input. Figures 4-15 and 4-16 show the saturation study plots for the Xe I 834.9 nm and the Xe II 835.0 nm experiments, respectively. The Xe I 834.9 nm transition is unsaturated, which is exactly what is needed, but the Xe II 835.0 nm transition is slightly saturated. This is because the relatively low density of the singly-charged xenon ion requires a much higher laser power to probe when compared to the neutral xenon. The red elongated dot on Figure 4-16 indicates the laser power level ultimately chosen for the Xe II 835.0 nm experiment. This laser power level strikes a balance between saturation broadening and SNR. By comparing Figures 4-15 and 4-16, one may also notice that the operating regime spans a greater range of laser power for the neutral saturation study than for the singly-charged saturation study, and that there are error bars on only the singly-charged study. Recall that the singly-charged study used an improved diode laser when compared to the neutral study. This new laser also undergoes significantly lower power drift when

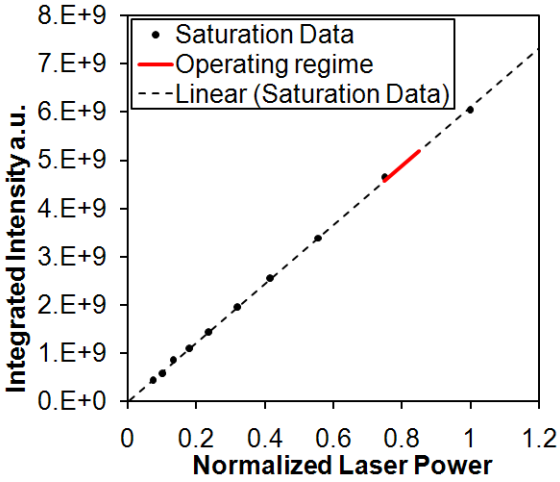

Figure 4-15. Saturation study plot for the Xe I 834.9 nm Zeeman effect study.

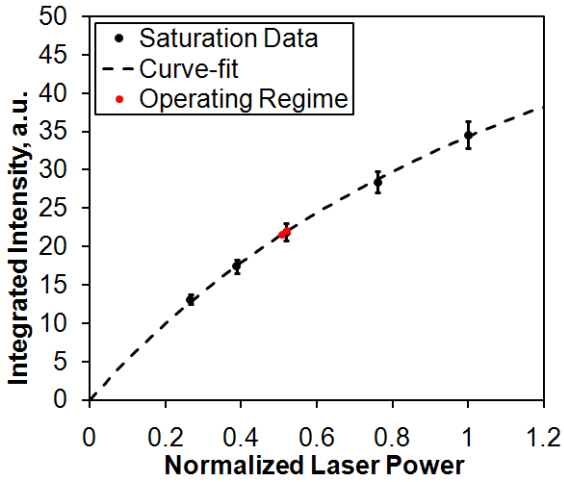

Figure 4-16. Saturation study plot for the Xe II 835.0 nm Zeeman effect study.

operating. However, even with the new equipments, the SNR is much lower for the singly-charged traces, which leads to error bars that are actually visible in the figure. Error bars plotted for the neutral saturation study would not show up very well due to how small the uncertainties are.

Since all of the hyperfine structure constants for the Xe I 834.9 nm transition are known, the purpose of the neutral xenon Zeeman effect study is to validate the use of the existing nonlinear Zeeman model to predict the lineshape. The model was developed by Bailo Ngom for his dissertation work [77] and its original function was to curve-fit raw Zeeman-split traces to back out the magnetic field strength. It works equally well as a predictive tool for simulating the lineshape given a priori information about the external magnetic field. Since none of the hyperfine structure constants for the Xe II 835.0 nm transition are known, the purpose of the singly-charged xenon Zeeman effect study is to find an engineering model that roughly approximates the broadening due to the Zeeman effect.

Figure 4-17 shows a series of σ-polarized Xe I 834.9 nm traces taken at externally applied magnetic field strengths varying from 0 to 300 Gauss. Although traces were taken every 30 Gauss, only traces taken every 60 Gauss are shown to avoid clutter. Figure 4-18 shows the same for the π-polarized traces. It becomes immediate apparent from these results that the amount of hyperfine splitting for π-polarized configuration is negligible in the presence of the Doppler broadening present in the optogalvanic cell. Since the xenon

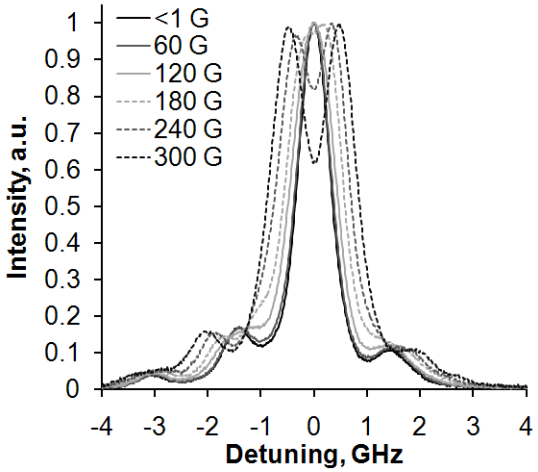

Figure 4-17. Zeeman effect study σ-polarized results for the Xe I 834.9 nm transition.

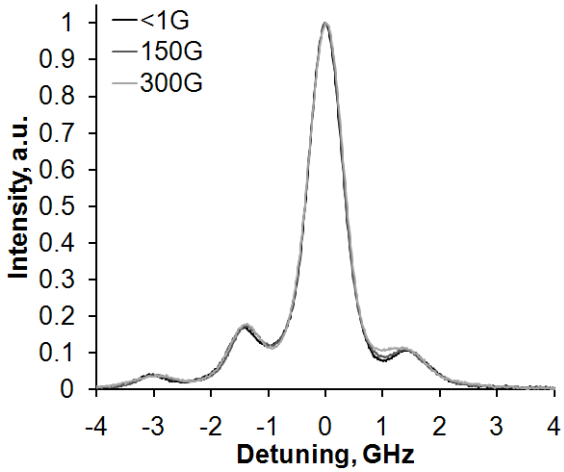

Figure 4-18. Zeeman effect study π-polarized results for the Xe I 834.9 nm transition.

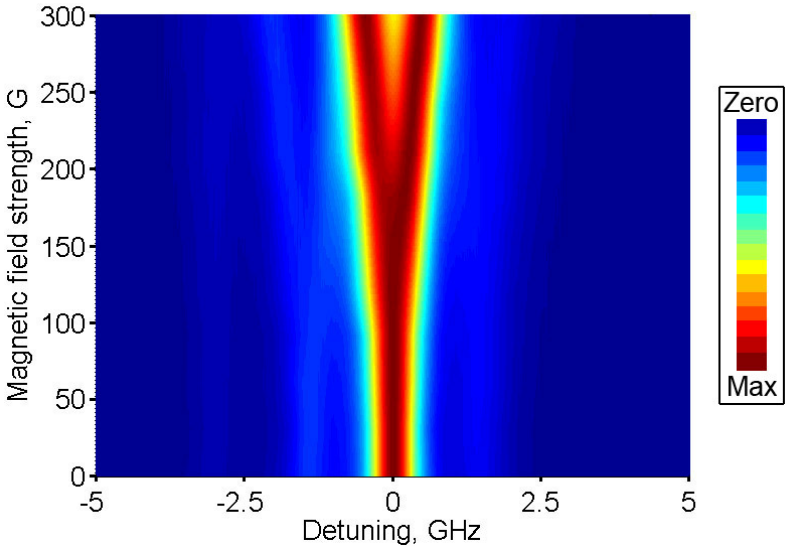

Figure 4-19. Intensity plot showing how the Xe I 834.9 nm peaks split as the σ-component magnetic field increases.

gas in the optogalvanic cell is only a few hundred K in temperature, lower than what is found in Hall thrusters, we can assume the Zeeman effect at magnetic strength up to a few hundred Gauss is negligible compared to Doppler broadening effect if the injected laser beam is π-polarized.

Figure 4-19 shows an intensity plot of how the Xe I 834.9 nm peaks split as the σ-component magnetic field increases. Validation of the nonlinear Zeeman model against the σ-polarized traces were carried by Bailo Ngom as part of his dissertation [77]. The lineshapes predicted by the model were shown to be in excellent agreement with the experimentally obtained optogalvanic traces. Figure 5.22 of Ngom's dissertation [77] is reproduced in Figure 4-20 showing the excellent agreement between the model and experimental data. The success of this model led to a collaborative project to try to map the in-situ magnetic field of the 6-kW Hall thruster [58]. However, for the purpose of this project, it turns out that no meaningful population of high energy

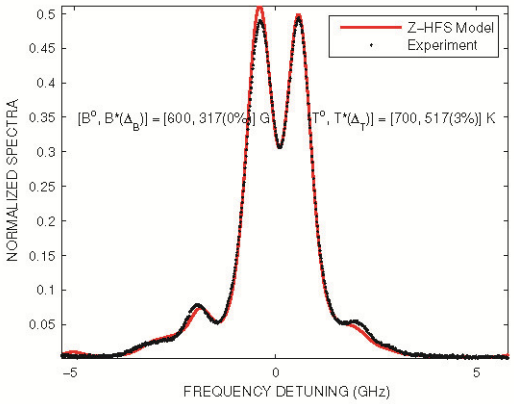

Figure 4-20. Reproduction of figure from Ngom's dissertation [77] showing excellent agreement between nonlinear Zeeman model and optogalvanic trace.

neutral xenon was detected in the main two-axis LIF experiment. Thus there was no need to try to extract the xenon neutral VDFs for the purpose of this dissertation. Nevertheless, this neutral xenon Zeeman study gave valuable insights into the Zeeman-effect-related behavior of the Xe II 835.0 nm transition. In particular, the fact that π-polarized Zeeman effect produce negligible splitting in a Hall thruster setting is an invaluable piece of knowledge.

As a side note, with no relation to this dissertation, it should be mentioned that the sensitivity of xenon Zeeman splitting to polarization direction can in theory be used to determine the direction of the in-situ magnetic field. We begin by assuming that the magnetic field in a Hall thruster is azimuthally uniform and then take two LIF trace at the same spatial location. The first trace should be taken such that the laser polarization is in the azimuthal direction and therefore purely σ-polarized. The second trace should be taken such that the laser polarization is parallel with the local radial-axial plane of the channel. This scan will contain both σ- and π-components. By using the first trace to back out the in-situ magnetic field strength, one can in theory use the second trace to calculate how much of that field strength is perpendicular to the polarization direction of the second trace and therefore back out the magnetic field direction. The development of this idea is left for some other brave soul to pursue in the future.

4.3.4 Singly-Charged Xenon Zeeman Effect Study Results and Validation

The primary purpose of the Xe II 835.0 nm Zeeman effect study is to find an approximation for the Zeeman effect in the absence of published hyperfine structure constants. Figure 4-21 shows a series of σ-polarized Xe I 835.0 nm traces taken at externally applied magnetic field strengths varying from 0 to 300 Gauss. Although traces were taken every 30 Gauss only traces taken every 60 Gauss are shown to avoid clutter. Figure 4-22 shows the same for the π-polarized traces. Figure 4-23 shows an intensity plot of how the Xe II 835.0 nm peaks split as the σ-component magnetic field increases. Just like in the Xe I 834.9 nm Zeeman effect study, the amount of hyperfine splitting for π-polarized configuration is negligible in the presence of the Doppler broadening present in the optogalvanic cell, and, by extension, the Hall thruster plasma.

In the absence of hyperfine structure constants, a simple engineering approach is used to model the Zeeman effect. To simplify the process, the IMOG kernel shown in Figure 4-14 is approximated as a Gaussian of equivalent FWHM frequency. This simplification is advantageous for three reasons. One, a simple Gaussian has the best known stability property when used in devolution operation. The meaning of this will become more apparent in Section 5.3.3. For now, it is sufficient to know that this reduces the amount of smoothing needed when performing post-processing. Two, the IMOG kernel does not look all that different from a Gaussian. Since it is not possible to make

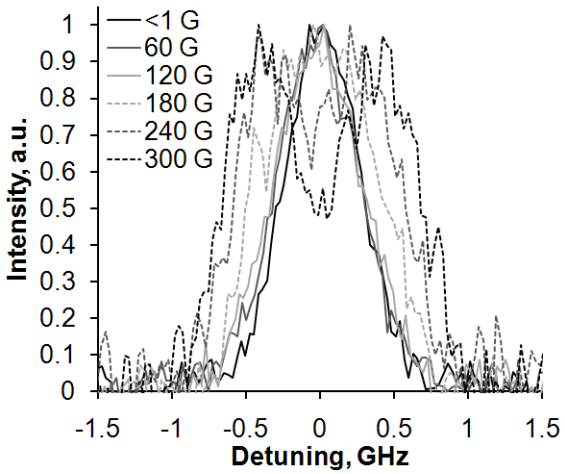

Figure 4-21. Zeeman effect study σ-polarized results for the Xe II 835.0 nm transition.

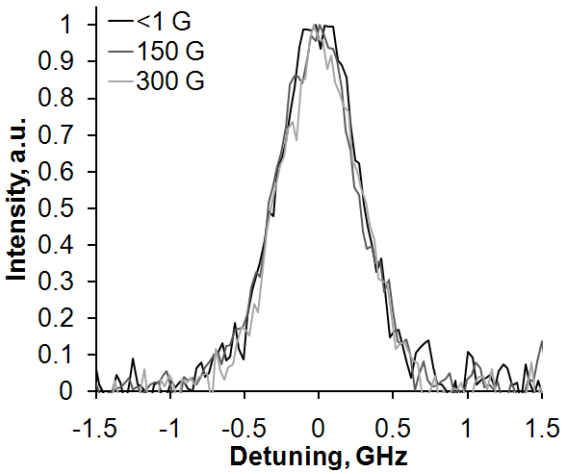

Figure 4-22. Zeeman effect study π-polarized results for the Xe II 835.0 nm transition.

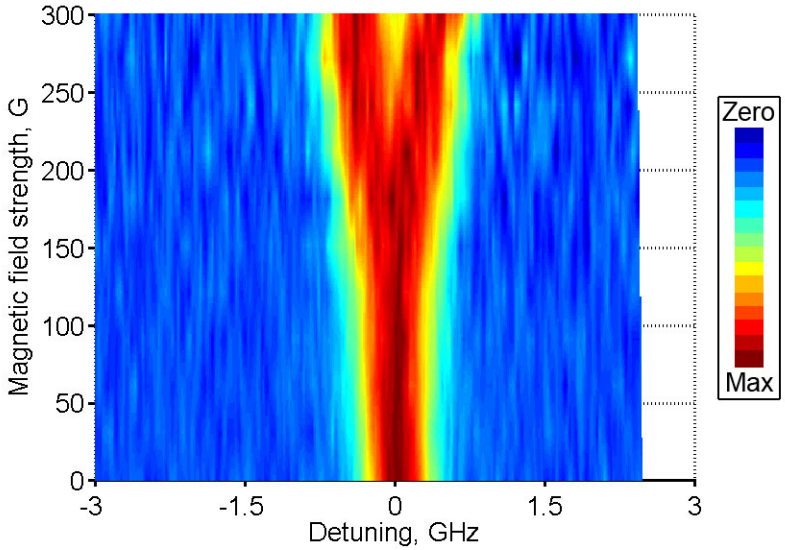

Figure 4-23. Intensity plot showing how the Xe II 835.0 nm peaks split as the σ-component magnetic field increases.

out the individual hyperfine structure lines in the IMOG kernel, the main result of the IMOG spectroscopy is actually to establish an upper limit of how broad the hyperfine structure is. Three, it is easier pass on the data for a 450-MHz-wide Gaussian than it is to pass on the IMOG data for future work. The disadvantage of using the Gaussian approximation is that you lose the little bit of lineshape information that is built into the IMOG kernel. This lineshape information is typically smoothed out by even room temperature Doppler broadening and is not a great loss. Figure 4-24 shows a comparison of the IMOG kernel to a 450-MHz-wide Gaussian.

Next, the σ-component Zeeman effect for the Xe II 835.0 nm transition is assumed to be the same as a single Dirac delta function splitting into two. The frequency difference between the two Dirac delta function is assumed to be a linear function of the magnetic field strength. Or, in other words, the 450-MHz-wide Gaussian simply splits into two equal

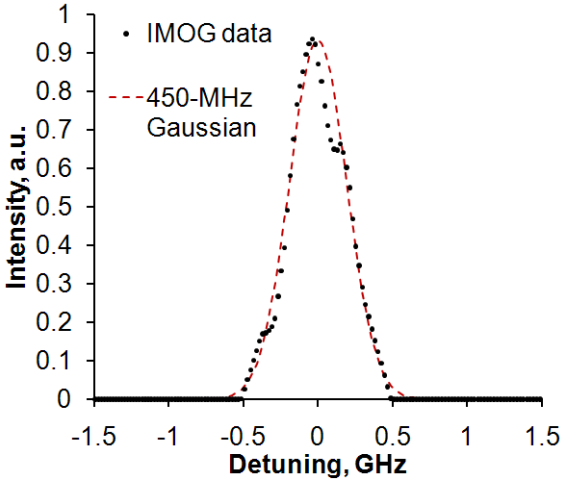

Figure 4-24. Comparison between IMOG kernel and 450-MHz Gaussian.

strength copies with their center being separated by an amount that scales linearly with the magnetic field strength. For the remainder of this dissertation, this model will be referred to as the simple linear Zeeman model (SLZ model).

Figure 4-25 shows a graphical example of how the SLZ model would work if the external magnetic field causes a 1-GHz split in the hyperfine structure. If we now further convolve this hyperfine structure with the Doppler broadening function of the optogalvanic cell, we can predict the raw optogalvanic lineshape that we should obtain for a given applied σ-component magnetic field strength. To obtain the key constants for the SLZ model, we begin by estimating the Doppler broadening function of the optogalvanic cell. Since the Doppler broadening function of a stationary plasma is a Maxwell-Boltzmann distribution, which is just a special form of the Gaussian, and a Gaussian convolved with another Gaussian produces a third Gaussian, we can approximate both the un-split hyperfine structure and the Doppler broadening together. This was found to be a Gaussian with an FWHM frequency of ~610 MHz.

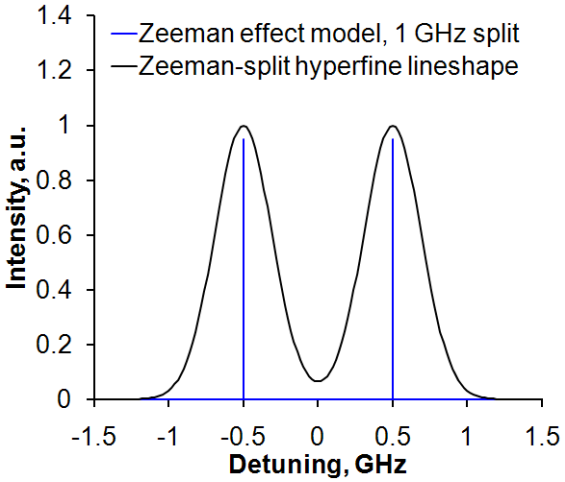

Figure 4-25. Graphical example of a simple linear Zeeman model for the Xe II 835.0 nm transition.

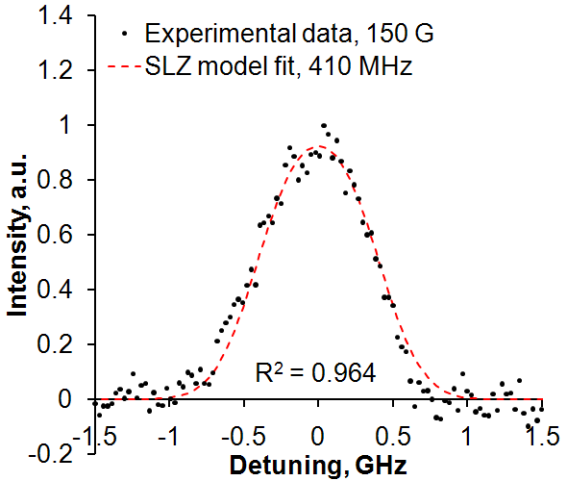

Figure 4-26. Performing curve-fit of the SLZ model to 150-G experimental data.

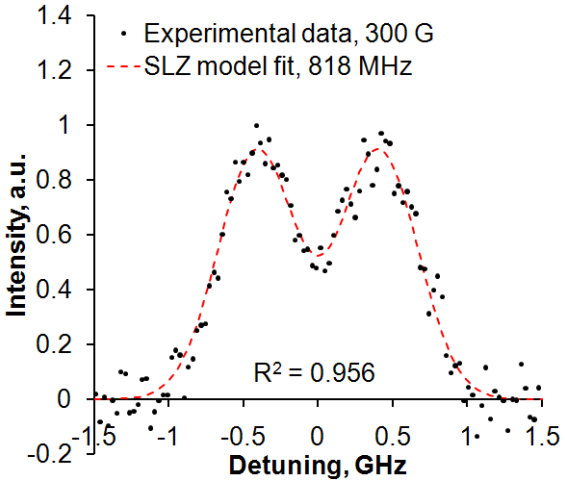

Figure 4-27. Performing curve-fit of the SLZ model to 300-G experimental data.

Figures 4-26 and 4-27 show the results of curve-fitting the SLZ model against the σ-component experimental data at field strengths of 150 G and 300 G, respectively. The curve-fit is carried out using Matlab v7.3.0 cftool.m. The specific algorithm is the Trust-Region non-linear least-squares algorithm. These figures show that the σ-component Zeeman effect on the Xe II 835.0 nm transition can be adequately approximated by the SLZ model.

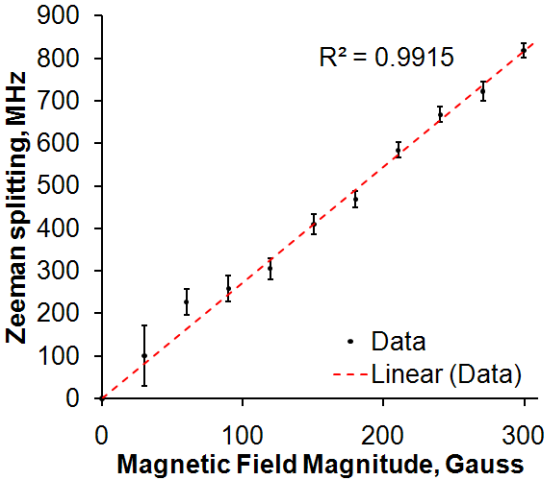

Figure 4-28. Amount of Zeeman splitting based on the SLZ model versus σ-component magnetic field strength.

By performing SLZ model curve fits on all of the collected Zeeman effect study data, we can plot the amount of Zeeman splitting as a function of the σ-component magnetic field strength. This plot is shown in Figure 4-28. As this figure shows, the amount of Zeeman splitting is indeed approximately linear over the range of 0 to 300 Gauss. The uncertainty in the amount of Zeeman splitting is determined by standard residual analysis built into the Matlab program. The final result is shown in equation (4-1).

$$\text{Amount of Zeeman splitting in MHz} = 2.7273 \times \text{magnetic field strength in Gauss} \quad (4\text{-}1)$$

This engineering estimate will serve as the basis for approximating the hyperfine structure of the Xe II 835.0 nm transition when under the influence of an external magnetic field.

Chapter 5
Laser-Induced Fluorescence Velocimetry

The main purpose of performing laser-induced fluorescence (LIF) velocimetry is to learn more about the particles that bombard and erode away the ceramic channel walls of the Hall thruster. In particular, we are interested in studying three characteristics; one, the energy and angle with which xenon particles bombard the walls; two, the location where the xenon particles are energetic enough to cause erosion; and three, the influence of varying the discharge voltage and the anode mass flow rate, the two main control parameters in this dissertation, on number one and two.

To quantify these characteristics, a series of LIF experiments were performed. This section will describe the two experiments that are most central to answering the dissertation questions. The first experiment is a one-axis singly-charged xenon LIF study to map the acceleration of the Hall thruster along the channel centerline. The second experiment is actually a multi-part study to extract two-axis VDFs of both neutral and singly-charged xenon inside the discharge channel of the Hall thruster with emphasis on near-wall data. The Xe I 834.9 nm transition is used to probe the neutral xenon while the Xe II 835.0 nm transition is used to probe the singly-charged xenon.

5.1 Principles of Laser-Induced Fluorescence Velocimetry

The basic principles of LIF velocimetry were described in Section 3.2.1 and will be briefly summarized here for reference.

When an atom is moving in the direction of a photon's wave vector, the atom absorbs the photon at a shifted frequency due to the Doppler effect. To actually take advantage of this effect for velocimetry purposes, one must vary the frequency of the injected laser light and measure the intensity of the collected fluorescence. For example, if one measures 3 units of fluorescence at some frequency, υ_1, and 5 units at the

stationary frequency, υ_0, then the population at velocity, v, is 3/5 times as large as the population that is at rest at the interrogated location. The velocity, v, is calculated using equation (5-1) where positive velocity means traveling toward the photon source.

$$\frac{\upsilon_1 - \upsilon_0}{\upsilon_0} = -\frac{v}{c} \tag{5-1}$$

The physics of LIF velocimetry leads to some practical limitations that must be observed. First, LIF generally produces only relative population results, which is all that is required for mapping VDFs but not sufficient to determine density. Second, good velocimetry measurement also requires a good stationary reference since there would otherwise be no absolute reference for converting photon frequency to particle velocity. Third, it is only possible to map the velocity in one direction for each laser beam used. Fourth, only particles that interact with the laser beam are probed. In other words, the population measured is of the lower state of the transition. Since there are no physical forces in the Hall thruster that distinguish between the various energetic states of a species, one can assume that the lower state of the transitions has the same VDF as all of the other states of the associated species.

5.2 Experimental Setup for the LIF Experiments

This section describes the experimental setup for both the one-axis acceleration zone study and the two-axis near-wall study. If both studies used the same setup for a particular part, then only one description will be given for that part.

5.2.1 Facility

Experiments are performed in the Large Vacuum Test Facility (LVTF) of the Plasmadynamics and Electric Propulsion Laboratory (PEPL) at the University of Michigan. The LVTF is a ϕ6 m × 9 m stainless steel-clad cylindrical steel chamber. The thruster plume is allowed to expand un-impeded until termination at a carbon beam dump ~4 m downstream. Pumping is provided by seven cryopumps with a nominal xenon pumping speed of 245,000 l/s. Facility pressure is monitored by two hot-cathode ionization gauges and any reported pressure reading is the average of the two. Typical base pressure is ~2×10^{-7} Torr. The pressure during thruster operation at 30 mg/s anode mass flow rate is ~2×10^{-5} Torr, corrected for xenon [79]. During the experiments, the 6-kW Hall thruster is mounted on two cross-mounted stepper-motor-driven translation stages. The stages provide 1 m of axial travel, 2 m of radial travel, and have a positioning accuracy of ±0.1 mm. Laser injection and fluorescence collection optics are fixed to the floor of the chamber.

5.2.2 Thruster Settings and Test Matrices

Seven operating conditions spanning discharge voltages of 150-600 V and anode mass flow rates of 10-30 mg/s are considered. These operating conditions are listed in Table 2-2, and the discharge voltage, anode mass flow rate, and discharge current are re-listed here in Table 5-1 for quick reference.

Table 5-1. Thruster operating conditions for the LIF experiments.

Discharge voltage, V	Anode mass flow rate, mg/s	Discharge current, A
150	10	9.08
150	20	21.55
150	30	35.35
300	10	8.94
300	20	20.11
300	30	33.25
600	10	9.36

As mentioned in Section 2.4, the magnetic

field strength is set to maximize anode efficiency while maintaining symmetry across the channel centerline and the cathode flow fraction is 7%. Research-grade xenon is supplied to both the anode and the cathode with separate commercial flow meters and controllers from MKS at an accuracy of ±1%.

The main discharge is powered by an HPS1000-100-K027 100-kW power supply capable of 1000 V and 100 A with a 0.4 V_{rms} ripple. This supply is isolated from the strong oscillations that can occur in the Hall thruster discharge by a 95-μF capacitor that is in parallel with the anode and cathode. Separate Sorenson power supplies power the inner magnetic coil, the outer magnetic coil, the cathode heater, and the cathode keeper. The trim coil is not in use.

The discharge voltage is monitored with a BK Precision 2880B multimeter at an accuracy of ±0.05%. Discharge current is monitored using a high-precision calibrated shunt with a BK Precision 5491A multimeter for a net accuracy of ±0.2%. This accuracy is driven by the shunt.

For the experiments, the thruster is baked out at 300 V, 20 mg/s prior to testing. The bake out time typically takes 40-50 minutes and is characterize by an initial rise in discharge current followed by a slow drop to steady-state. Steady-state status is considered to be reached when the discharge current change by less than 0.05 A in 10 minutes. When switching between operating conditions, the thruster is allowed to run for at least 10 minutes at the new condition or however long it takes to reach steady-state again before data is taken.

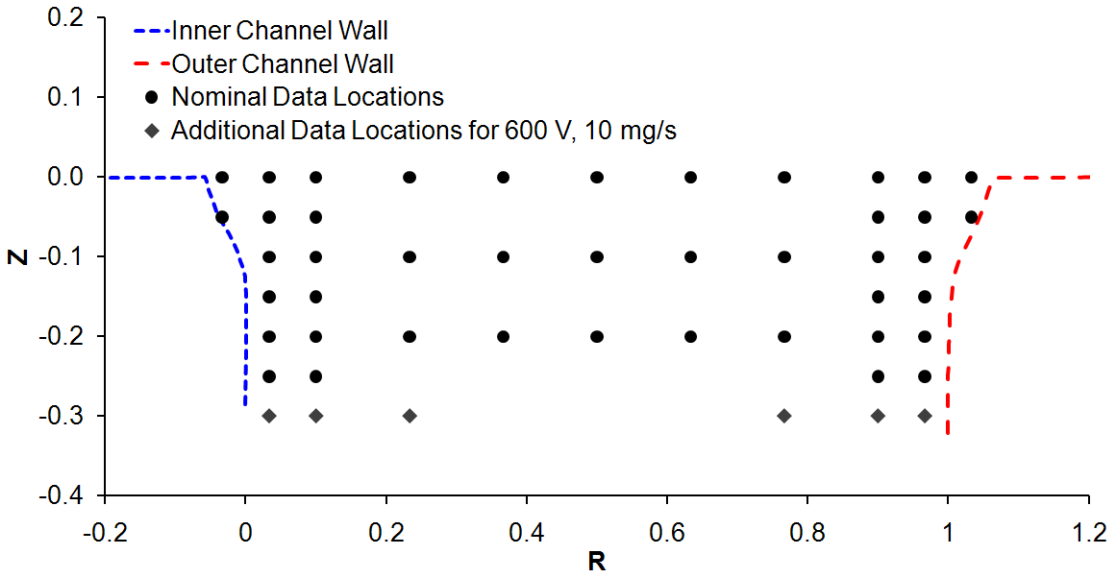

Figure 5-1. Data location matrix for the two-axis near-wall LIF experiment.

For the one-axis acceleration zone study, data was taken along the channel centerline from Z = -0.375 to Z = 0.375 at 0.075 spacing. Additional data was taken at Z = 0.5, 0.75, 1.5, and 2.5 if SNR is high enough at those locations. Figure 5-1 shows the data location matrix for the two-axis near-wall study along with the channel profile obtained from the first profilometry measurement. Z is the non-dimensional axial coordinate and is the axial coordinate normalized by the channel length. Z = 0 at the channel exit plane and Z = -1 at the anode exit. R is the non-dimensional radial coordinate and is the radial coordinate normalized by the nominal (non-eroded) channel width. R = 0 at the inner channel wall and R = 1 at the outer channel wall. Note that calculations using data from Shastry's dissertation work [61], which was carried out on the same 6-kW thruster at many of the same operating settings, shows the average Debye length is ~0.05 mm. Assuming 95% of the sheath drop is within 5 Debye lengths, the sheath is ~0.25 mm thick. This means the near-wall LIF interrogation zones negligibly overlap with the sheath except at the Z = -0.05, R = -0.033 and the Z = -0.05, R = 1.033 data locations.

5.2.3 Air-Side Laser Injection Setup

Figure 5-2 shows the air-side laser and optics setup for the one-axis acceleration zone study. The laser system used is a TOPTICA TA-100/830 tapered-amplifier diode laser. This system has a nominal linewidth of ~10 MHz and a mode-hop-free range of

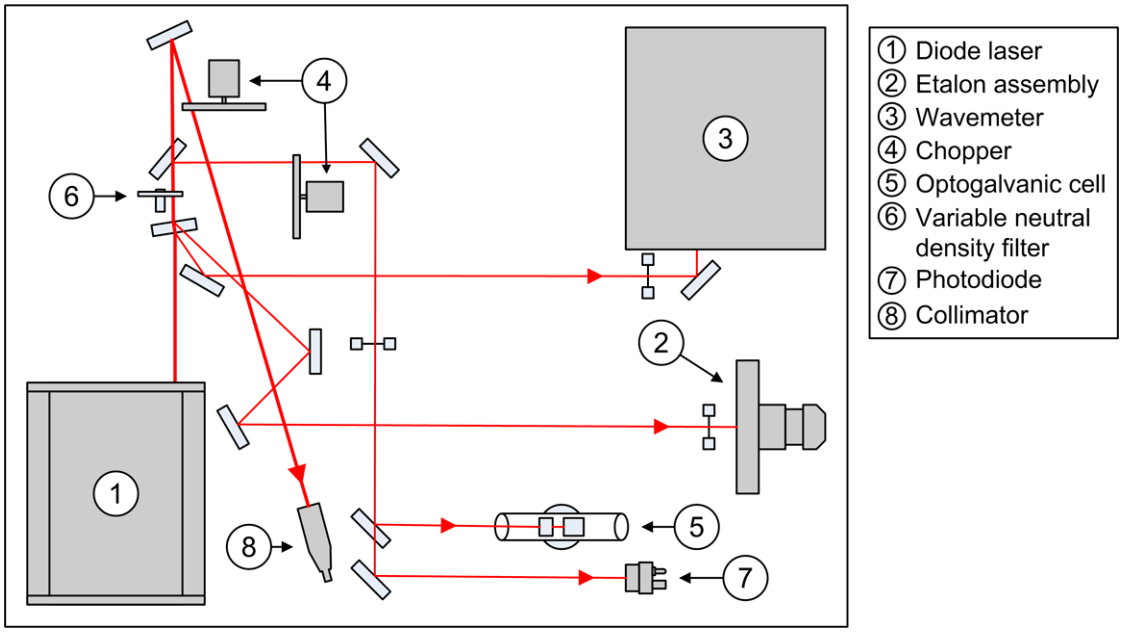

Figure 5-2. Air-side laser injection setup for the one-axis acceleration zone LIF study.

~12 GHz. The total output power is ~250 mW. The setup uses a Burleigh SA-91 etalon assembly (2-GHz free spectral range, finesse >250) and a WA-1000 wavemeter (±1 pm) for reference information.

A relatively weak laser beam is split off and sent into the optogalvanic cell for stationary reference. This beam is mechanically chopped at ~1.1 kHz. The optogalvanic cell is the same as the one used in IMOG spectroscopy and Zeeman effect studies. It is a Hamamatsu L2783-42 XeNe-Mo galvatron. The core of this cell is a pair of ϕ6.25 mm cylindrical Mo tube electrodes centered in a ϕ25 mm x 120 mm glass cylinder. The cylinder is filled with approximately 3 Torr of xenon and 4 Torr of neon. The ends of the cylinder are angled at approximately 10 degrees from being perpendicular to the electrode axis to eliminate retro-reflection. The output current is passed through a ballast resistor and the voltage drop across the resistor is connected to an SR-810 lock-in amplifier through a high-pass RC filter. The optogalvanic cell is operated at 250 V for all LIF studies.

Additionally, a Thorlabs DET-110 photodiode is used to monitor laser power drift. This reading is fed into another SR-810 lock-in amplifier for signal collection. For an unsaturated LIF trace, this laser power reading is used to remove the effect of laser power drift from the trace.

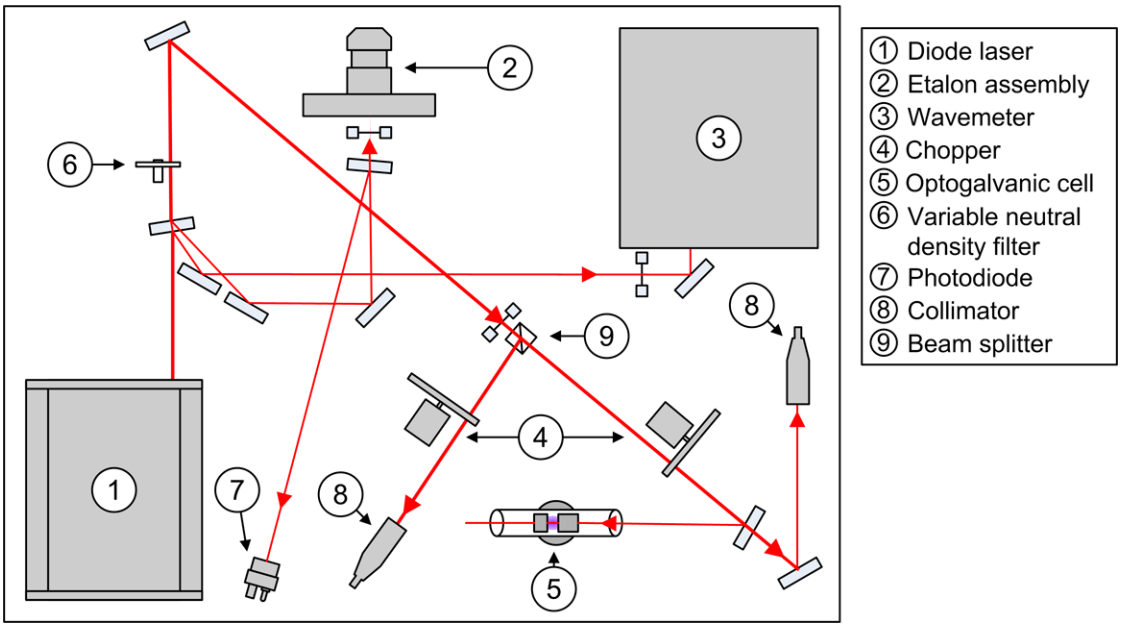

Figure 5-3. Air-side laser injection setup for the two-axis near-wall LIF study.

The main laser beam is mechanically chopped at ~1.8 kHz via an SR540 chopper, and sent into a fiber collimator. The fiber delivers the light through an optical feedthru into the vacuum chamber. The power of the laser injected into the collimator is ~100 mW. The laser power drops to ~20 mW, measured via thermopile, by the time it arrives at the interrogation, corresponding to a beam intensity of ~25 mW/mm^2. When doing saturation study, an additional variable neutral density filter is added so that the laser power being injected into the chamber can be controlled.

Figure 5-3 shows the air-side laser and optics setup for the two-axis near-wall study. The laser system used is the same one used for the one-axis acceleration zone study. The SA-91 etalon assembly, the WA-1000 wavemeter, the L2783-42 optogalvanic cell, and the DET-110 photodiode are also re-used for reference information.

Unlike the one-axis setup, the main laser beam is split roughly 50:50 by a non-polarizing beam splitter cube. The transmitting beam is chopped by a SR540 mechanical chopper at ~1.2 kHz before being collimated into the axial injection fiber. The reflected beam is chopped by another SR540 chopper at ~1.7 kHz before being collimated into the off-radial injection fiber. Total laser beam power (from both the axial and the off-radial beams) delivered to the interrogation volume was measured to be ~20 mW giving an beam intensity of ~25 mW/mm^2.

During the experiments, all scans were taken with a lock-in amplifier time constant of 100 ms and a scan-duration of 160 s. Amplifier settings and PMT bias voltage are fixed throughout for each experiment to ensure data are taken under comparable conditions.

5.2.4 Vacuum-Side Setup

Figure 5-4 shows a diagram of the vacuum-side experimental setup for the one-axis acceleration zone study. The injection optics sends the laser axially into the thruster. This beam is focused down to a point with a 1-mm diameter circular cross section via an anti-reflect-coated plano-convex lens. A polarizer (not shown in the figure) is placed between the optical fiber output and the lens to provide horizontal polarization with respect to the chamber floor. The injected beam interrogates the thruster plasma at the nine o'clock position when viewed face-on, where the local magnetic field is approximately parallel to the laser polarization vector. This configuration gives rise to purely π-polarized Zeeman effect, which has been shown in Section 4.3.4 to be negligible. The injection beam focusing cone has a half-angle of ~0.5° and the injection axis is aligned to within 0.1° of the thruster firing axis so cosine losses are negligible. To prevent excess thermal drift during thruster operation, a shield optic is installed to allow the laser through while blocking the incoming xenon particles. Optical fiber beam delivery is chosen over mirror-based delivery because it is easy to deliver light to different places in the vacuum chamber and at varying angles. Also, mirrors can become coated by sputter products easily and require protection. The fiber used for delivering near-infrared laser is the 50-μm low-OH multi-mode fiber from Thorlabs.

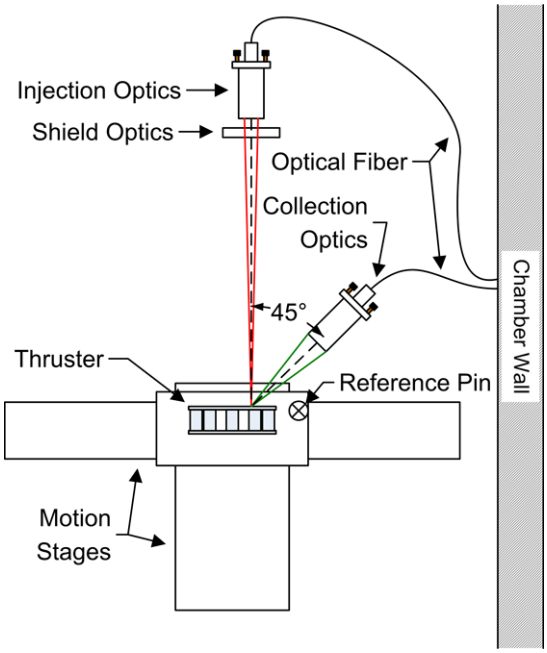

Figure 5-4. Vacuum-side setup for the one-axis acceleration zone LIF study.

The collection optics is built from a matching pair of anti-reflect-coated

achromatic lenses. Thermal shielding is not necessary for this optics based on thermocouple measurements taken during the test. Nevertheless, the overall LIF system may undergo thermal drift over the course of the test so a reference pin is installed to the left of the thruster when viewed face on. The exact position of this pin relative to the thruster is measured and the reflected laser signal from this pin is used to compensate for thermal drift. The collection optics is set at a 45±1° angle with respect to the thruster firing axis to allow the interrogation of the inside of the thruster channel. The resulting intersection of the injection and collection optics forms a diagonally sliced cylinder ~1 mm in diameter in the radial direction and ~1.5 mm long in the axial direction. This interrogation zone size sets the spatial resolution of the one-axis acceleration zone study. The optical fiber used to bring the fluorescence out of the vacuum chamber is the 1000-μm low-OH multi-mode fiber from Thorlabs.

Figure 5-5 shows a diagram of the vacuum-side experimental setup for the two-axis near-wall study. Subplot (a) shows how the two sets of injection optics are aligned with respect to the thruster for probing near the inner channel wall. Shown in dash outline is the position of the off-radial injection optics when probing near the outer channel wall. The axial injection optics and the shield optics are identical to those in the one-axis acceleration zone study setup. The off-radial injection optics is so named because it sends in laser light at a 60±1° angle with respect to the thruster axis, which is close to but

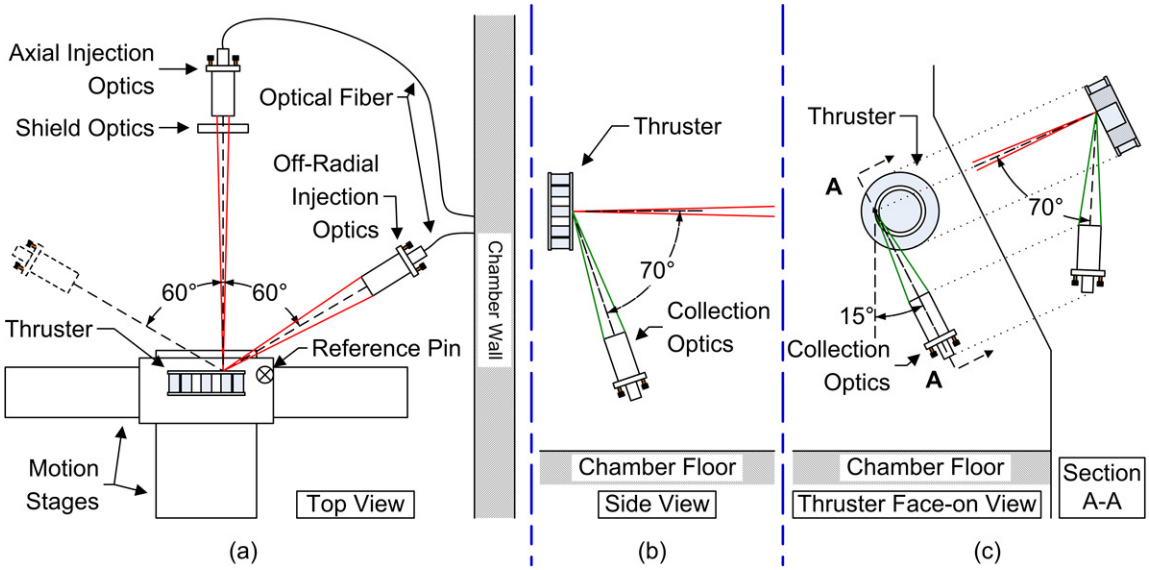

Figure 5-5. Vacuum-side setup for the two-axis near-wall LIF study.

slightly off from perfectly radial injection. Although analysis will be more complicated with an off-radial scheme, it is necessary for the laser light to reach to inside of the thruster channel. The beams are focused down to a point with a 1-mm diameter circular cross section via anti-reflect-coated plano-convex lenses. The injection beam focusing cones have half-angles of ~0.5° and the injection axes are aligned to within 0.1° using laser tools so cosine errors are negligible. The interrogation zone is at the 9 o'clock position when viewing the thruster face on.

Polarizers (not shown in the diagram) are placed between the optical fiber outputs and the lenses to control the beam polarization with respect to the local magnetic field direction. Note that a photon's polarization direction is always perpendicular to the photon's travel direction. For this reason, the polarization of the axial beam can be tuned to match the local magnetic field direction in order to minimize the Zeeman effect, but it is not possible to do the same for the off-radial beam. Instead, we opt to make the polarization of the off-radial beam perpendicular to the local magnetic field in order to excite only the σ-component splitting. We will use the SLZ model described in Section 4.3.4 to remove the broadening caused by the Zeeman effect.

Figure 5-5 subplots (b) and (c) show the collection optics positions used to collect data from near the inner and the outer channel walls, respectively. The axial injection beam is shown for reference. The collection optics was built from a matching pair of anti-reflect-coated achromatic lenses. Thermal shielding is not necessary for this optics based on thermocouple measurements taken prior to testing. To counter the overall system thermal drift, a reference pin with a 1-mm-diameter hemispherical head is installed to the left of the thruster when viewed face on. The exact position of the head of this pin relative to the thruster is known and the reflected laser signal from this pin is used to compensate for thermal drift. Additionally, the yaw movement of the axial injection optics is motorized with a Thorlabs Z812 micro-stepper motor because this movement axis was found to undergo the most thermal drift during testing. The collection optics axis is set at a 70±0.5 degree angle with respect to the thruster axis to allow the interrogation of the inside of the thruster channel. The interrogation volume resulting from the intersection of the injection and collection optics is a sphere of ~1.5 mm in

diameter. The size of this interrogation volume sets the spatial resolution of the two-axis near-wall study.

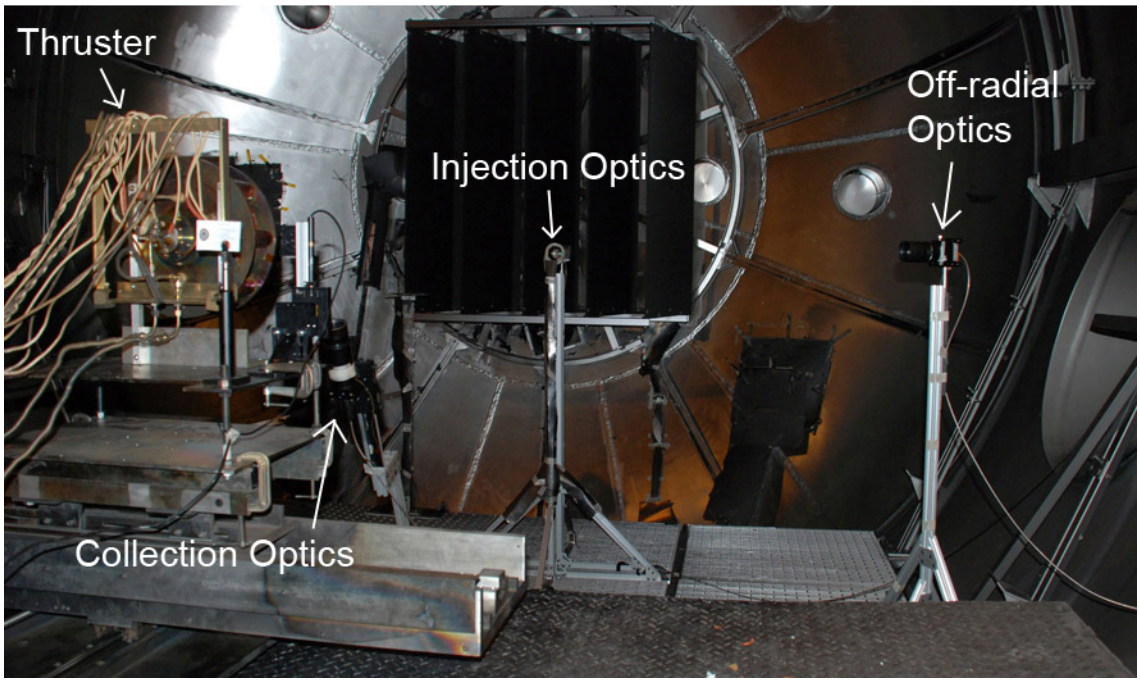

Figure 5-6. Photograph of the vacuum-side experimental setup for taking LIF measurements near the inner wall of the channel.

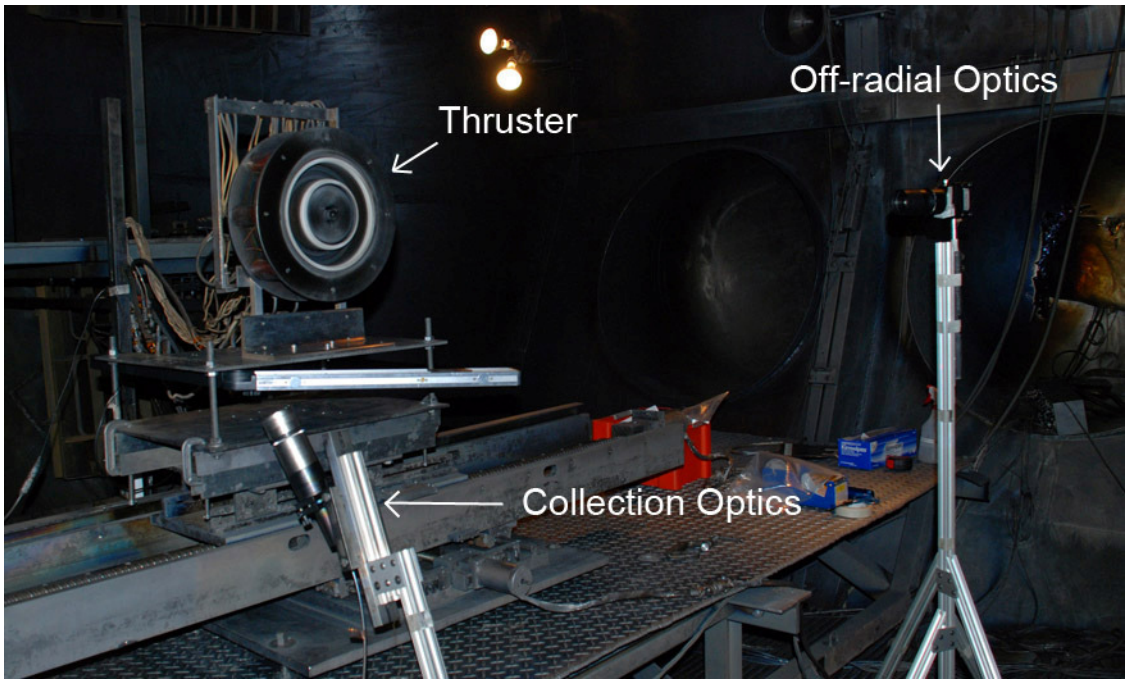

Figure 5-7. Photograph of the vacuum-side experimental setup for taking LIF measurements near the outer wall of the channel.

Figures 5-6 and 5-7 show photographs of the two different configurations used for two-axis near-wall LIF. In Figure 5-7, the tool attached to the bottom of the platform on which the 6-kW Hall thruster sat is equipped with a laser beam that helped precisely position the off-radial optics at 60° with respect to the firing axis to within 0.1°. Figure 5-8 shows a photograph of the thruster in operation during the two-axis near-wall LIF test.

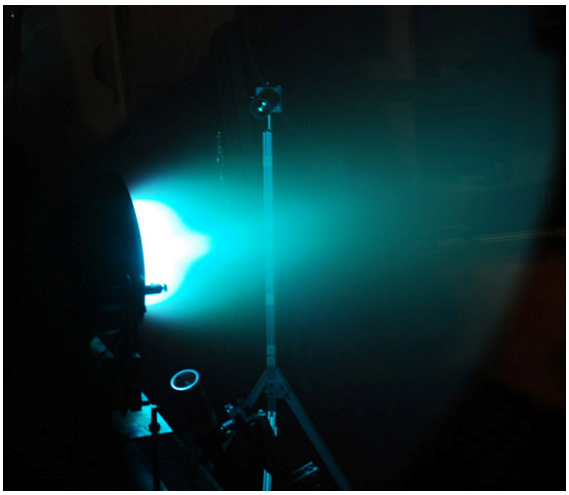

Figure 5-8. Photograph of the thruster in operation during the two-axis near-wall LIF test.

5.2.5 Air-Side Data Collection Setup

For all LIF studies, an optical fiber brings the fluorescence signal from the chamber to a SPEX-500M monochromator, the output of which is amplified by a Hamamatsu R928 photomultiplier. The current output of the photomultiplier is converted to a voltage signal via an Oriel 70710 current preamplifier. For the one-axis test, the amplified signal is sent to a SR850 lock-in amplifier, while for the two-axis test, the signal is also sent to a SR810 lock-in amplifier; the SR850 locks onto the chopping frequency of the axial injection beam while the SR810 locks onto the chopping frequency of the off-radial injection beam. The monochromator slit size is 1 mm, which corresponds to an optical bandwidth of ~1 nm. This size is selected to allow a fair amount of signal to reach the PMT while filtering out natural emission at other wavelengths. Based on information from the NIST atomic spectra database, ver. 4, the wavelength range formed by a bandwidth of 1 nm around the 541.9 nm (air) transition includes one other xenon ion line and two other xenon neutral lines. The 541.9 nm (air) transition is listed on the database as at least two orders of magnitude stronger than these three other lines so we can be reasonably assured that most of the signal, and noise, comes only from the target transition.

All lock-in amplifiers communicate with the data acquisition (DAQ) computer via GPiB commands. The auxiliary output of one of the lock-in amplifiers is also used to

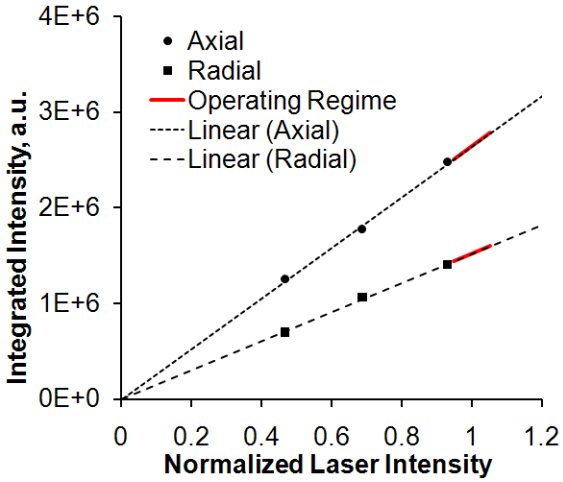

Figure 5-9. Saturation study plot for the Xe I 834.9 nm transition of the two-axis near-wall study.

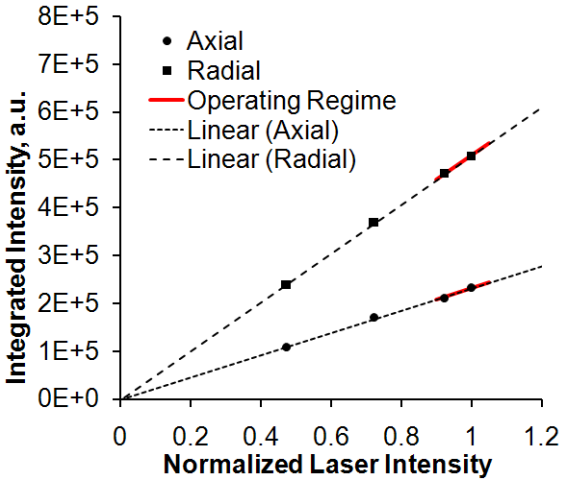

Figure 5-10. Saturation study plot for the Xe II 835.0 nm transition of the two-axis near-wall study.

control the piezoelectric scanning of the laser's wavelength. The DAQ computer uses a custom Labview 7.1 program to control the laser and to assemble the data from the various lock-in amplifiers and monitoring devices.

5.3 Data Reduction for the LIF Experiments

5.3.1 Saturation Studies

Figures 5-9 and 5-10 show the saturation study plots for the two-axis near-wall study. The uncertainties on the intensities are ~1% each and the associated error bars are too small to show.

Figure 5-11 shows the saturation study plot for the one-axis acceleration zone study. Unlike the two-axis studies, the laser power was measured with a thermopile because the photodiode was not in place for the saturation study. This increases the uncertainty on the laser intensity and associated error bars are shown.

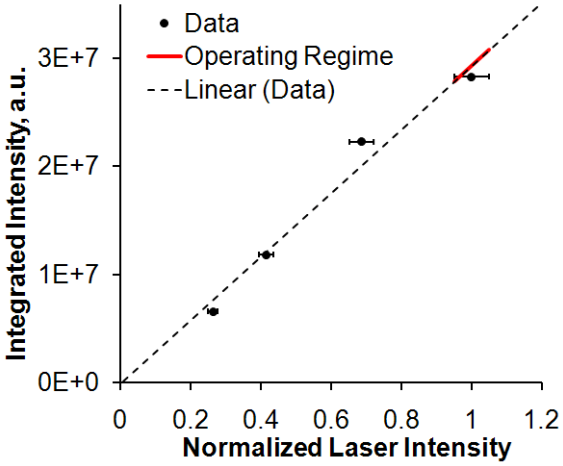

Figure 5-11. Saturation study plot for the one-axis acceleration zone study.

5.3.2 Stitching Scans

As previously mentioned, the maximum mode-hop-free range of the TA-100/830 laser system is ~12 GHz but the VDFs of the singly-charged ion at many locations can exceed this range. To obtain complete VDFs at these locations, it is necessary to stitch multiple LIF scans together. This section describes how the stitching is done.

Whenever multiple LIF scans are needed at a data location, the scans are taken within 15 minutes of each other to avoid excess drift between scans. This is accomplishable by making sure multiple mode-hop-free ranges covering different frequency spans are available to the diode laser before the test. These different ranges are accessed by using different diode currents the value of which must be found beforehand. When properly tuned, up to five different 12-GHz mode-hop-free spans that overlap slightly are available on the TA-100/830, though usually only four is achievable. Up to 40 GHz of total scanning range is typically available with proper tuning. For the two-axis near-wall LIF study, the biggest number of scans that had to be stitched together was three and this was done only a few times. The stitching of two scans together is more common. Stitching, in general, was only needed for 300 and 600 V operating conditions.

Figure 5-12 illustrates how the stitching of two LIF scans is done. The data comes from axial scans of the Xe II 835.0 nm transition taken at Z = -0.1, R = 0.633 while the thruster operates at 300 V, 10 mg/s. The black dots and the blue dots represent data from two different scans that overlapped. One set of data is scaled by an adjustable multiplier and shifted by an adjustable frequency. In the overlap region, data from two sets are

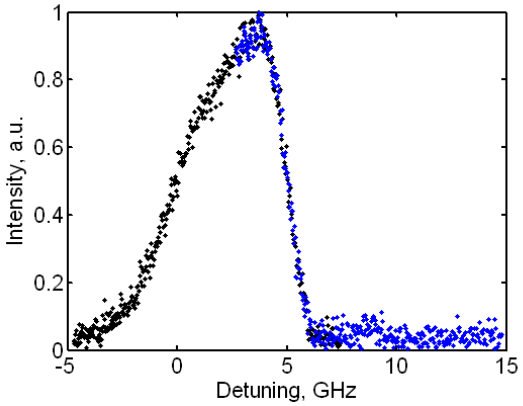

Figure 5-12. Example of how two LIF scans are stitched.

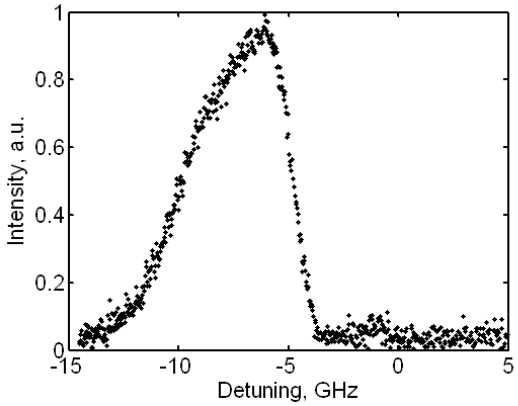

Figure 5-13. Example of the result of stitching two LIF scans.

paired and an error parameter is calculated by taking the root mean square of the difference. The adjustable parameters are iterated until the error is minimized then human judgment is applied to make sure the two scans look like they actually overlap. The final result is produced by averaging pairs of data points in the overlap region so the data density in frequency space is still roughly the same everywhere. Figure 5-13 shows the result of the stitching operation applied to the data shown in Figure 5-12.

5.3.3 Deconvolution Methods

The deconvolution operation is used to remove broadening and distortion in the lineshape caused by effects other than the Doppler effect. The main sources of broadening and distortion that must be accounted for are hyperfine structure and the Zeeman effect. As mentioned in the experimental setup section, the axial laser beam excites only the π-component Zeeman effect so only hyperfine structure needs to be removed from axial scans. The off-radial beam, on the other hand, excites only the σ-component Zeeman effect so both hyperfine structure and the Zeeman effect must be accounted for. In order to use the SLZ model to predict the Zeeman-split hyperfine lineshape, we need to know a priori the in-situ magnetic field strength. This information is unfortunately not available. However, the magnetic field strength obtained from MagNet v6 simulations is available, and it has been shown in a previous in-situ magnetic field measurement experiment that the in-situ field strength is at most 30 Gauss off from simulated field strength [58]. Using the SLZ result shown in equation (4-1), the maximum error in the FWHM frequency of the hyperfine lineshape is ~80 MHz, or ~70 m/s when converted to velocity. Considering that the typical FWHM velocity of Xe II lineshapes are 1000+ m/s this amount of error is considered acceptable.

Note that when a 1000 m/s wide lineshape is broadened by a 70 m/s wide effect, the percentage of broadening is not 70 divided 1000, or 7%. The broadening is actually the result of a convolution function. When two Gaussian functions are convolved the FWHM value of the result is equal to the square root of the sum of the squares of the individual FWHM values. In order words, the actual amount of broadening is square root of the sum of 70^2 and 1000^2 subtracted by 1000, which is 2.4 m/s, or only 0.24% broadening. For this reason, any Gaussian-type broadening factor that is not at least of the

same order of magnitude as the source function being broadened will produce negligible amount of broadening.

We will now go over two methods of carrying out deconvolution that have been considered for this dissertation and the advantages and disadvantages of each.

5.3.3.1 Fourier Transform with Inverse Gaussian Filtering

Mathematically, the convolution operation can be carried out using Fourier transforms. The convolution theorem states that the Fourier transform of the result of a convolution operation is equal to the point-wise product of the Fourier transform of the individual components. This is shown in equation (5-2),

$$c(\upsilon) = a(\upsilon) \otimes b(\upsilon) \Leftrightarrow C(\omega) = A(\omega) * B(\omega) \tag{5-2}$$

where A, B, and C are the Fourier transform of a, b, and c, respectively. Similarly, the deconvolution operation is simply point-wise division in Fourier space.

Although the idea of using Fourier transform to perform fast deconvolution seems simple, actual implementation of the technique is more complicated. This is because the division operation is inherently unstable, which is another way of saying the extracted lineshape is highly sensitive to noise in the raw trace. To put it another way, the relative uncertainty in the raw trace is greatly magnified through the deconvolution operation and will overwhelm the signal. Tim Smith's dissertation work includes a very detailed treatment of this problem and how to suppress the noise [80]. In particular, it was concluded that the addition of an inverse Gaussian filter is the best way to stabilize the

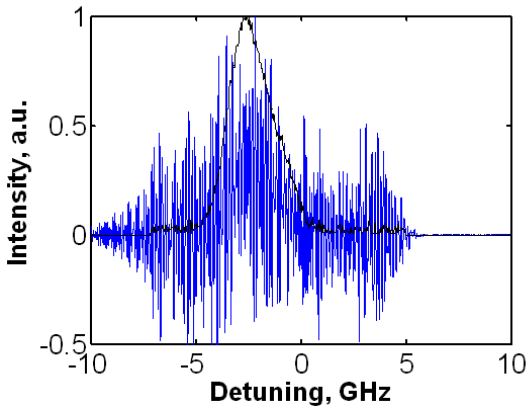

Figure 5-14. Example of how sensitive the solution is to noise in the raw trace.

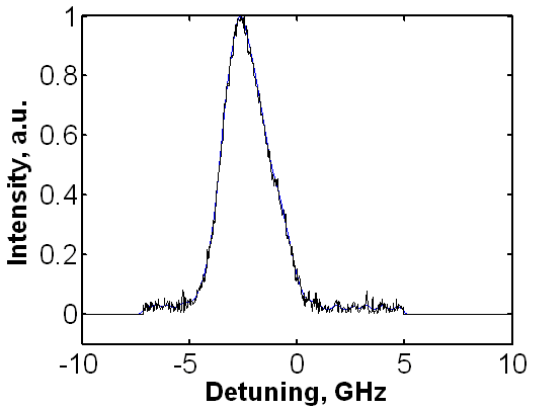

Figure 5-15. Example of improved solution stability with the application of an inverse Gaussian filter.

solution and is the approach used in this dissertation.

Figure 5-14 demonstrates the inherent instability in the Fourier transform method by applying Fourier transform deconvolution on a raw LIF trace with an SNR of 85 (exact definition of SNR to be defined later, >10 is consider acceptable) while using insufficient stabilization. The black line is the raw data and the blue line is the result of the deconvolution operation. The inverse Gaussian filter used in this operation is 73 MHz wide in frequency space, which is close to not having a filter. Figure 5-15 demonstrates the same deconvolution procedure when a 527-MHz inverse Gaussian filter is applied. Note that strictly speaking an inverse Gaussian filter only exists in Fourier space. However, the Fourier transform of a Gaussian is still a Gaussian and so one way to gauge the smoothing power of an inverse Gaussian filter is to calculate the FWHM frequency of the inverse Fourier transform of the said filter. This will give us an effective measure of how wide the smoothing filter is in the real solution space. Also note that the smoothing filter being a Gaussian of finite width also acts as an artificial broadening factor on the extracted solution. Thus, increasing the smoothing power of the Gaussian stabilizes but also broadens the solution. Fourier transform with an inverse Gaussian filter is not ideal when the size of the smoothing filter needed becomes comparable in width to the extracted solution.

The primary advantage of using a Fourier transform with an inverse Gaussian filter to perform deconvolution is speed. The amount of computational resources needed to carry out this method scales approximately as n, the domain size, because the deconvolution step is performed using point-wise multiplication. The typical run time for this deconvolution method measured using a Matlab implementation on a decent home pc built in 2007 is ~0.3 s. The run time is not dominated by the deconvolution algorithm but by the generation of diagnostic plots.

5.3.3.2 Tikhonov Regularization

The deconvolution operation performed on a raw trace containing noise is an ill-posed problem. Here, ill-posed refers to the fact that there is no exact solution to the problem because the noise in the raw trace makes it impossible to form a single set of values to satisfy the associated convolution equation. An analogy to this would be trying

to draw a straight line through three points but the location of the three points no longer lines up due to noise. While it is impossible to draw a straight line that goes through all three points, it is possible to draw one that fits the general direction formed by the three points. The Tikhonov method tries to find the best approximate solution to the deconvolution problem by finding a solution that has the same general behavior as the noiseless true solution.

The Tikhonov method to solve ill-posed problems was first proposed by Andrey N. Tikhonov back in 1963. The original paper is, however, in Russian and appears not to be available to the public, so instead of citing the original paper Petrov's application of the Tikhonov method toward spectral line deconvolution is cited [81]. The Tikhonov method when applied to spectral deconvolution treats the broadening function as an n-by-n matrix that is multiplied by the source function column array (what we are trying to extract) to form the raw data column array. Here, n is the domain size of discretization (i.e. how many discrete points are in the domain). Equation (5-3) shows the mathematical representation of the convolution operation in the Tikhonov treatment,

$$a(\upsilon) \otimes b(\upsilon) = c(\upsilon) \Leftrightarrow \mathbf{A} \cdot \underline{b} = \underline{c} \tag{5-3}$$

where **A** is the broadening function in a square matrix form, $\underline{b}$ is the source function in column array form, and $\underline{c}$ is the raw data in column array form. Note that **A**, $\underline{b}$, and $\underline{c}$ are generated from discretized forms of the original a(υ), b(υ), and c(υ) functions, and have the sizes n-by-n, n-by-1, and n-by-1, respectively. The multiplication between **A** and $\underline{b}$ is a simple matrix multiplication. In this formulation, the deconvolution operation is represented in equation (5-4),

$$b(\upsilon) = c(\upsilon) \varnothing a(\upsilon) \Leftrightarrow \underline{b} = \mathbf{A}^{-1}\underline{c} \tag{5-4}$$

where $\mathbf{A}^{-1}$ is the matrix inverse of **A**. Just as in the Fourier representation of the deconvolution operation, the matrix representation shown here is also unstable, i.e. highly sensitive to noise. To stabilize the solution, Tikhonov begins by multiplying both sides of the matrix equation by the transpose of the broadening function matrix as follows,

$$\mathbf{A}^{T}\mathbf{A}\underline{b} = \mathbf{A}^{T}\underline{c} \tag{5-5}$$

where $\mathbf{A}^{T}$ is the transpose of **A**. Then a stabilizing function is added to the left-hand side of the equation as follows,

$$\mathbf{A}^T\mathbf{A}\underline{b} + \alpha\mathbf{I}\underline{b} = \mathbf{A}^T\underline{c} \tag{5-6}$$

where **I** is the identify matrix matching in size to **A**, and α is called the regularization parameter and is a constant. For a given set of raw noisy data, $\underline{c}$, equation (5-6) yields a single solution for $\underline{b}$. Increasing α decreases the sensitivity of the solution to the noise in the data. When compared to the Fourier method, the Tikhonov approach does not introduce filter-related broadening. However, the lineshape can still be artificially broadened due to over-smoothing. So α should be chosen to be as low as possible while still keeping the noise in the extracted lineshape lower than the noise in the original raw data.

The main advantage of the Tikhonov regularization method is that it provides excellent fidelity but in return requires a large amount of computational resource. The computational requirement for the Tikhonov method is driven by the need to perform matrix inversion. Test runs on a Matlab implementation of the method shows that the needed computational time scales as $n^{8/3}$, where n is the discretized domain size. This test was carried out using a decent home computer built in 2007. The LIF data on which the test is carried out is the radial trace for the 300 V, 30 mg/s condition at Z = -0.1, R = 0.033. Figure 5-16 shows the computation time plotted against the domain size on a log-log scale along with a power-law curve-fit. The largest n tested in this study was 4501 (n has to be an odd number). n = 6001 was tested but the computer ran out of memory (3 GB of RAM) and could not complete the matrix inversion.

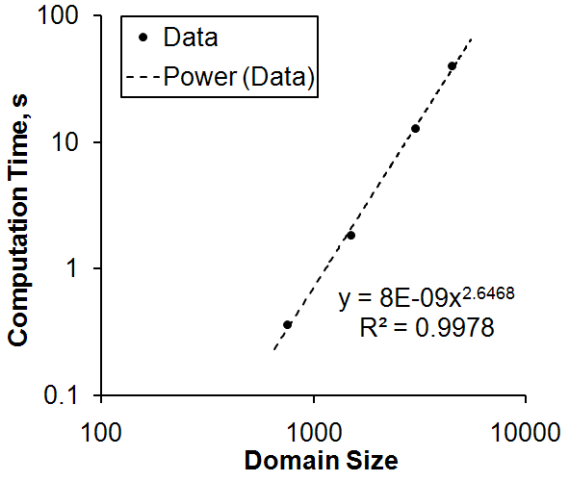

Figure 5-16. Computation time study for the Tikhonov method applied to spectral deconvolution.

5.3.3.3 Comparison of Deconvolution Methods

From the description of the Fourier and Tikhonov approaches to deconvolution it is clear that each method is better suited for different situations. Ideally, the Tikhonov method would be used for all deconvolution operations, but, as it turns out, certain

deconvolution problems are computationally too expensive when using the Tikhonov approach. In particular, any lineshape that contains a large number of hyperfine splittings also requires a large domain size to accurately capture. For example, accurate representation of the hyperfine structure of the Xe I 834.912 nm transition requires $n \approx 30000$ for a typical Hall thruster LIF scan. From the computational resource study presented earlier, in the section on the Tikhonov method, it is clear that such a matrix inversion is computational not feasible on an average computer. For such traces, if it is also true that the Doppler broadening is at least several times wider than the inverse Gaussian filter needed to obtain a stable solution, then the Fourier transform method is preferable over the Tikhonov method.

In terms of calculating the mean velocity, there is actually very little difference between using the Fourier approach versus the Tikhonov approach. This is to be expected since the advantage of the Tikhonov approach lies in reduced artificial broadening, which has little influence on the mean velocity.

All Xe II 835.0 nm LIF data presented in this dissertation uses the Tikhonov approach to perform deconvolution. Recall that the hyperfine structure of this transition is approximated as a 450-MHz Gaussian due to the lack of established hyperfine structure constants. This approximation has the side benefit of requiring a much lower domain resolution than other xenon transitions (with known hyperfine structure constants) to represent. Typical domain size needed is ~1500 so that a typical deconvolution run takes ~2 s to carry out. The use of the Tikhonov approach is also important because parts of the plasma upstream of the acceleration zone has VDFs that are about 1 GHz wide, and a 500-MHz inverse Gaussian filter is typically necessary to stabilize the solution of the Fourier approach. A 500-MHz Gaussian adds about 12% broadening to a 1-GHz VDF, which is a small but not negligible amount.

The value of the regularization parameter, α, is allowed to take on the values of 200, 400, 800, and 1600. The value for each scan is chosen such that the lowest value that makes the noise in the extracted VDF lower than the noise in the raw data is chosen. In other words, the SNR of the extracted VDF must be lower than the SNR of the raw data. The values of α presented may seem to span a large range but in reality the solution is not

that sensitive to small changes in α. Typically, one order of magnitude increase in α is equivalent to halving the amplitude of the noise in the VDF.

Any deconvolution of the Xe I 834.9 nm LIF data is handled using Fourier transform with inverse Gaussian filtering. Since the neutrals have relatively narrow VDFs, the extracted VDF results will be slightly broadened. Figure 5-17 shows a representative example of VDF extraction from raw Xe I 834.9 nm LIF data. Although the resulting VDF (shown in red line) has a slight unphysical dip on the left shoulder, it does not contain any hyperfine structure humps. This indicates the deconvolution process works the way it should.

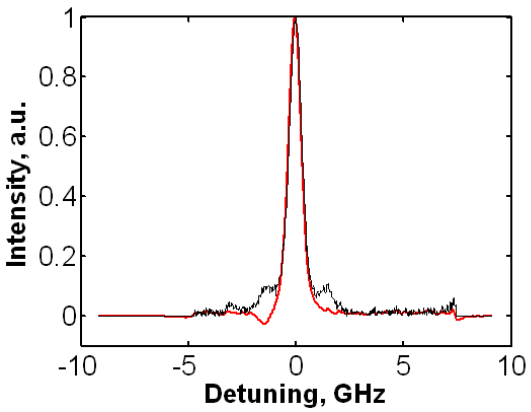

Figure 5-17. Example of applying Fourier deconvolution to a Xe I 834.9 nm LIF lineshape.

As it will turn out, no highly energetic neutrals were found as a part of this dissertation so there is no need to extract neutral xenon VDFs as the raw lineshape is sufficient for calculating mean velocity. Nevertheless, the experimental setup and analytical tools needed to study neutral xenon have been demonstrated as a part of this dissertation. Further detail on the use of Fourier transform with inverse Gaussian filtering to study the Xe I 834.9 nm transition in the environment of a Hall thruster can be found in [57].

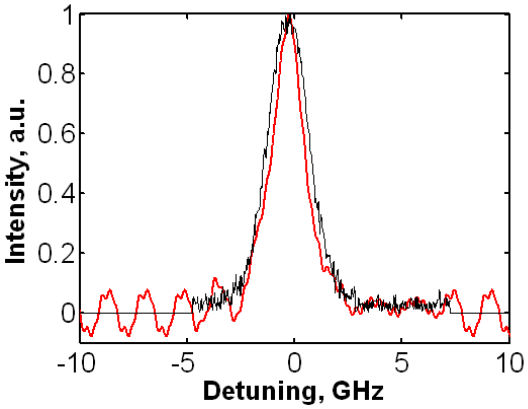

Figure 5-18. Example of problem associated with using the SLZ model with Fourier deconvolution. *Magnetic field strength is 320 G and Gaussian filter is 527 MHz wide.*

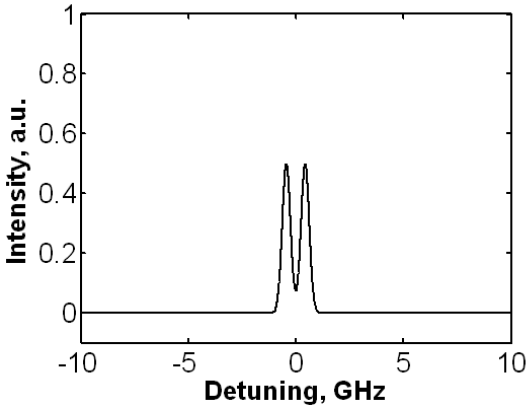

Figure 5-19. SLZ model approximation of the hyperfine structure lineshape for a magnetic field strength of 320 Gauss.

5.3.3.4 Modification to the SLZ Model for Deconvolution Purposes

While processing the off-radial Xe II 835.0 nm traces, an unforeseen problem was encountered. Recall that the off-radial traces have σ-component Zeeman broadening and that the original plan was to use the simple linear Zeeman (SLZ) model to approximate the hyperfine structure lineshape. It turns out the fact that the SLZ model has two distinct humps causes higher than normal instability in the solution for both the Fourier and the Tikhonov approaches.

Figure 5-18 shows what happens when the Fourier approach with a 527-MHz Gaussian filter is used to perform deconvolution on a trace at a location where the simulated magnetic field strength is 320 Gauss. Figure 5-19 shows the corresponding hyperfine structure lineshape produced by the SLZ model. Figure 5-18 shows that extra ringing oscillation shows up in the extracted VDF, which is the red line. It is possible to increase the width of the Gaussian to decrease the amount of noise in the VDF but doing so will introduce additional artificial broadening. The added artificial broadening effectively negates the purpose of modeling the Zeeman splitting, which is to remove Zeeman-effect-related broadening.

When the same analysis was performed with the Tikhonov approach, a different problem arose. For reasons unknown, the FWHM frequency of the extracted VDF sometimes increases with increasing SLZ model magnetic field strength. This is

unphysical since if the hyperfine structure being removed from the raw trace is increasing in width, the extracted solution should be decreasing in width. This problem may have something to do with the choice of the stabilization function $\alpha\mathbf{I}$. Normally, the stabilization function is picked based on the type of ill-posed problem being tackled. However, a survey of the literature did not turn up any specialized stabilization function for spectral deconvolution so the original default function proposed by Tikhonov was used. In any event, the author by chance tried a different representation of the SLZ model and found that it gave the right trends though loses out on certain lineshape details in the process.

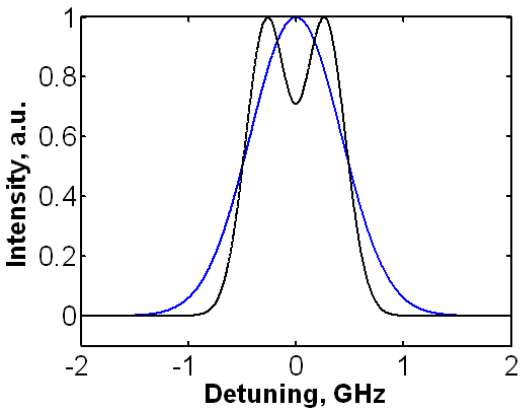

Figure 5-20. Comparison of the SLZ and the simple Gaussian model for representing the Zeeman-split hyperfine structure of the Xe II 835.0 nm transition at 200 G.

The modification to the SLZ model is to approximate the hyperfine structure as a single Gaussian function with the same FWHM frequency as the original SLZ model. This model is referred to as the simple Gaussian model. Figure 5-20 compares the approximate lineshapes generated by the SLZ model and the simple Gaussian model for a magnetic field strength of 200 G.

The simple Gaussian model has a number of advantages. One, the Gaussian produces less instability in the solution than most other functions of equivalent FWHM frequency and thus requires lower α (or, in the case of the Fourier transform approach, narrower Gaussian filter) to stabilize. Two, it still correctly models the approximate width and, by extension, the amount of broadening introduced by the Xe II 835.0 nm hyperfine structures. In exchange, the hyperfine structure lineshape no longer has two peaks but only one. So if the Doppler broadening in the sample is about the same in width as or lower in width than the Zeeman-split hyperfine structures, the resulting distortion to the lineshape will not be removed. Luckily, the Doppler broadening is at least 2.5 times wider than the hyperfine structure broadening everywhere in the off-radial Xe II 835.0 nm data set collected during this dissertation.

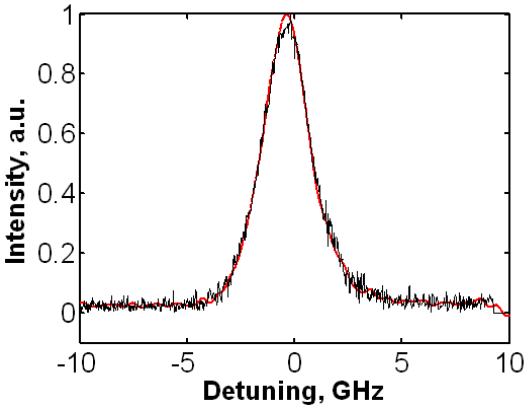

Figure 5-21. Example of extracted VDF using the Tikhonov approach with the SLZ model.

Magnetic field strength is 261 G.

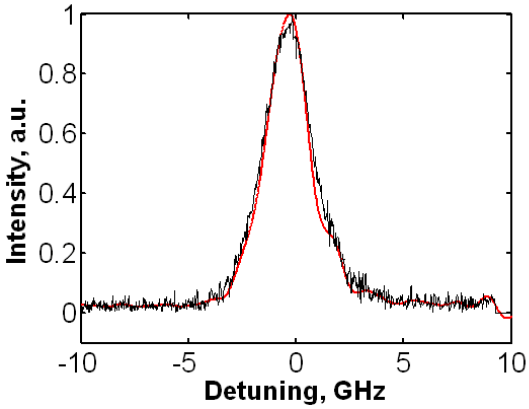

Figure 5-22. Example of extracted VDF using the Tikhonov approach with the simple Gaussian model.

Magnetic field strength is 261 G.

Figures 5-21 and 5-22 show the extracted VDF using the Tikhonov approach for deconvolution with the SLZ model and the simple Gaussian model, respectively. The data is from the off-radial LIF trace at Z = -0.05, R = 0.033 for the 300 V, 20 mg/s condition. The local magnetic field strength is 261 G, the SNR is 39, and the raw data has an FWHM frequency of 2.58 GHz. The extracted VDF using the SLZ model has an FWHM frequency of 2.52 GHz while the same using the simple Gaussian model has an FWHM frequency of 2.27 GHz. As can be seen from these figures, the actual difference between using the SLZ versus the simple Gaussian model is very small. Nevertheless, deconvolution tests with multiple scans across multiple operating conditions show that using the SLZ model with the Tikhonov method sometimes gave an FWHM frequency that is too large to be physical. So, to avoid this unknown artificial broadening, the simple Gaussian model is adopted for analyzing the off-radial Xe II 835.0 nm traces.

5.3.4 Calculating the Bulk Quantities

The calculation of the bulk quantities can actually be a little tricky for LIF traces that contain a lot of noise. Note that bulk refers to the average property of the species being studied. Figure 5-23 illustrates the approach to calculating bulk quantities. This figure is outputted by a custom Matlab implementation of the Tikhonov approach to deconvolution. The red line represents the extracted VDF. The black dots in this figure indicate data points. It is important to note that there are regions (the left and right edges)

in the trace where even though there is no signal the "noise floor" does not average out to zero. This is an artifact of how the lock-in amplifiers reported the data. The phase-lock-in amplifier compares the raw signal to a reference sine wave that has the same frequency as the corresponding chopper frequency. However, there are always two components to this signal, the component in-phase with the sine wave and the component that is 90° out-of-phase with the sine wave. Ideally, the lock-in amplifier would report only the in-phase component, but this approach to data extraction only works when the phase lag between the lock-in amplifier and the chopper is known. Unfortunately, since the chopper and the lock-in amplifier drifts over time, the phase difference is also drifting. It was decided early on in the test that keeping track of the phase difference for two laser beams added too much complexity to an already complex experiment. Instead, the lock-in amplifiers are instructed to report the magnitude of the signal at the chopping frequency. This is equivalent to taking a square-root of the sum of the squares of the in-phase and out-of-phase components. In this mode, phase is not an issue, but the signal-to-noise ratio goes down by a factor of $\sqrt{2}$ compared to in-phase only lock-in amplification. The reason for this drop in SNR stems from the fact that the noise is now the square-root of sum of square of the noise in the in-phase and out-of-phase components. As a result, the data can never go negative, and the noise seen in the raw trace is always the magnitude of the noise.

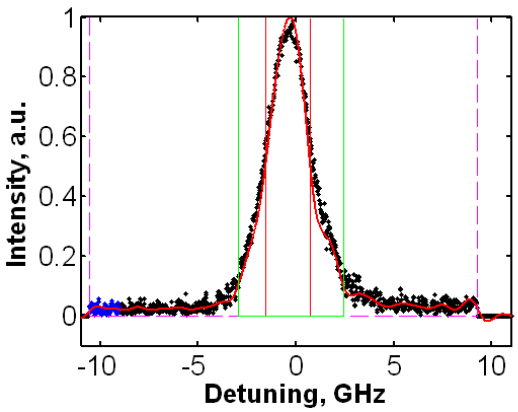

Figure 5-23. An example of how bulk quantities are extracted from the VDF.

The existence of this "noise floor" means that one cannot simply perform an intensity-average across the entire trace to find the average velocity. Such an average would depend on the exact frequency range chosen to average over since the noise floor averages to a non-zero value. To get around this problem, the analysis program finds the frequency range where the bulk signal resides by finding all data with intensities that are above a certain percentage of the maximum intensity. This percentage, called the signal threshold, is allowed to vary from 10% to 40% in increments of 10%. For a given LIF

trace, the lowest signal threshold that exceeds the "noise floor" is used. In Figure 5-23, the signal threshold is 10% and the exact location of the boundaries for what is considered signal is marked in green vertical lines. This approach to calculating the average velocity has the side effect of neglecting low-intensity tail populations. As it will turn out, most of these tail populations are on the low energy side. So the use of signal threshold produces bulk quantities that correlate better with the high energy populations, which are responsible for erosion.

For the LIF data, SNR is defined as the peak signal divided by the root-mean-square of noise. A study of the extracted data after post-processing revealed that, typically, a signal threshold of 10% is used for traces with SNR >20, 20% is used for traces with 10< SNR <20, 30 and 40% are used for traces with SNR <10. Section 5.3.6 will give a more detailed description of SNR as well as show traces at various levels of SNR and what type of information can be accurately extracted at each level.

To estimate the width of the extracted VDF, the FWHM velocity of the extracted VDF is calculated. The exact locations where the boundaries of the FWHM velocity calculations are shown in Figure 5-23 as red vertical lines. Also shown in Figure 5-23 is a pair of pink dashed vertical lines on the left and right edges of the plot. These represent the boundaries of the raw data. Since deconvolution typically requires a calculation domain that is larger than the domain of the raw data, the extracted VDF contains artificial data beyond these pink boundaries. The VDF exported from the Matlab program is cut off at these boundaries.

Once the bulk velocity is extracted from both the axial and the off-radial traces, geometric calculations are applied to calculate the bulk radial velocity and the average angle of motion with respect to the thruster firing axis. The travel angle is then combined with profilometry data to calculate the angle of incidence before the singly-charged ions enter the sheath. For neutral xenon, this angle of incidence is the angle with which they bombard the walls because they are not affected by the sheath.

5.3.5 Approximating the Angular Energy Distribution Function

Composing the angular energy distribution function from two one-axis VDFs requires a number of assumptions and approximations. The first assumption is that the

plasma is azimuthally uniform, and the azimuthal velocity component of the plasma carries negligible amount of energy in comparison to the axial and radial components. The azimuthal uniformity of the 6-kW thruster (which has a center-mounted cathode) is shown to be excellent in Reid's dissertation work [46]. Smith's dissertation work included an LIF study of the azimuthal component of velocity for a P5 Hall thruster, which is similar in power level and magnetic field strength to the 6-kW Hall thruster. He found that the azimuthal component of velocity is no greater than 1000 m/s at 100 mm downstream of the exit plant [80]. Comparing to the velocity component in the axial-radial plane, which reaches 5000+ m/s at the exit plane, we can conclude that the azimuthal component of velocity does not carry significant bombardment energy. (Recall energy scales as velocity squared so 1000 m/s is 0.7 eV while 5000 m/s is 17 eV)

The second assumption is that the axial and radial components of velocity are independent of each other. Or in other words, the angular energy distribution function is separable into an axial component and a radial component. For the singly-charged xenon in the LIF studies, this assumption holds true so long as the net force applied is purely axial. The main force acting on the singly-charged xenon is the electric force. From the bulk velocity results that will be shown later, it can be inferred that the electric field is largely axial. However, a non-trivial radial component of velocity is observed at select near-wall locations suggesting that the forces acting on the ions are not purely axial at those locations. For the ions found at these locations, the axial and radial components of velocity are not un-correlated and the constructed angular energy distribution functions will contain some errors.

To further simplify the results, a third assumption is made. The third assumption is that the singly-charged xenon population can be modeled as a Maxwell-Boltzmann distribution with different axial and radial temperatures. The reason for using two different temperatures is because the VDF results often show very different spread in the axial and radial directions. This is speculated to be due to a combination of difference in forces between the two directions and the fact that the ions in the acceleration zone are in the free-molecular flow regime. Equation (5-7) shows the mathematical form of this distribution.

$$G(u,w) = A\exp\left[-\frac{m(u-u_0)^2}{2kT_Z} - \frac{m(w-w_0)^2}{2kT_R}\right] \tag{5-7}$$

where A is a normalizing constant, m is the mass of xenon, u is the axial velocity, u_0 is the bulk axial velocity, w is the radial velocity, w_0 is the bulk radial velocity, k is the Boltzmann constant, T_Z is the axial temperature, and T_R is the radial temperature. The two temperatures are found by using the FWHM velocities from the axial and off-radial traces. In fact, temperatures are not physically meaningful for a non-equilibrium plasma like that found in the acceleration zone of the Hall thruster. So for the most part, this dissertation will focus on FWHM velocity instead. Next, the mathematical derivation for the formulas needed to calculate the radial FWHM velocity from the axial and off-radial FWHM velocities is shown.

We begin by first inserting the mathematical relationship between the FWHM velocity and the corresponding temperature, which is shown in equation (5-8).

$$\Delta u = \sqrt{8\ln 2\frac{kT_Z}{m}} \tag{5-8}$$

Substituting equation (5-8) into equation (5-7) and setting the value of G to 0.5*A yields the following relationship,

$$1 = \frac{u^2}{\left(\frac{1}{2}\Delta u\right)^2} + \frac{w^2}{\left(\frac{1}{2}\Delta w\right)^2} \tag{5-9}$$

where Δu and Δw are the axial and radial FWHM velocities, respectively. Equation (5-9) is an ellipse. This equation can be used to relate the off-radial FWHM velocity to the axial and radial FWHM velocities.

Figure 5-24 shows a diagram of an ellipse in velocity space. This diagram represents the FWHM velocity of the Maxwell-Boltzmann distribution with different axial and radial temperatures. The variable q represents the off-radial axis at which LIF data was taken and Δq is the FWHM velocity of the off-radial VDF. In this diagram, we know the values of Δu, Δq, and θ, and we wish to calculate the value of Δw. Note, that a blue dot marks the location of one of the intersections between the q-axis and the FWHM velocity ellipse. This dot has a magnitude of R_θ and is located at the coordinates (u_θ, w_θ).

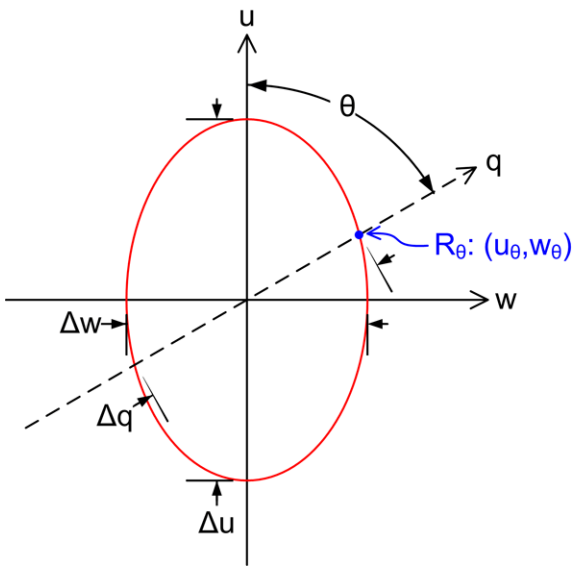

Figure 5-24. Diagram of the FWHM velocity ellipse for a two-temperature Maxwell-Boltzmann distribution in velocity space.

We begin the calculation by using trigonometry to relate u_θ, w_θ, and θ, shown in equation (5-10).

$$u_\theta = \frac{w_\theta}{\tan\theta} \tag{5-10}$$

Substituting equation (5-10) into equation (5-9) and solving for w_θ yields equation (5-11).

$$w_\theta^{\,2} = \frac{1}{\frac{1}{\left(\frac{1}{2}\Delta u \tan\theta\right)^2} + \frac{1}{\left(\frac{1}{2}\Delta w\right)^2}} \tag{5-11}$$

Next, we use another trigonometric relationship this time relating w_θ, R_θ, and θ, shown in equation (5-12).

$$w_\theta = R_\theta \sin\theta \tag{5-12}$$

Square both sides of equation (5-12), substitute in $R_\theta = (1/2)*\Delta q$, substitute the result into equation (5-11), and then simplifying yields a relationship between the various FWHM velocities and θ. The end equation is shown in (5-13).

$$\frac{1}{(\Delta w)^2} = \frac{1}{(\Delta q \sin\theta)^2} - \frac{1}{(\Delta u \tan\theta)^2} \tag{5-13}$$

There is a problematic limitation with equation (5-13). Since $\theta = 60°$ for the two-axis near-wall study, this equation cannot be solved if $\Delta q \geq 2\Delta u$. Mathematically, it is not possible to draw such an ellipse as the boundaries of the ellipse along the w-axis will go to infinity. However, particularly for the 600 V, 10 mg/s operating condition, there are several data locations where the off-radial FWHM velocity, Δq, is quite a bit larger than the axial FWHM velocity, Δu. There are several possible physical explanations for the discrepancy. First, even if the VDF is close to a two-temperature Maxwell-Boltzmann distribution, it is possible that the directions of the two temperatures do not line up with the axial and radial direction thereby invalidating equation (5-13). Second, it is possible for the real VDF to have multiple peaks, which is not very well captured by the two-temperature Maxwell-Boltzmann distribution. Third, the VDF may not be anything like a Maxwell-Boltzmann distribution. The two-temperature Maxwell-Boltzmann fit is really only meant to capture the bulk velocity and spread of the VDF in two directions.

To make sure the value of Δw does not become unrealistically large when equation (5-13) is applied to data where Δq is much larger than Δu, an engineering capping equation is used to limit the value that Δw can take on. This capping equation is shown in (5-14).

$$\Delta w = \Delta q + \Delta u, \; \Delta q > 1.68 * \Delta u \tag{5-14}$$

This capping equation basically says that the real spread in the radial VDF cannot realistically be greater than the sum of the spread in the off-radial VDF and the axial VDF. There is no physical basis for this cap; it is just an engineering equation. The criterion for this cap becoming active is set to $\Delta q > 1.68*\Delta u$ because when $\Delta q = 1.68*\Delta u$ for $\theta = 60°$, equations (5-13) and (5-14) give the same answer. Any radial FWHM velocity that is calculated from equation (5-14) instead of equation (5-13) will be specially marked.

5.3.6 Signal-to-Noise Ratio and Uncertainty

Signal-to-noise ratio (SNR) is defined in this dissertation as peak signal strength divided by root-mean-square of noise. It turns out the SNR is very easy to calculate when the lock-in amplifier is set to report the signal magnitude. The recorded intensity is the root-mean-square of the in-phase and out-of-phase components. In parts of the LIF trace that contain only noise, both components are noise and the recorded intensity is equal to the root-mean-square of noise. Thus, to get an accurate estimate of the amount of noise, one simply averages over a part of the LIF trace where there should be only noise. Figure 5-25 shows an example of how the SNR is extracted from the raw LIF trace. The blue dots, highlighted in a blue circle, are the data points used to compute the SNR of this particular trace.

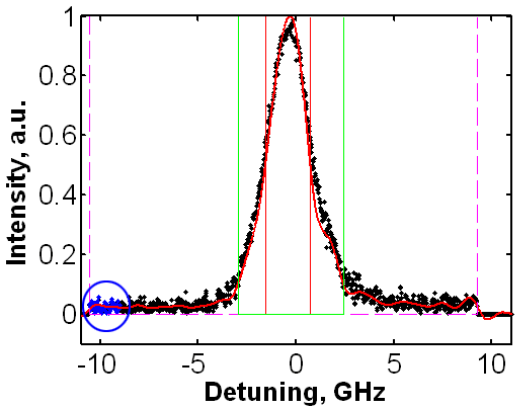

Figure 5-25. An example of how the SNR is extracted.
The SNR is 39.

To simplify matters even more, the SNR calculation is carried out after first normalizing the entire trace by the peak intensity. Thus the peak signal strength is always 1 and SNR can be calculated using equation (5-15).

$$\mathrm{SNR} = \frac{1}{\mathrm{Noise}_{\mathrm{RMS}}} \tag{5-15}$$

Figures 5-26, 5-27, and 5-28 show example LIF traces with SNR of roughly 5, 10, and 20, respectively. From these traces we can see that the VDF shows up very cleanly for SNR > 20 while it is decent for SNR > 10. The VDF is not clear for SNR of 5 and a lot of details are buried by the noise. For the purpose of this dissertation only VDFs from raw traces with SNR > 10 are shown. However, bulk velocity and temperature extraction is done for all scans with SNR > 5. The rationale is that even though the SNR is low, we are extraction only two parameters using hundreds of data points so the accuracy of the results should be decent.

The velocity resolution of the extracted VDF is ~25 m/s and is set by the piezo-voltage step size. The corresponding resolution in frequency space is ~30 MHz, well above the laser linewidth of ~10 MHz. The accuracy of the extracted bulk velocity is limited by the accuracy with which the stationary reference is calculated and the SNR in the LIF scan. In practice, the SNR in the optogalvanic scan is generally higher than the SNR in the LIF scan so the noise in the LIF scan drives the overall accuracy. For an SNR of ~5, the relative uncertainty is ~6%. For an SNR of ~10, the relative uncertainty is ~4%. These numbers are calculated by taking the average relative uncertainty of each data point that is a part of the raw trace (can be found by taking the inverse of the SNR and divided by two), finding the typical number of data points that makes up the signal part of each trace (~100), and then applying Student's t-distribution statistics. All reported uncertainties are with a 95% confidence-interval. On top of these relative uncertainties, the Matlab program performs intensity-weighted averaging to find the stationary reference marker and is determined to be accurate to ~60 MHz, or ~50 m/s. So the total uncertainty in the bulk velocity for SNR 5-10 is determined to be ±50 m/s or ±6%, whichever is greater. For SNR 10+, the uncertainty is ±50 m/s or ±4%, whichever is greater.

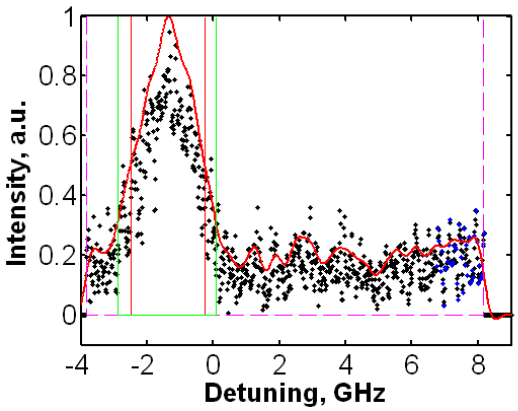

Figure 5-26. An example LIF trace with an SNR of ~5.

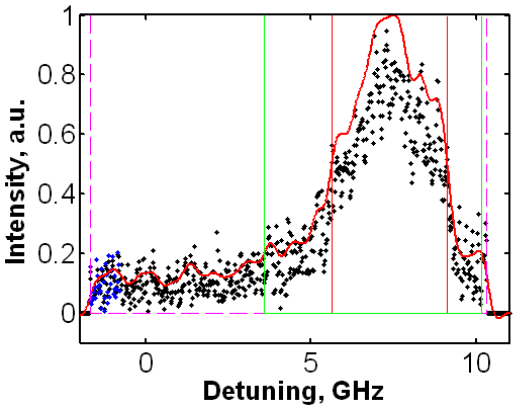

Figure 5-27. An example LIF trace with an SNR of ~10.

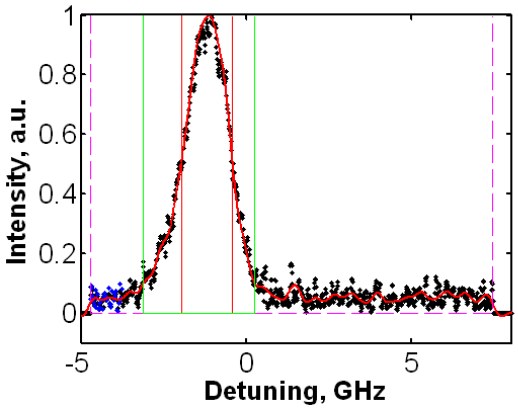

Figure 5-28. An example LIF trace with an SNR of ~20.

Individual LIF traces may, and often do, have lower uncertainty than what is specified here. However, they will not have a higher uncertainty unless otherwise specified.

The uncertainty in the angles calculated from the bulk velocities can be calculated using standard error propagation theories. If the SNR for both the axial and the off-radial trace are 5-10, the uncertainty in the associated angle is ±5°. If the SNR for both the axial and the off-radial trace are 10+, the uncertainty in the associated angle is ±3°.

The uncertainty in the extracted FWHM velocity is tricky to quantify. Since the FWHM velocity is calculated after the deconvolution operation, it is, statistically speaking, a result of averaging a fair number of data points. The problem lies in determining how many data points are effectively contributing to this calculation. Since application of the standard error propagation theory to deconvolution operation is outside the author's expertise, a simple assumption will be made instead. Each of the two boundary locations that contribute to the calculation of the FWHM will be assumed to use half of the typical number of data points that make up the signal. Since the FWHM velocity is the difference between these two locations in velocity space, both locations contribute to the total uncertainty. Applying Student's t-distribution statistics, the final result is that for SNR 5-10, the uncertainty in the extracted FWHM velocity is ±13%. For SNR 10+, the uncertainty in the extracted FWHM velocity is ±9%.

5.4 LIF Results for Singly-Charged Xenon

All results are shown on a non-dimensionalized spatial grid. The axial coordinate, Z, is non-dimensionalized by the channel length with Z = 0 at channel exit plane and Z = -1 at the anode exit plane. The radial coordinate, R, is non-dimensionalized by the channel width with R = 0 at the nominal inner wall radius and R = 1 at the nominal outer wall radius. All intensities are given in arbitrary units (abbreviated a.u.).

5.4.1 Channel Centerline Acceleration Zone Results

Originally, the goal of the one-axis acceleration zone study was to develop various techniques to prepare for the two-axis near-wall study. However, it was realized shortly after processing the data that the results had implications for channel wall erosion trends.

Figure 5-29 show a comparison of channel centerline axial LIF data converted to equivalent plasma potential compared to emissive probe data collected by Reid [46] while accounting for the cathode-to-ground voltage in the experimental setup. The two sets of data are in good agreement with the 300 V, 10 mg/s data showing the most disagreement. The cause of the disagreement is unknown as of writing. Figures 5-30 to 5-36 show the extracted axial VDFs along the channel centerline of the 6-kW Hall thruster operating at various operating conditions.

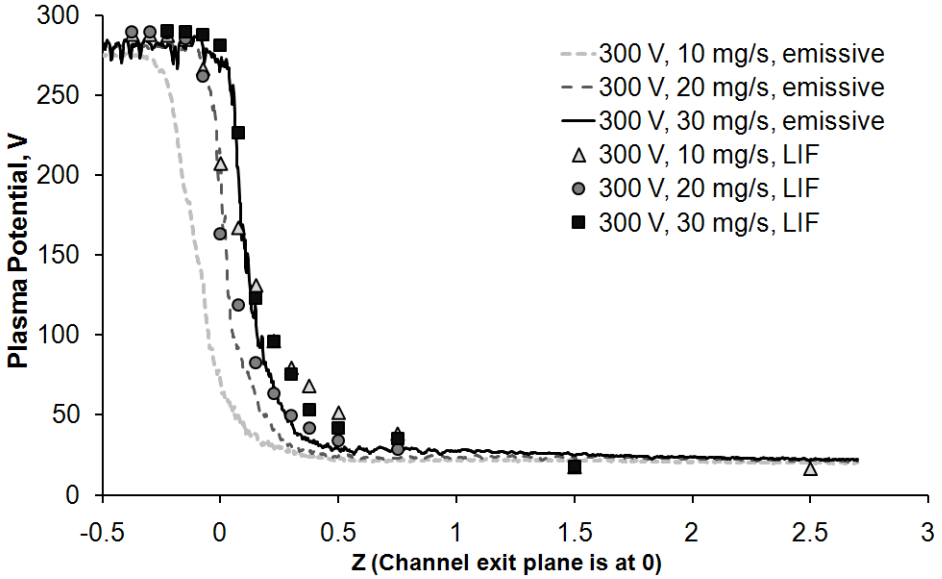

Figure 5-29. Validation of LIF data against emissive probe data [46].

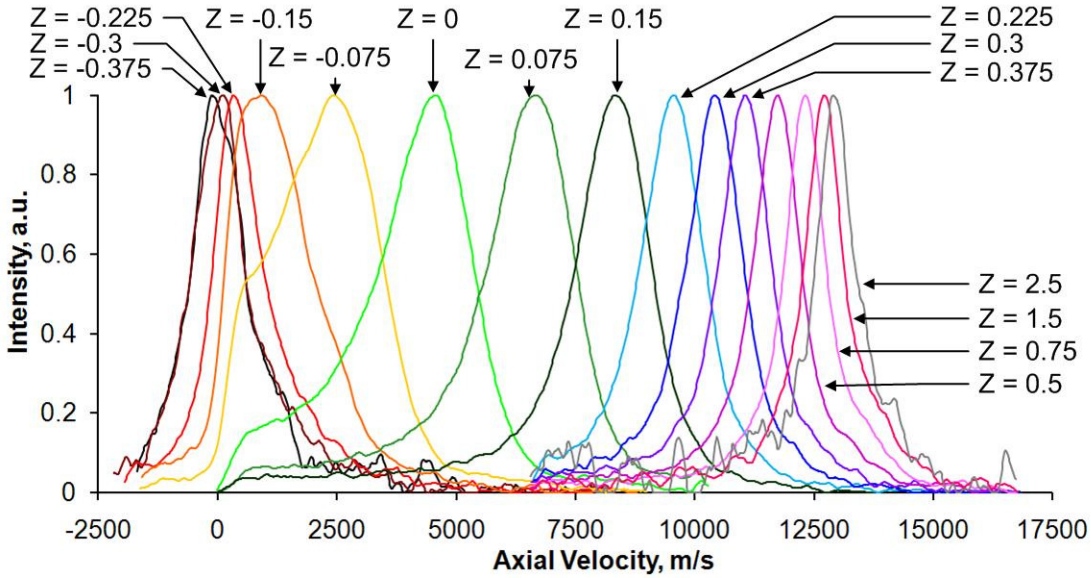

Figure 5-30. Axial VDFs along the channel centerline of the 6-kW Hall thruster operating at 150 V, 10 mg/s.

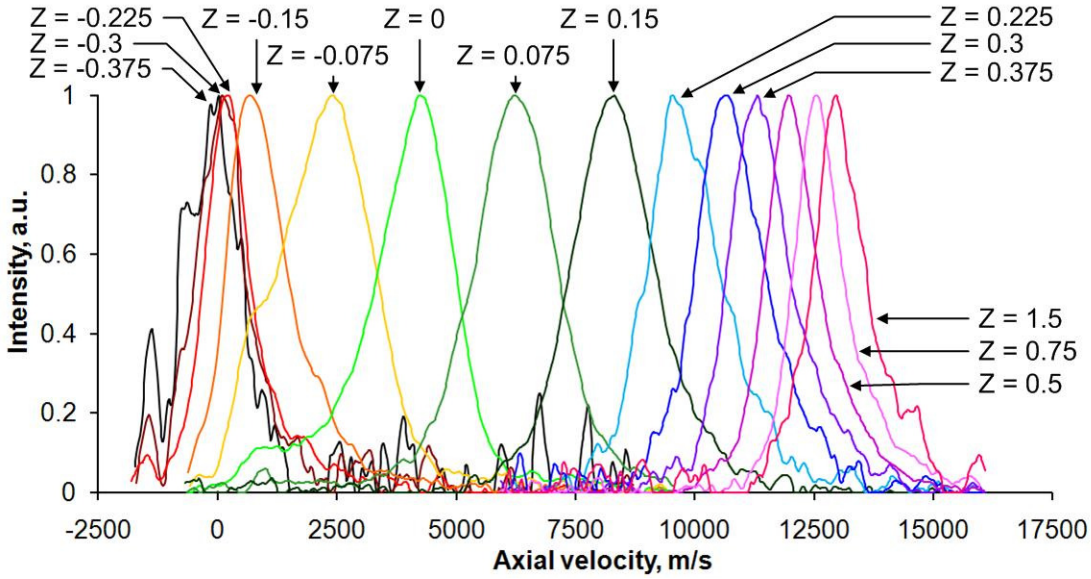

Figure 5-31. Axial VDFs along the channel centerline of the 6-kW Hall thruster operating at 150 V, 20 mg/s.

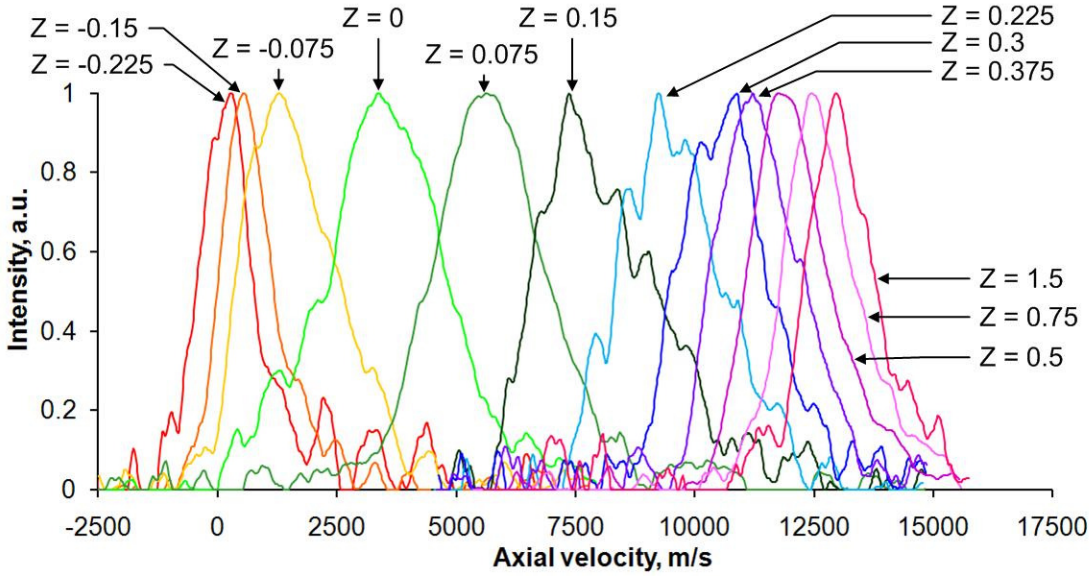

Figure 5-32. Axial VDFs along the channel centerline of the 6-kW Hall thruster operating at 150 V, 30 mg/s.

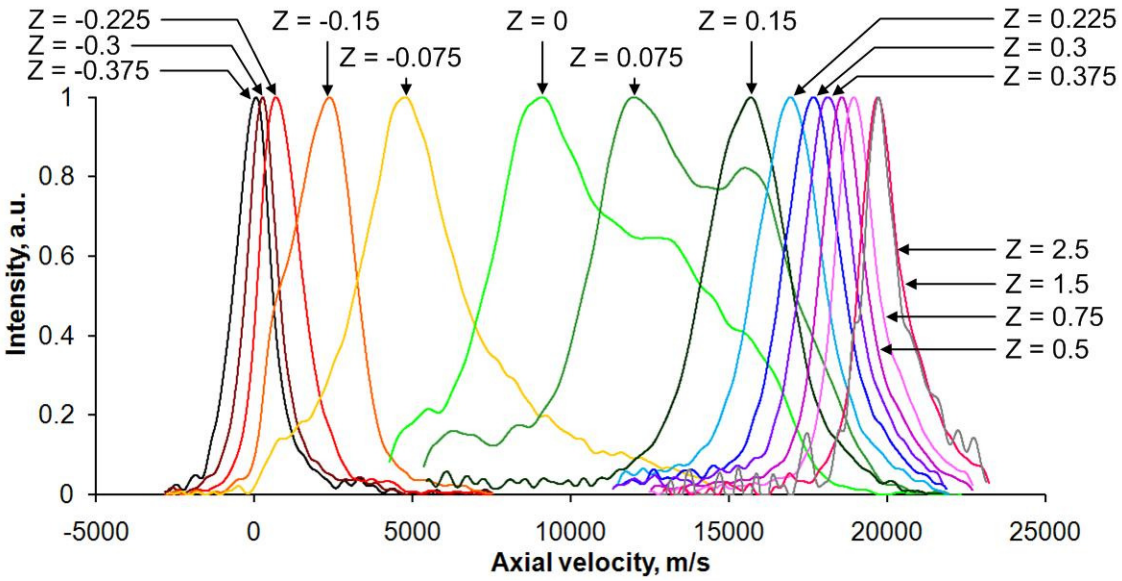

Figure 5-33. Axial VDFs along the channel centerline of the 6-kW Hall thruster operating at 300 V, 10 mg/s.

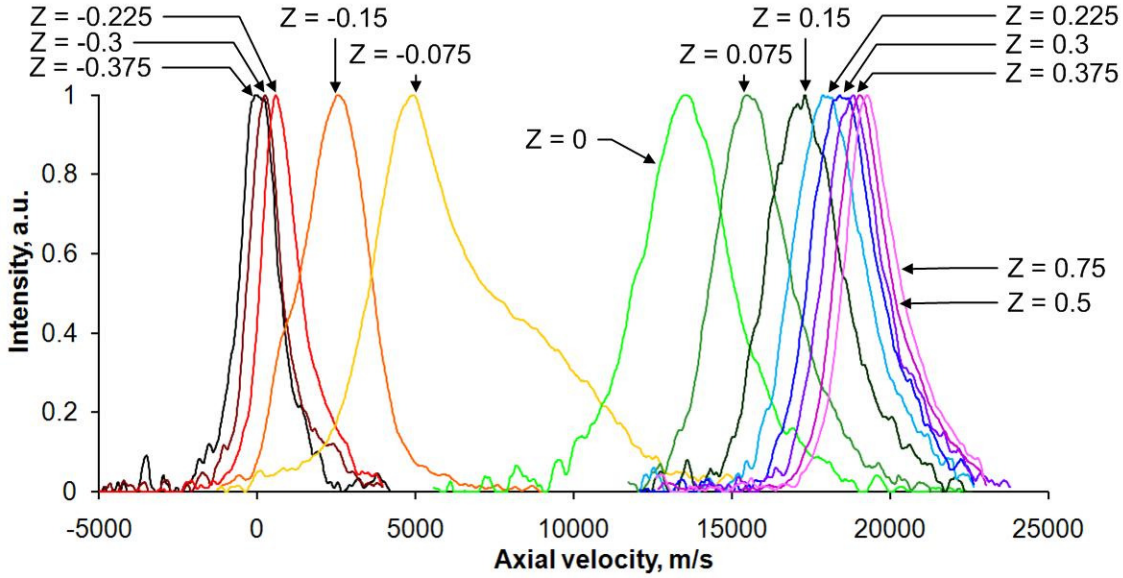

Figure 5-34. Axial VDFs along the channel centerline of the 6-kW Hall thruster operating at 300 V, 20 mg/s.

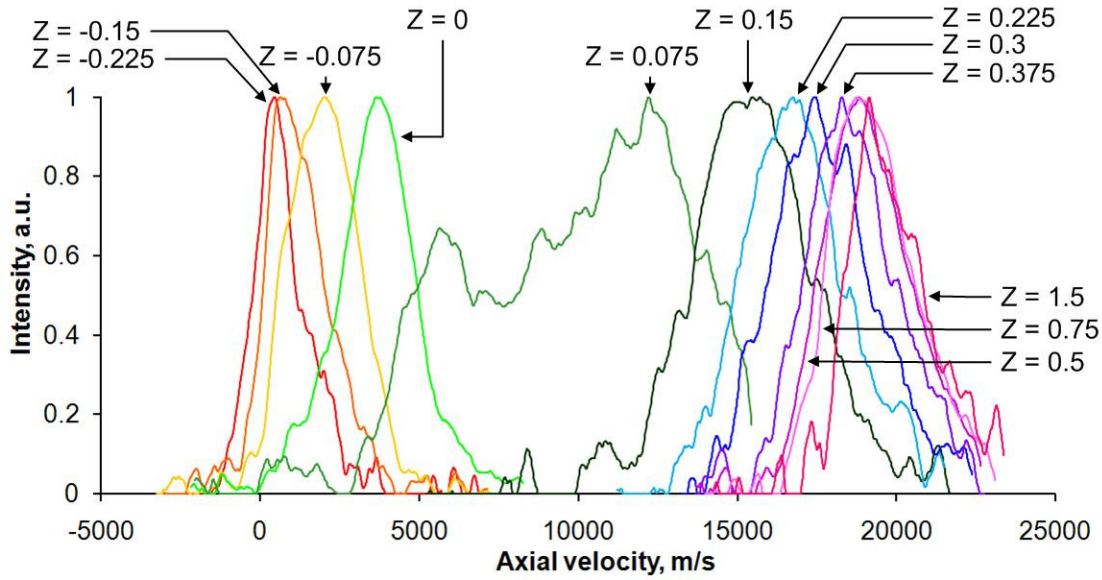

Figure 5-35. Axial VDFs along the channel centerline of the 6-kW Hall thruster operating at 300 V, 30 mg/s.

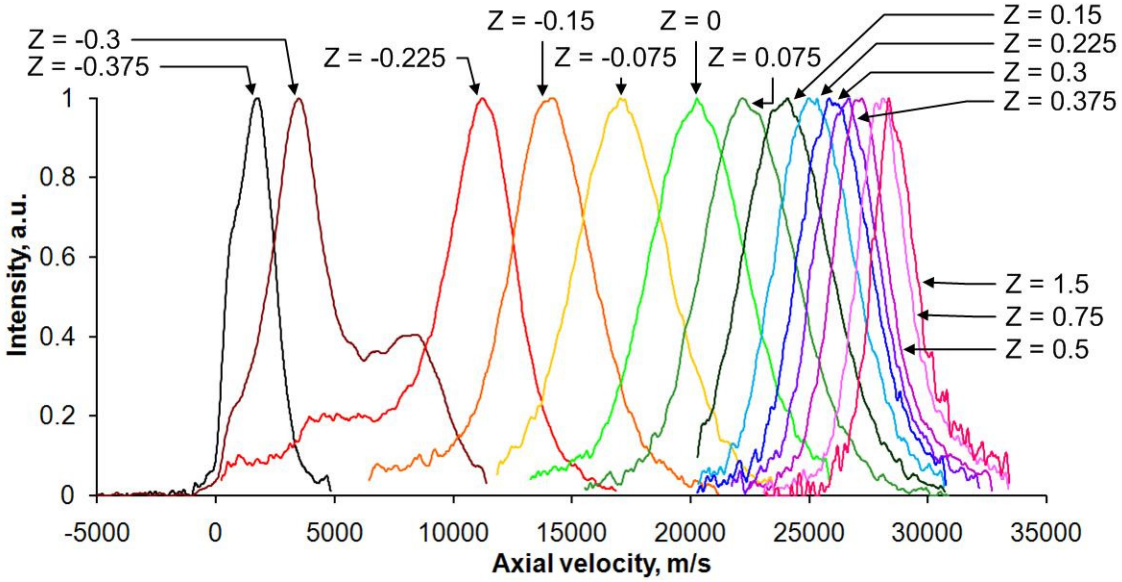

Figure 5-36. Axial VDFs along the channel centerline of the 6-kW Hall thruster operating at 600 V, 10 mg/s.

The VDFs appeared to be much wider inside of the acceleration zone than both upstream and downstream of the acceleration zone. This is speculated to be due to plasma oscillations. In particular, the breathing mode is the prime suspect. Rough calculations show that the effect of kinematic compression is too small to explain the rapid reduction in the width of the VDF post acceleration. The only plausible explanation is that the true instantaneous VDF is not that wide to begin with, and that the apparent width of the VDF in the acceleration zone is due to temporal variations in the instantaneous bulk velocity. Physically, one can think of the breathing mode as causing the whole acceleration zone to move back and forth in the axial direction [82]. Since the interrogation zone is fixed, the local plasma potential and associated bulk velocity is observed to fluctuate with time. Since the LIF experiments take time-averaged data, the VDF will appear smeared out and artificially broadened. More analysis on how the breathing mode can cause the VDF to broaden and even take on bi-modal appearance can be found in [55].

Figure 5-37 shows the bulk axial velocity as a function of normalized axial coordinate. Note that the SNR for these axial traces are higher than the ones obtained for the two-axis near-wall LIF most likely because the laser beam power is split into two probe beams for the two-axis test. The higher SNR also means lower uncertainty. Using

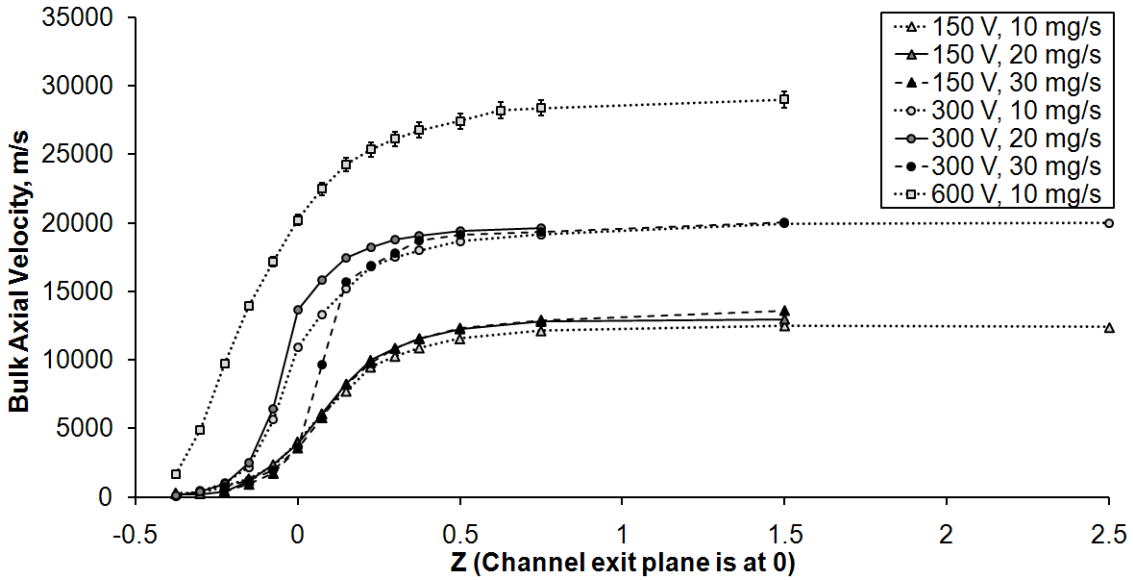

Figure 5-37. Bulk axial velocity as a function of the normalized axial coordinate along the channel centerline of the 6-kW Hall thruster.

One set of representative error bars are drawn on the 600 V, 10 mg/s data set.

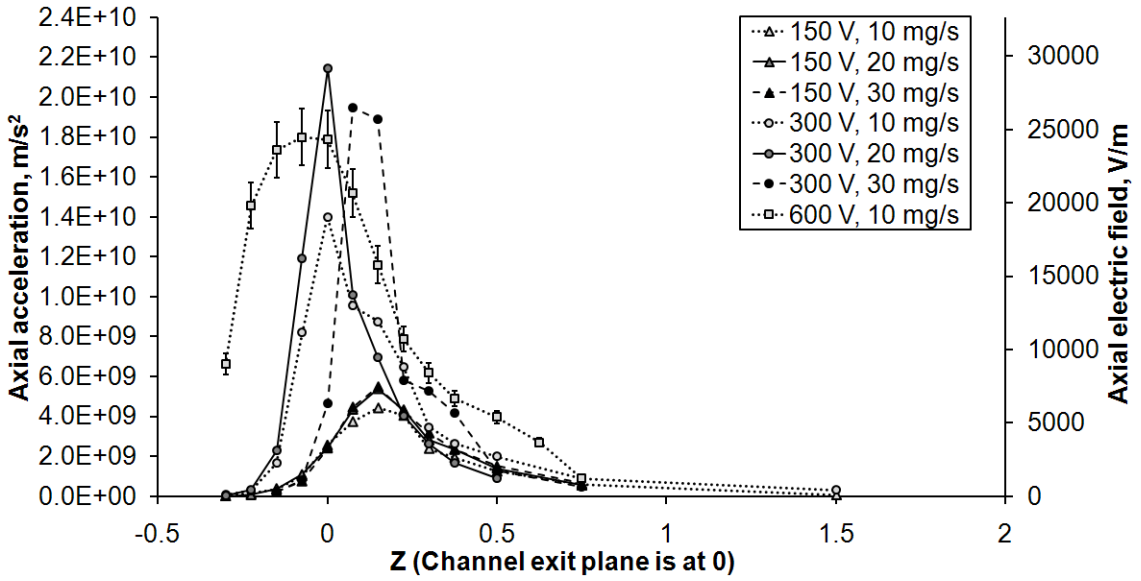

Figure 5-38. Axial acceleration as a function of the normalized axial coordinate along the channel centerline of the 6-kW Hall thruster.

One set of representative error bars are drawn on the 600 V, 10 mg/s data set.

Student's t-distribution statistical methods, the calculated uncertainty is ±50 m/s or ±2%, whichever is greater.

Figure 5-38 shows a plot of the axial acceleration as a function of position for all tested conditions. The acceleration is calculated by using the bulk axial velocity of the current data location with the bulk axial velocity at the next downstream data location along with the distance between the two locations. We can further assume the acceleration of the singly-charged ion is due entirely to the electric field and that charge exchange effects make negligible contributions to calculate the electric field value associated with the acceleration. The result is that the electric field scales with the acceleration and can be plotted as a second axis on the same figure. The associated uncertainty, accounting for variance in the bulk axial velocity and positioning accuracy of the axial motion stage, is calculated to be ±8%.

Table 5-2 lists the acceleration zone length, which is defined as the distance it takes for the bulk velocity to rise from 10% to 90% of the maximum bulk velocity. Also listed is the acceleration zone location, defined as the average of the 10%- and 90%-of-maximum bulk velocity locations. Lastly, the position where the peak acceleration occurs is listed. Non-dimensionalized values are normalized by the discharge channel length. In general, the peak acceleration position tends to be further into the channel than the

Table 5-2. Acceleration zone length and location at various operating conditions for the 6-kW Hall thruster.

ND in this table means the value is non-dimensionalized by the discharge channel length.

Discharge voltage, V	Anode mass flow rate, mg/s	Length of acceleration zone (ND)	Location of acceleration zone (ND)	Peak acceleration position (ND)
150	10	0.607	0.143	0.15
150	20	0.542	0.130	0.15
150	30	0.596	0.188	0.15
300	10	0.537	0.108	0
300	20	0.344	-0.005	0
300	30	0.395	0.123	0.075
600	10	0.643	-0.025	-0.075

location of the acceleration zone, suggesting that more of the potential drop occurs nearer upstream than downstream. The acceleration zone is more recessed into the channel at higher discharge voltages and lower anode mass flow rate though the relationship is difficult to quantify. Physically, we may expect an increase in the discharge voltage to cause a nonlinearly greater amount of erosion to the channel walls due to rising particle energy and the acceleration zone recessing into the discharge channel. Analogously, the fact that the acceleration zone is pushed downstream when the anode mass flow rate is increased may mean comparatively lower erosion rates when a higher anode mass flow rate is used. Although the trends cannot be clearly quantified, the exact location of the acceleration should, in theory, play an important role in determining erosion rate.

With the information from Table 5-2, we can now re-process Figure 5-37 to look for universal features in the acceleration zone. Figure 5-39 is a specially normalized plot of bulk velocity versus position. The bulk velocity is normalized against the maximum recorded bulk velocity for each condition. The axial position is normalized against the acceleration zone length as previously defined. The zero of the normalized axial position is set at the position where the bulk velocity reaches 10% of the maximum bulk velocity, which is also the start of the acceleration zone as previously defined. Note that the seven curves overlap each other very well suggesting a near-identical acceleration-zone structure for all tested operating conditions. The fact that these acceleration profiles can be collapsed on top of each other can be related to Morozov's equipotential theory [45].

Since the seven operating conditions have near identical magnetic field topology, the electric potential contours, particularly along the channel centerline, should also be very similar. This similarity should give rise to self-similar acceleration profiles.

To extend equipotential idea further, if the electric potential extends along the magnetic field line all the way to the channel walls, then the near-wall acceleration profile should also be similar to what is plotted in Figure 5-39. The entire acceleration profile can be determined by only three parameters, the discharge voltage, the acceleration zone length, and the acceleration zone start location. This knowledge is highly desirable because it allows the formation of an engineering model to make first-order estimates of the ion energy along the channel walls. The idea of such a model is explored in Section 8.5.

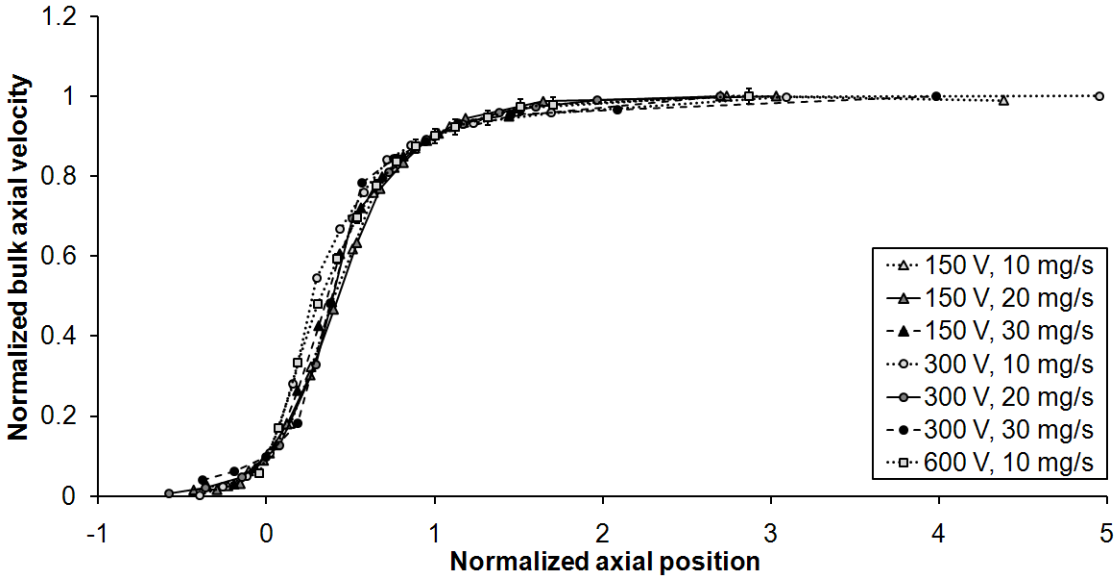

Figure 5-39. Maximum-velocity-normalized bulk axial velocity versus acceleration-zone-length-normalized axial position.

5.4.2 Two-Axis Ion Velocity Vector Map

Instead of starting with the VDF results first as was done with the one-axis acceleration zone study, we will start with the extracted bulk velocity results. Figures 5-40 to 5-46 show the two-axis singly-charged ion velocity vector maps for the 6-kW Hall thruster at various operating conditions. Empty spots in the velocity vector maps are due to low SNR for the LIF traces obtained at those spots. The exact number of empty spots can be determined by comparison against Figure 5-1, the data location matrix for the two-axis near-wall study. In general, there are many more empty spots for 30 mg/s operating conditions than for the 10 mg/s and 20 mg/s operating conditions. The number of empty spots is indicative of the fact that SNR is generally lower everywhere in the acceleration zone for the 30 mg/s operating conditions. This is in agreement with the data from the acceleration zone LIF study. Since signal is fairly constant, the decrease in SNR is due to increase in noise, and the increase in noise is believed to be driven by plasma oscillations. Unfortunately, the amount of plasma oscillation was never quantified. The only direct evidence of increased plasma oscillation is that, at 30 mg/s conditions, the discharge current measurement, which is time-averaged to some extent, fluctuates by a few tenths of an Ampere, while at 10 and 20 mg/s conditions, the discharge current measurement fluctuates by a few hundredths of an Ampere.

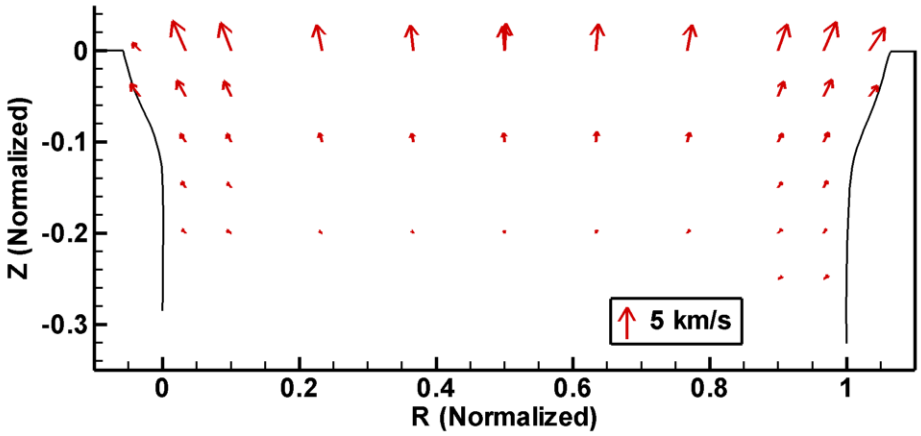

Figure 5-40. Singly-charged ion bulk velocity vector map of the 6-kW Hall thruster operating at 150 V, 10 mg/s.

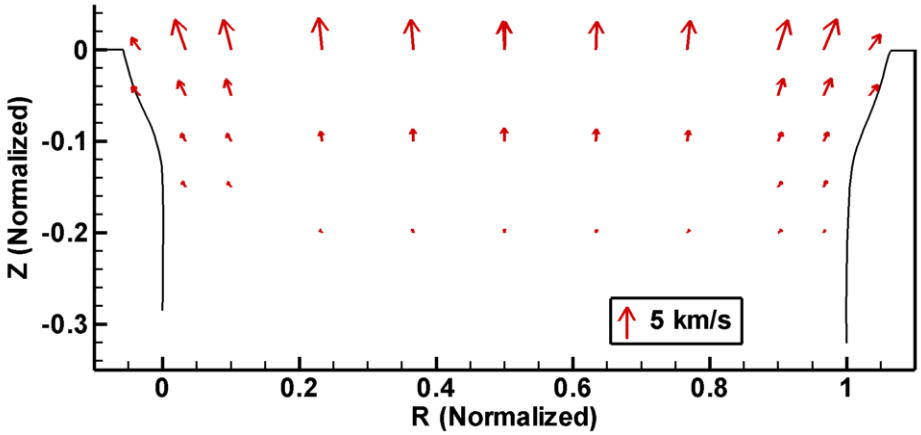

Figure 5-41. Singly-charged ion bulk velocity vector map of the 6-kW Hall thruster operating at 150 V, 20 mg/s.

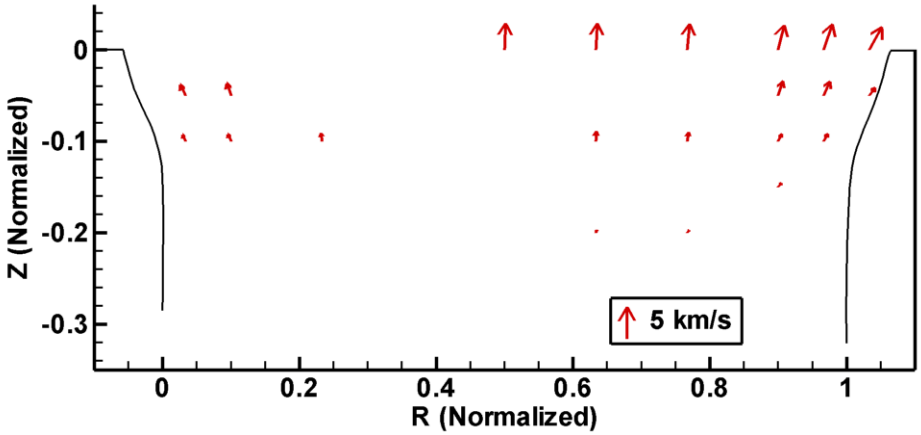

Figure 5-42. Singly-charged ion bulk velocity vector map of the 6-kW Hall thruster operating at 150 V, 30 mg/s.

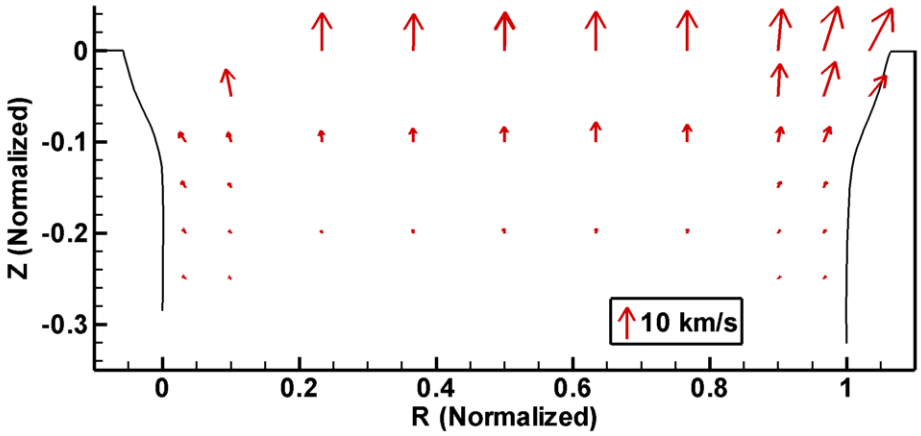

Figure 5-43. Singly-charged ion bulk velocity vector map of the 6-kW Hall thruster operating at 300 V, 10 mg/s.

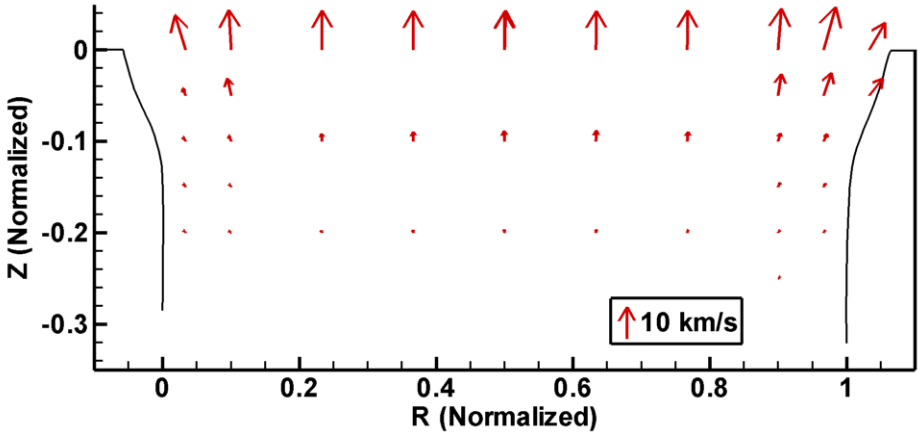

Figure 5-44. Singly-charged ion bulk velocity vector map of the 6-kW Hall thruster operating at 300 V, 20 mg/s.

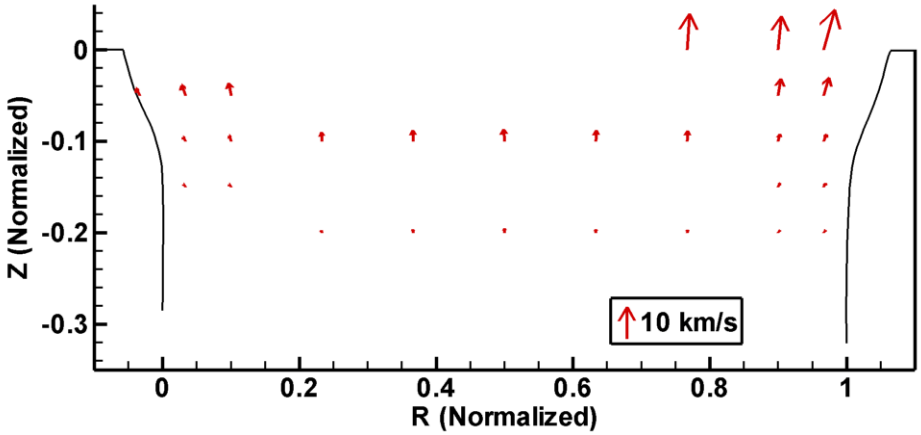

Figure 5-45. Singly-charged ion bulk velocity vector map of the 6-kW Hall thruster operating at 300 V, 30 mg/s.

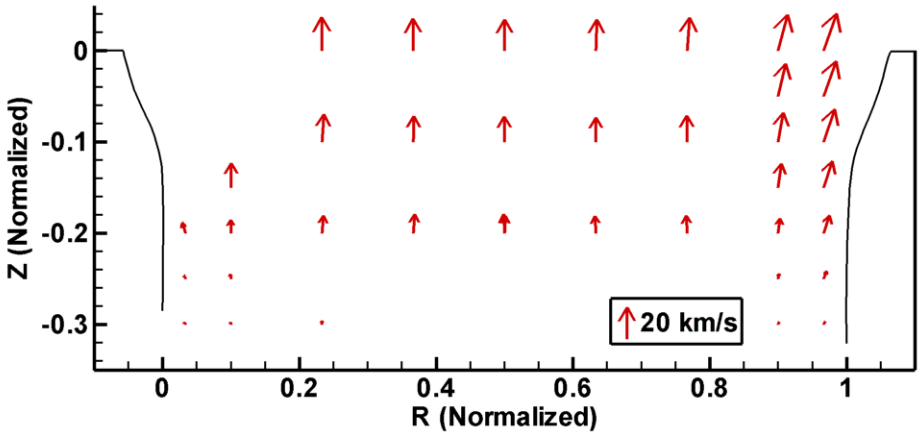

Figure 5-46. Singly-charged ion bulk velocity vector map of the 6-kW Hall thruster operating at 600 V, 10 mg/s.

From these bulk velocity vector maps, we can see that the bulk speed of the singly-charged ions is very low until they reach the acceleration zone near the exit plane. The velocity vectors are generally axial near the center of the channel but diverge near the two channel walls. The near-wall velocity vectors appear to be tilted towards the wall at all measured axial locations for all operating conditions. Two effects are speculated to be contributing to this phenomenon. One, ion populations have a spread of radial velocity and its density is much lower near the wall than at the center. Where the density gradient is strong, there are more ions moving away from the region of high density than away from the region of low density. The near-wall ion density is relatively low because the Hall current is confined to the center of the channel and the ions that make contact with the wall tends to become neutralized. Furthermore, to balance the electron current going to the channel walls, some ion current must also go to the walls. This is the basis for ambipolar diffusion. Two, there could be deviation of the contours of electrical potential from the magnetic field line due to the strong gradients in plasma density. It is unclear how much each of the above effects contributes to the defocusing of the ions near the wall.

It is also interesting to note that the sizes of the arrows, indicative of the bulk speeds of the ions, are roughly the same across the channel at any given axial location. These data give empirical backing to the idea that the near-wall acceleration profiles are similar to the centerline acceleration profile. A quantitative comparison of the profiles will be made later in Section 5.4.5 after more of the data has been shown.

5.4.3 Ion Angular Energy Distribution Functions

As previously mentioned in Section 5.3.5, the singly-charged ion angular energy distribution functions (AEDFs) are assumed to be two-temperature Maxwell-Boltzmann distribution functions. However, in keeping with the idea that temperature is physically inaccurate, the axial and calculated radial FWHM velocities are used instead. Tables 5-3 to 5-9 show the key parameters that describe the angular energy distribution functions (AEDF) at each data location for the 6-kW Hall thruster at various operating conditions. These tables list, in order, the non-dimensional axial coordinate, Z, the non-dimensional radial coordinate, R, the average speed, the average angle, the FWHM velocity in the axial direction, the FWHM velocity in the radial direction, and the average SNR. Z is the axial position normalized by the channel length where Z = 0 at the channel exit plane and Z = -1 at the anode exit plane. R is the radial position normalized by the channel width where R = 0 at the nominal inner wall radius and R = 1 at the nominal outer wall radius. The average speed is calculated from the two bulk velocities extracted at each location, taking into account the geometries of the setup. The average angle is the angle of the bulk flow with respect to the firing axis, where a positive angle indicates radially outward flow and a negative angle indicates radially inward flow. The calculations for the FWHM velocities were previous explained in Section 5.3.5. The average SNR is the average of the SNR found in the axial and the off-radial traces. This average SNR is given to help assign uncertainties to the other listed values.

Each table contains two portions, separated by a dash line. The first portion is data collected nearer to the inner wall of the discharge channel while the second portion is data collected nearer to the outer wall. Each portion of data is taken on a different day from the other portion because they require a different experimental setup to collect. Some data locations are overlapped between the two portions to make sure approximately the same results are found on different days. The near-wall data locations are highlighted with a grey background for easier identification. A * denotes that equation (5-14) is used to calculate the corresponding value instead of equation (5-13).

Table 5-3. Key parameters describing the angular energy distribution functions for the 6-kW Hall thruster operating at 150 V, 10 mg/s.

Z	R	Average speed, m/s	Equivalent energy, eV	Average angle, deg.	Axial FWHM velocity, m/s	Radial FWHM velocity, m/s	Avg. SNR
0	-0.033	2035	3	-52	9724	2667	18
0	0.033	5392	20	-30	3123	1413	29
0	0.100	4909	16	-26	2941	1761	25
0	0.233	4268	12	-16	2493	1219	23
0	0.367	4075	11	-9	2010	1194	24
0	0.500	3893	10	-6	1883	1158	29
-0.05	-0.033	2821	5	-45	6481	4730	9
-0.05	0.033	3082	6	-35	3909	1251	41
-0.05	0.100	2788	5	-31	2828	1201	37
-0.1	0.033	1638	2	-39	2136	1356	43
-0.1	0.100	1589	2	-34	1776	1037	39
-0.1	0.233	1532	2	-22	1686	908	30
-0.1	0.367	1436	1	-15	1475	966	32
-0.1	0.500	1475	1	-7	1366	1106	31
-0.15	0.033	1191	1	-38	1609	2398	20
-0.15	0.100	1089	1	-47	1279	1389	26
-0.2	0.033	929	1	-47	1491	4524*	11
-0.2	0.100	932	1	-63	1175	2288	13
-0.2	0.233	740	0	-67	874	1159	20
-0.2	0.367	463	0	-39	779	940	21
-0.2	0.500	450	0	-16	812	850	18
0	0.500	4333	13	4	1906	1258	34
0	0.633	4281	12	8	1969	1323	43
0	0.767	4207	12	14	2251	1379	42
0	0.900	4589	14	24	3002	1670	44
0	0.967	5184	18	29	3344	1300	44
0	1.033	4851	16	41	3445	2845	15
-0.05	0.900	2649	5	27	2624	1135	58
-0.05	0.967	3163	7	32	3552	1283	46
-0.05	1.033	2419	4	46	6604	4983	20
-0.1	0.633	1684	2	9	1551	1124	36
-0.1	0.767	1645	2	21	1870	994	47
-0.1	0.900	1623	2	34	1943	954	54
-0.1	0.967	1812	2	34	2359	1486	40
-0.15	0.900	1123	1	41	1406	1314	40
-0.15	0.967	1263	1	34	1714	2416	23
-0.2	0.633	587	0	28	909	823	26
-0.2	0.767	776	0	49	995	1073	26
-0.2	0.900	969	1	51	1128	2177	20
-0.2	0.967	848	0	36	1490	4105*	17
-0.25	0.900	957	1	72	1186	2010	10
-0.25	0.967	1075	1	64	1315	4788*	7

Table 5-4. Key parameters describing the angular energy distribution functions for the 6-kW Hall thruster operating at 150 V, 20 mg/s.

Z	R	Average speed, m/s	Equivalent energy, eV	Average angle, deg.	Axial FWHM velocity, m/s	Radial FWHM velocity, m/s	Avg. SNR
0	-0.033	2615	5	-44	11500	2869	11
0	0.033	5217	19	-25	4256	1850	13
0	0.100	5016	17	-19	3857	2114	10
0	0.233	5090	18	-8	2496	1409	10
0	0.367	4787	16	-6	1836	1346	13
0	0.500	4520	14	-3	1733	1409	14
-0.05	-0.033	2222	3	-49	4821	4425	10
-0.05	0.033	2801	5	-32	3688	1558	21
-0.05	0.100	2759	5	-22	3699	1616	16
-0.1	0.033	1484	1	-36	2049	1430	13
-0.1	0.100	1510	2	-37	1831	1295	13
-0.1	0.233	1603	2	-11	2342	1183	11
-0.1	0.367	1994	3	-3	2093	1344	12
-0.1	0.500	2130	3	-2	2393	1432	12
-0.15	0.033	1062	1	-38	2054	3972	6
-0.15	0.100	1068	1	-45	1294	1656	9
-0.2	0.233	685	0	-51	1014	2124	7
-0.2	0.367	635	0	-20	927	885	10
-0.2	0.500	604	0	-1	970	1137	10
0	0.500	4483	14	1	1655	1247	18
0	0.633	4501	14	2	1632	1130	24
0	0.767	4642	15	7	2105	1360	24
0	0.900	5105	18	23	3006	1948	23
0	0.967	5544	21	29	3604	1746	19
0	1.033	3357	8	44	9026	3686	15
-0.05	0.900	2905	6	22	2906	1498	25
-0.05	0.967	3134	7	31	3924	1399	28
-0.05	1.033	2437	4	48	5801	6489	11
-0.1	0.633	1982	3	3	1828	1250	17
-0.1	0.767	1792	2	11	1949	1241	21
-0.1	0.900	1676	2	29	2102	1071	23
-0.1	0.967	1717	2	31	2400	1521	19
-0.15	0.900	1131	1	42	1413	1388	15
-0.15	0.967	1163	1	30	1554	3125	10
-0.2	0.633	564	0	26	1008	758	14
-0.2	0.767	683	0	48	1030	942	13
-0.2	0.900	932	1	61	1277	1910	8
-0.2	0.967	491	0	28	2019	2817	6

Table 5-5. Key parameters describing the angular energy distribution functions for the 6-kW Hall thruster operating at 150 V, 30 mg/s.

Z	R	Average speed, m/s	Equivalent energy, eV	Average angle, deg.	Axial FWHM velocity, m/s	Radial FWHM velocity, m/s	Avg. SNR
-0.05	0.033	2016	3	-31	2823	1694	15
-0.05	0.1	2181	3	-23	2788	1593	9
-0.1	0.033	1186	1	-30	2008	1332	9
-0.1	0.1	1283	1	-30	1429	1294	8
-0.1	0.233	1277	1	-10	1635	998	6
0	0.5	4158	12	3	1646	1378	8
0	0.633	4148	12	4	2305	1281	8
0	0.767	4138	12	7	2636	1556	8
0	0.9	4060	11	18	3159	1745	11
0	0.967	4425	13	24	3976	1617	10
0	1.033	4443	13	36	6379	2137	10
-0.05	0.9	2490	4	23	2442	1414	13
-0.05	0.967	2561	4	30	2901	1325	12
-0.05	1.033	1740	2	45	5248	3231	9
-0.1	0.633	1576	2	3	1863	1023	8
-0.1	0.767	1460	1	14	1568	1268	9
-0.1	0.9	1360	1	32	1868	1247	8
-0.1	0.967	1426	1	36	1796	1455	7
-0.15	0.9	894	1	49	1469	1485	5
-0.2	0.633	531	0	28	830	1291	6
-0.2	0.767	633	0	50	1121	1066	6

Table 5-6. Key parameters describing the angular energy distribution functions for the 6-kW Hall thruster operating at 300 V, 10 mg/s.

Z	R	Average speed, m/s	Equivalent energy, eV	Average angle, deg.	Axial FWHM velocity, m/s	Radial FWHM velocity, m/s	Avg. SNR
0	0.233	11985	98	0	3507	2442	10
0	0.367	11671	93	1	3619	2030	11
0	0.5	11443	89	0	3704	1907	13
-0.05	0.1	8835	53	-14	5703	2377	12
-0.1	0.033	4045	11	-40	5767	2363	12
-0.1	0.1	3237	7	-23	2808	1565	30
-0.1	0.233	3589	9	-7	2150	1333	27
-0.1	0.367	4418	13	-3	2334	1386	25
-0.1	0.5	4903	16	-3	2561	1339	27
-0.15	0.033	2481	4	-36	3255	2548	20
-0.15	0.1	1682	2	-31	1927	1126	45
-0.2	0.033	1666	2	-36	1784	3611	20
-0.2	0.1	1188	1	-42	1260	1077	31
-0.2	0.233	1005	1	-28	1215	812	34
-0.2	0.367	1180	1	-12	1192	803	28
-0.2	0.5	1533	2	-16	1627	1072	16
-0.25	0.033	1461	1	-49	1476	3363	10
-0.25	0.1	1113	1	-56	1016	1515	24
0	0.5	12314	103	1	2964	1486	27
0	0.633	12383	104	0	2997	1655	30
0	0.767	12700	110	0	3158	1885	29
0	0.9	13251	119	7	2941	2356	21
0	0.967	15003	153	23	2834	8126*	17
0	1.033	16282	180	35	3065	7213	7
-0.05	0.9	10341	73	5	4119	2302	25
-0.05	0.967	11886	96	23	4215	4254	21
-0.05	1.033	9579	62	47	13710	7852	7
-0.1	0.633	6330	27	0	4327	2490	31
-0.1	0.767	5630	22	2	3847	2426	41
-0.1	0.9	4926	17	14	3229	2548	53
-0.1	0.967	5471	20	27	4572	2355	29
-0.15	0.9	2093	3	24	1915	1203	62
-0.15	0.967	2198	3	31	2571	1211	52
-0.2	0.633	1579	2	7	1171	1044	47
-0.2	0.767	1413	1	12	1289	914	55
-0.2	0.9	1330	1	33	1549	897	54
-0.2	0.967	1372	1	34	1679	1174	43
-0.25	0.9	1077	1	52	1141	1038	32
-0.25	0.967	1193	1	45	1216	1586	21

Table 5-7. Key parameters describing the angular energy distribution functions for the 6-kW Hall thruster operating at 300 V, 20 mg/s.

Z	R	Average speed, m/s	Equivalent energy, eV	Average angle, deg.	Axial FWHM velocity, m/s	Radial FWHM velocity, m/s	Avg. SNR
0	0.033	11680	93	-22	4173	4275	11
0	0.1	12571	108	-4	3719	2370	7
0	0.233	12338	104	0	3812	2283	8
0	0.367	12232	102	0	3988	1953	7
0	0.5	11844	95	-1	3753	1960	9
-0.05	0.033	2562	4	-24	2830	1746	32
-0.05	0.1	5589	21	-16	4006	2075	17
-0.1	0.033	1702	2	-33	3787	1551	35
-0.1	0.1	2085	3	-25	2578	1321	40
-0.1	0.233	2495	4	-4	2094	1595	23
-0.1	0.367	3099	7	-2	1774	1550	20
-0.1	0.5	3398	8	-4	1786	1374	24
-0.15	0.033	1287	1	-33	2259	1799	23
-0.15	0.1	1225	1	-37	1503	1178	28
-0.2	0.033	1015	1	-39	2008	4475	8
-0.2	0.1	1006	1	-55	1256	1352	12
-0.2	0.233	820	0	-34	1087	1021	17
-0.2	0.367	886	1	-18	1203	840	19
-0.2	0.5	1096	1	-9	1334	853	20
0	0.5	12347	104	2	3341	1584	14
0	0.633	12248	102	1	3202	1885	17
0	0.767	12410	105	1	3262	2094	18
0	0.9	13005	115	6	3062	2201	19
0	0.967	14288	139	20	2966	8111*	12
0	1.033	10799	79	38	14280	4116	7
-0.05	0.9	6942	33	12	5412	2358	30
-0.05	0.967	7809	41	24	7740	2641	24
-0.05	1.033	7687	40	46	12490	8514	6
-0.1	0.633	3455	8	3	1676	1243	31
-0.1	0.767	3110	7	8	1853	1516	28
-0.1	0.9	2803	5	19	2395	1541	36
-0.1	0.967	2673	5	31	3191	1360	39
-0.15	0.9	1519	2	31	1898	1129	33
-0.15	0.967	1424	1	33	1945	1431	29
-0.2	0.633	1092	1	14	1286	935	22
-0.2	0.767	976	1	27	1230	922	22
-0.2	0.9	1046	1	48	1292	1399	17
-0.2	0.967	949	1	36	1772	2342	9
-0.25	0.9	946	1	64	1234	1388	9

Table 5-8. Key parameters describing the angular energy distribution functions for the 6-kW Hall thruster operating at 300 V, 30 mg/s.

Z	R	Average speed, m/s	Equivalent energy, eV	Average angle, deg.	Axial FWHM velocity, m/s	Radial FWHM velocity, m/s	Avg. SNR
-0.05	-0.033	2712	5	-34	6550	7608	5
-0.05	0.033	3770	10	-25	4319	1987	13
-0.05	0.1	4062	11	-12	3443	2318	8
-0.1	0.033	1659	2	-31	2328	1650	17
-0.1	0.1	1952	3	-20	2603	1579	12
-0.1	0.233	2872	6	-4	2295	1766	6
-0.1	0.367	3451	8	-4	2159	1878	6
-0.1	0.5	3811	10	-6	1862	2454	5
-0.15	0.033	1157	1	-36	1807	3073	7
-0.15	0.1	1128	1	-38	1553	1678	11
-0.2	0.233	800	0	-32	1136	7237*	6
-0.2	0.367	1030	1	-13	1480	1202	5
-0.2	0.5	1329	1	-3	1972	983	5
0	0.767	11319	87	5	9028	4111	6
0	0.9	10716	78	7	6480	3277	8
0	0.967	13392	122	20	4500	4624	8
-0.05	0.9	5413	20	13	2928	1956	17
-0.05	0.967	6009	25	22	4917	2237	15
-0.1	0.633	3533	8	4	1946	1241	10
-0.1	0.767	3321	8	4	2011	1670	14
-0.1	0.9	2413	4	17	2468	1385	16
-0.1	0.967	2494	4	29	3024	1449	19
-0.15	0.9	1331	1	29	1693	1121	16
-0.15	0.967	1452	1	34	1767	1529	12
-0.2	0.633	1091	1	8	1246	885	12
-0.2	0.767	947	1	21	1063	983	13
-0.2	0.9	975	1	45	1137	1125	10
-0.2	0.967	1095	1	46	1464	2437	8

Table 5-9. Key parameters describing the angular energy distribution functions for the 6-kW Hall thruster operating at 600 V, 10 mg/s.

Z	R	Average speed, m/s	Equivalent energy, eV	Average angle, deg.	Axial FWHM velocity, m/s	Radial FWHM velocity, m/s	Avg. SNR
0	0.233	21222	306	-1	6053	3660	7
0	0.367	20199	278	-1	4663	2656	10
0	0.5	19736	265	0	4507	2072	13
-0.1	0.233	17586	210	5	4448	3934	6
-0.1	0.367	16412	183	2	4138	2700	12
-0.1	0.5	15875	171	-1	4018	2028	16
-0.15	0.1	15422	162	1	8807	4349	6
-0.2	0.033	7044	34	-20	9371	2773	11
-0.2	0.1	8425	48	-2	8981	5430	9
-0.2	0.233	11151	85	6	4723	3746	10
-0.2	0.367	12049	99	5	3307	2358	12
-0.2	0.5	11903	96	-3	3356	1698	17
-0.25	0.033	2370	4	-24	3403	1432	36
-0.25	0.1	2870	6	-9	2923	1603	27
-0.3	0.033	1642	2	-37	2232	1287	24
-0.3	0.1	1437	1	-22	2236	1107	31
-0.3	0.233	2218	3	-4	1780	1459	25
0	0.633	19676	263	3	4016	2492	24
0	0.767	21035	301	6	3983	3017	20
0	0.9	23878	388	19	2924	12349*	13
0	0.967	25697	449	25	2917	13717*	7
-0.05	0.9	21452	313	17	3323	13105*	13
-0.05	0.967	24332	403	25	3360	13530*	9
-0.1	0.633	15429	162	0	3880	2167	18
-0.1	0.767	16637	188	-1	3685	2910	22
-0.1	0.9	19022	246	13	4481	14019*	12
-0.1	0.967	22425	342	24	4695	13291*	7
-0.15	0.9	15233	158	10	5448	10916	13
-0.15	0.967	17852	217	23	6710	6762	8
-0.2	0.5	10841	80	4	2963	1554	30
-0.2	0.633	10725	78	-4	3080	1950	29
-0.2	0.767	11050	83	-4	3436	2776	29
-0.2	0.9	9604	63	10	6187	6932	20
-0.2	0.967	12266	102	24	5675	5386	11
-0.25	0.9	3382	8	11	2434	1743	51
-0.25	0.967	5295	19	26	4761	1870	38
-0.3	0.767	2284	4	6	1553	1363	62
-0.3	0.9	1689	2	23	2120	1010	55
-0.3	0.967	2004	3	30	2287	1116	49

Figure 5-47 shows a graphical representation of the AEDF data in Table 5-3 for the 6-kW Hall thruster operating at 150 V, 10 mg/s. The background coordinate grid is the non-dimensionalized spatial grid R-Z. At each data location, a square with an intensity plot of the local AEDF in velocity space is plotted. For example, the top left square in Figure 5-47 represents the AEDF of the singly-charged xenon at Z = 0, R = -0.033. Each intensity plot is plotted with red being the highest intensity and blue being the lowest. The center of each intensity plot is the origin of the velocity space. Thus the distance between the red spot and the center of the square represents the bulk speed of the population, and the direction the spot resides in with respect to the center represents the bulk angle. The velocity values of the boundaries of each square are shown at the top of each sub-figure. For example, if a red spot lies on the left edge of its intensity square-plot and the lower velocity boundary is -12000 m/s, then the bulk velocity is 12000 m/s in the negative radial direction.

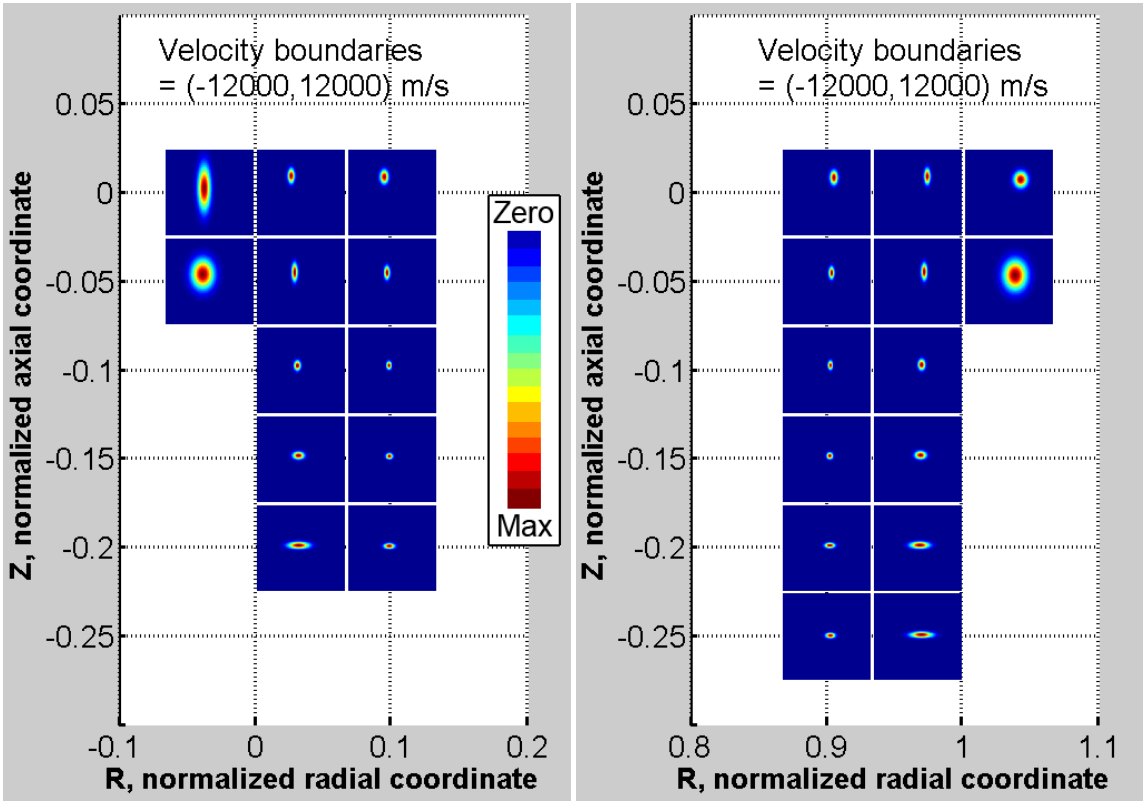

Figure 5-47. Graphical representation of near-wall angular energy distribution functions for the 6-kW Hall thruster at 150 V, 10 mg/s.

The most interesting feature of the graphical representation of the approximate AEDF is that the FWHM velocity is much larger near the walls than anywhere else. However, the direction of this broadening is different depending on how deep into the channel one travels. To be more precise, the AEDF is generally broadened in the radial direction all along the walls but is also broadened in the axial direction near the exit plane. The axial broadening is speculated to be a result of the breathing mode plasma oscillation, as seen previous in the one-axis acceleration zone study. The source of the radial broadening, on the other hand, is not clear from this graphical representation. We will investigate this further by looking at the individual VDFs.

Since the graphical representation of AEDF for other operating conditions show similar trends as Figure 5-47, they will not be shown.

5.4.4 Velocity Distribution Functions of Singly-Charged Xenon

Now we examine the individual VDFs from the two-axis near-wall LIF study. The mean speeds provided in Tables 5-3 to 5-9 do not give a complete picture of what the singly-charged ions are doing near the channel walls because ions with much higher energy than the average can exist. Erosion theories generally predict some energy threshold below which negligible erosion occurs. Yim's dissertation contains a good summary of existing models [10]. Though the specific values vary, the threshold energy is generally in the range of tens of eV. However, comparing Tables 5-3 to 5-9 and Figure 5-1 we find that at many near-wall data locations, the equivalent energy is much less than the threshold energy, and yet clear signs of channel erosion still appear in the measured profile. It is likely that the high-energy tail of the energy distribution function exceeds the threshold energy at these locations and is responsible for at least part of the observed erosion.

At 150 V discharge voltage the ions show equivalent energies of only a few eV near the channel walls, which is not of interest to erosion studies. The VDFs for 150 V conditions will not be shown.

Figures 5-48 and 5-49 show the axial VDFs of singly-charged xenon along the inner and outer channel walls, respectively, for 300 V, 10 mg/s operation. The location of the data points and the nearby channel wall profile are shown in a box at the lower right

of each figure. A pair of vertical dashed lines in the figures indicates the axial velocities that correspond to equivalent kinetic energies of 25 and 50 eV. It is assumed that the particle energies are dominated by their axial component. The values of 25 and 50 eV are chosen to be representative of typical energy threshold values and are not from any specific model. Also note that the sheath energy has not been accounted for.

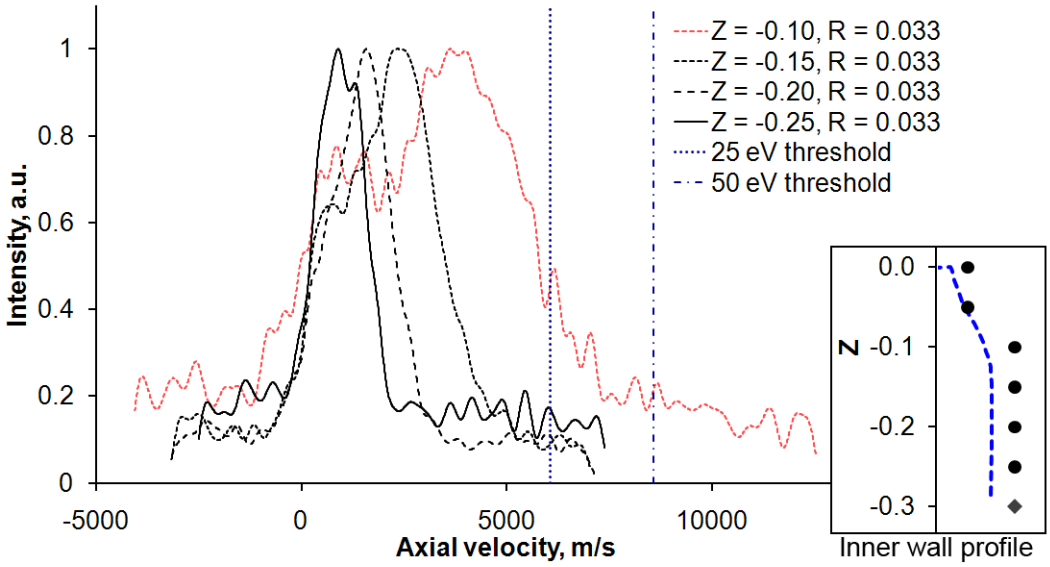

Figure 5-48. Axial ion velocity distribution functions along the inner wall of the 6-kW Hall thruster at 300 V, 10 mg/s.

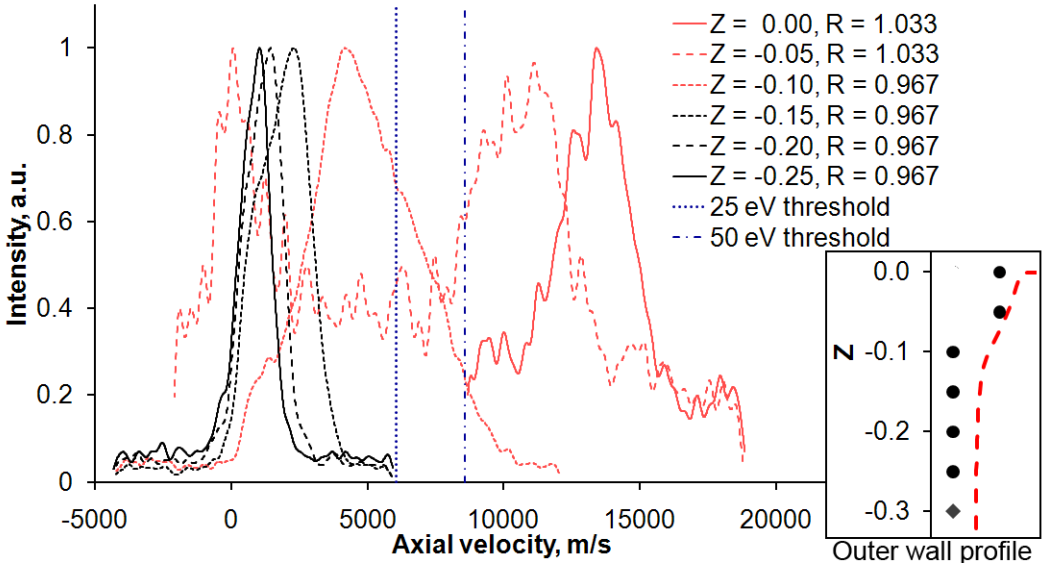

Figure 5-49. Axial ion velocity distribution functions along the outer wall of the 6-kW Hall thruster at 300 V, 10 mg/s.

As the figures show, axial locations showing discernable erosion (Z = -0.1 to Z = 0) also correspond to VDFs with measurable high energy populations. Axial VDFs for other 300 V operating conditions show similar trends. The 300 V, 10 mg/s condition was chosen as the representing example due to the completeness of clean data and the fact that ~70% of the 6-kW's operational history was occupied by 300 V operations.

Figures 5-50 and 5-51 show the axial VDFs of singly-charged xenon along the

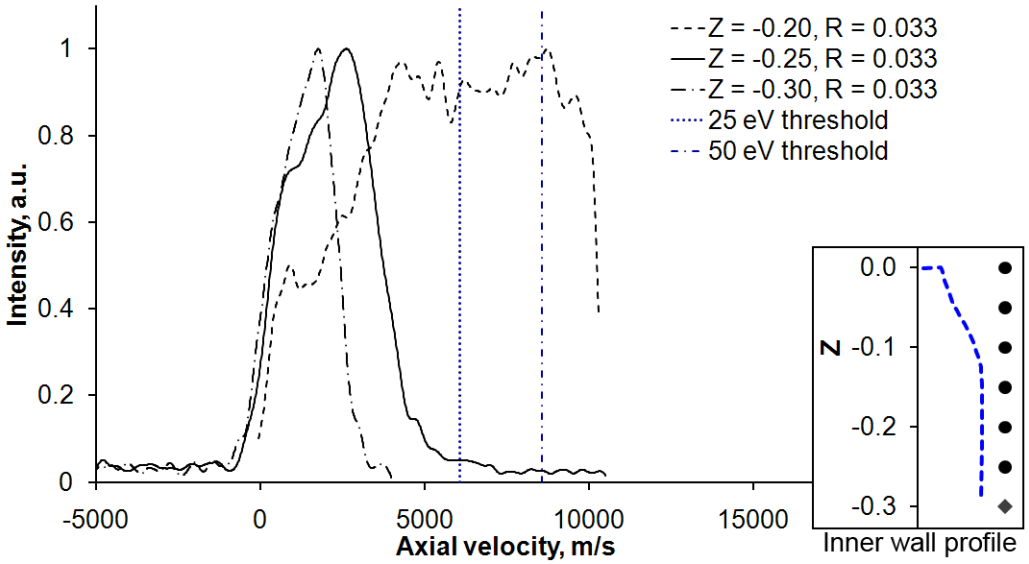

Figure 5-50. Axial ion velocity distribution functions along the inner wall of the 6-kW Hall thruster at 600 V, 10 mg/s.

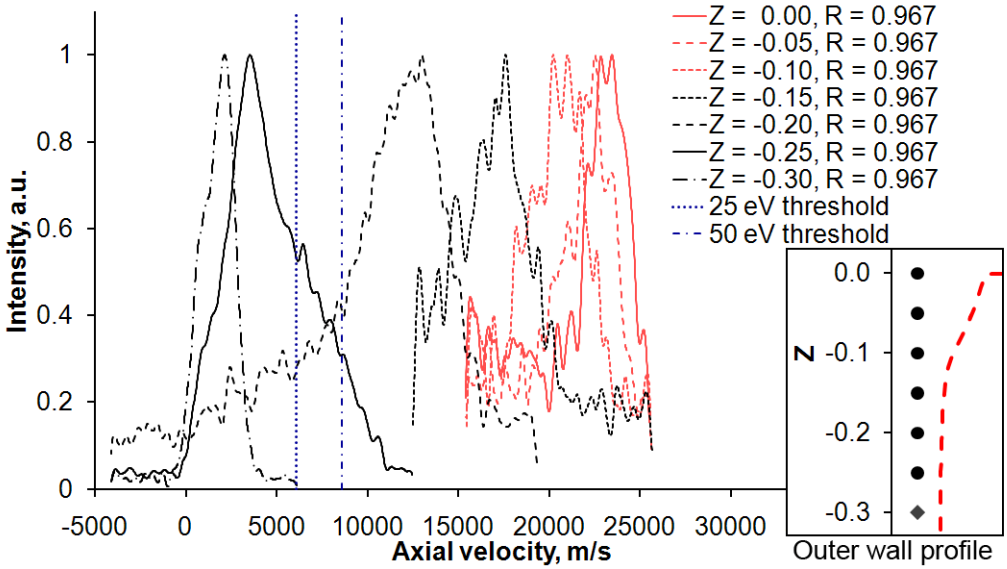

Figure 5-51. Axial ion velocity distribution functions along the outer wall of the 6-kW Hall thruster at 600 V, 10 mg/s.

inner and outer channel walls, respectively, for 600 V, 10 mg/s operation. At 600 V discharge voltage, high energy ions can be found much deeper into the discharge channel than other conditions. On the other hand, 600 V, 10 mg/s operation accounts for only ~7% of the total operational history of the 6-kW Hall thruster. It is possible, but not certain that operation at 600 V, 10 mg/s is the primary cause of the measurable erosion upstream of the Z = -0.1 axial position.

Note that Figures 5-48 to 5-51 also show large amounts of VDF broadening in the downstream locations. For example, in Figure 5-49, at Z = -0.1, the mean speed of the singly-charged ions near the walls is only 4-5.5 km/s, corresponding to 10-20 eV, but the FWHM velocity in the axial VDF is ~6 km/s, indicating a very large spread in particle energy (0-120 eV). There is also bi-modal behavior in some of the scans from data locations corresponding to the erosion zone. One possible explanation for these trends is that the breathing mode oscillation is causing the time-averaged VDF to be stretched out in much the same way as described in Section 5.4.1. If the breathing mode oscillation is indeed broadening the ion VDF, the resulting high energy ions may play a very important role in determining both the location and the rate of erosion in the Hall thruster.

Next, we move on to studying the off-radial VDFs. Once again, we will use the 300 V, 10 mg/s operation condition as the representative example for thruster operation at 300 V discharge voltage. Figures 5-52 and 5-53 show off-radial ion VDFs along the inner and outer channel walls, respectively, for the 6-kW Hall thruster operating at 300 V, 10 mg/s. Figures 5-54 and 5-55 show the same operations at 600 V, 10 mg/s. The common feature seen in all the traces that are obtained from the acceleration zone is that they are significantly broadened, typically by over an order of magnitude, when compared to the traces from upstream of the acceleration zone. Furthermore, they are all broadened in a way that a large part of the VDF hovers over 0 m/s velocity mark. One possible explanation is that there is significant contribution to the local singly-charged population from charge exchange events near the walls. When the fast ion enter the near-wall region where the neutral xenon density is known to be several times higher than on centerline [57], a sizable number of them may undergo charge exchange, leaving behind ions that were originally slow neutrals. Another possible explanation is that large numbers of electrons are gaining ionizing energy while traveling radially through the sheath and/or

crossing axially through the electric field. The large population of secondary electrons born at the walls would be a possible source of electrons that gain energy through the sheath. Ions born from impact by these electrons will have very low kinetic energy compared to ions from the bulk plasma. Both effect broadens the VDF by giving it a low energy tail.

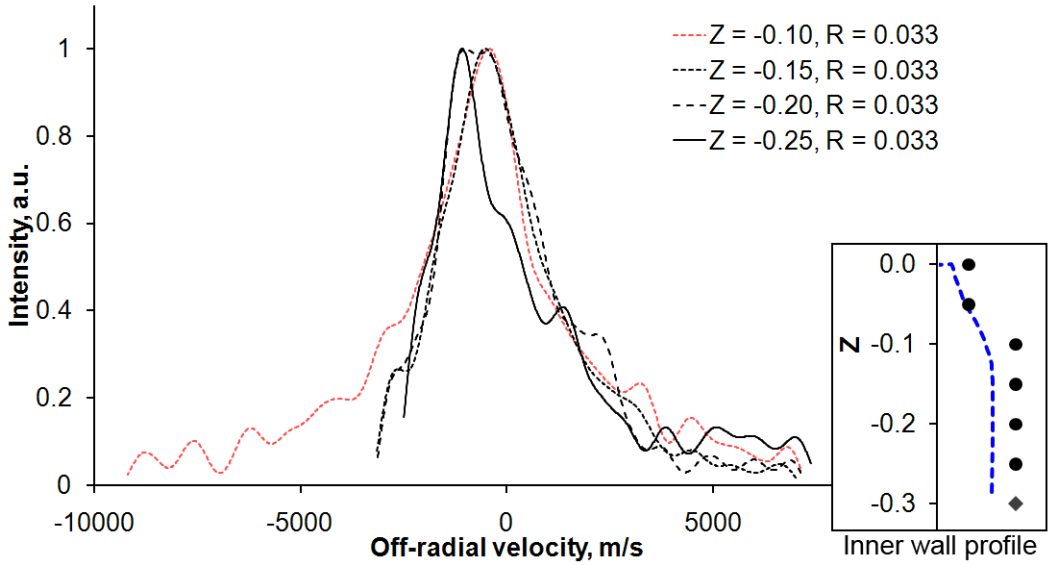

Figure 5-52. Off-radial ion velocity distribution functions along the inner wall of the 6-kW Hall thruster at 300 V, 10 mg/s.

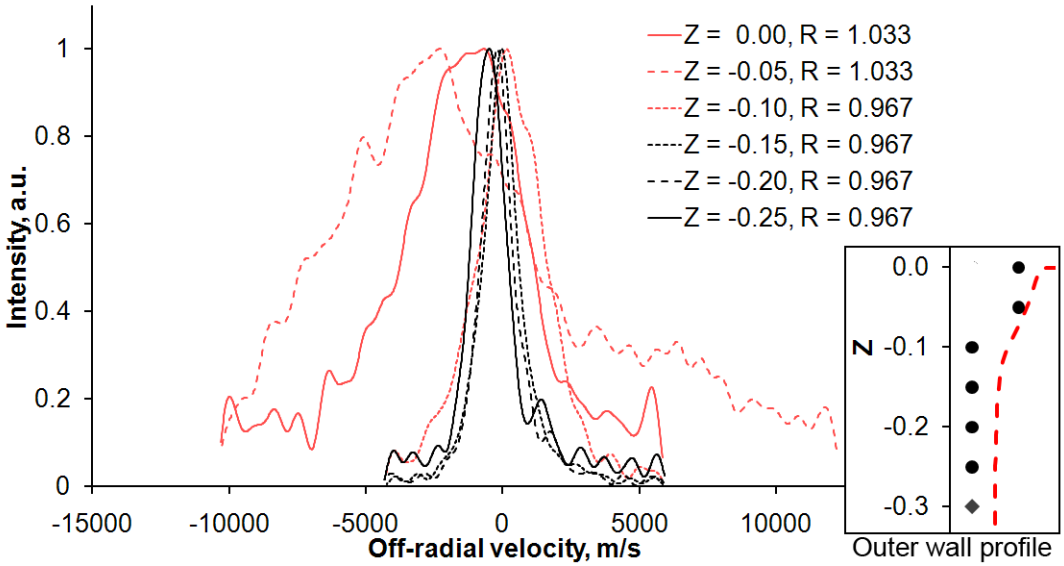

Figure 5-53. Off-radial ion velocity distribution functions along the outer wall of the 6-kW Hall thruster at 300 V, 10 mg/s.

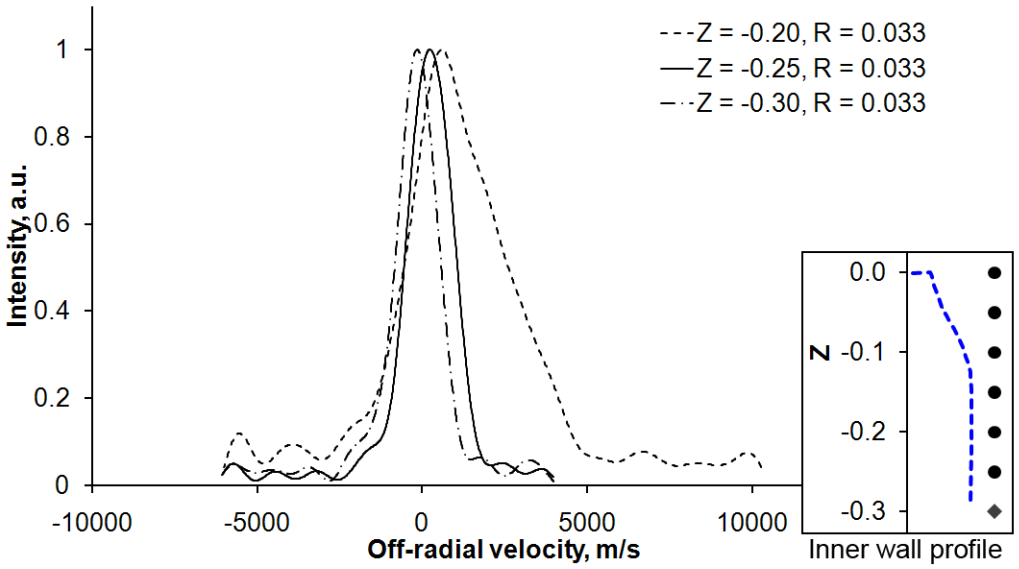

Figure 5-54. Off-radial ion velocity distribution functions along the inner wall of the 6-kW Hall thruster at 600 V, 10 mg/s.

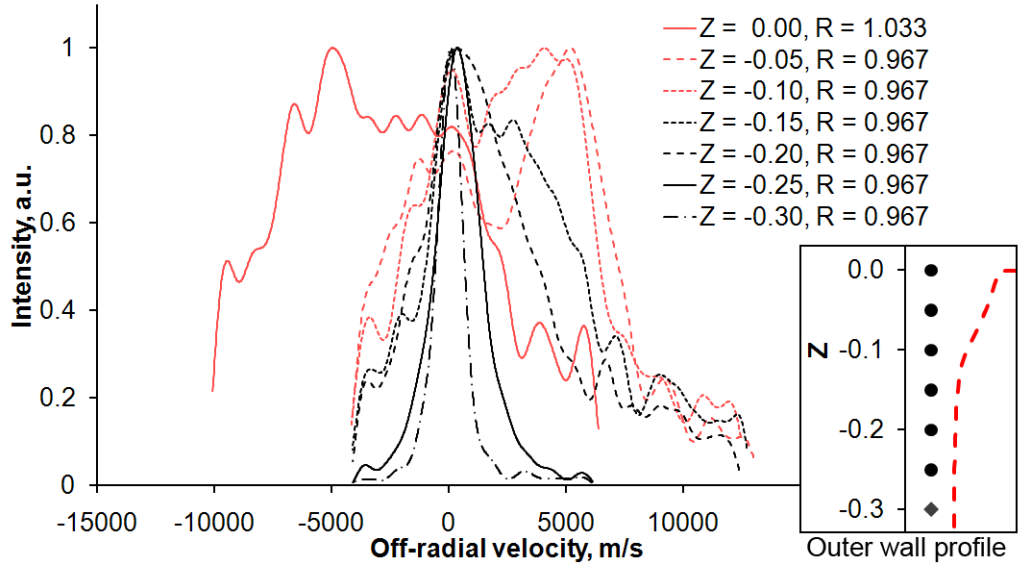

Figure 5-55. Off-radial ion velocity distribution functions along the outer wall of the 6-kW Hall thruster at 600 V, 10 mg/s.

5.4.5 Comparison of Centerline versus Near-Wall Ion Speed

Before we begin the comparison, it is important to establish a set of data to use as the reference for the singly-charged-ion speed along the centerline of the discharge channel of the 6-kW Hall thruster. Figure 5-56 compares the one-axis centerline data taken during the acceleration zone study against the axial velocity measured during the two-axis near-wall study. The motivation behind using the data set from the one-axis acceleration zone study, as opposed to the two-axis study, is that it has a finer spatial resolution and a higher signal-to-noise ratio, and, by extension, a lower uncertainty. In this figure, whenever there are two measurements at a given data location (due to overlap in the data locations of the inner and outer wall experiments), the bulk axial and radial velocities are averaged and a new speed and angle is calculated.

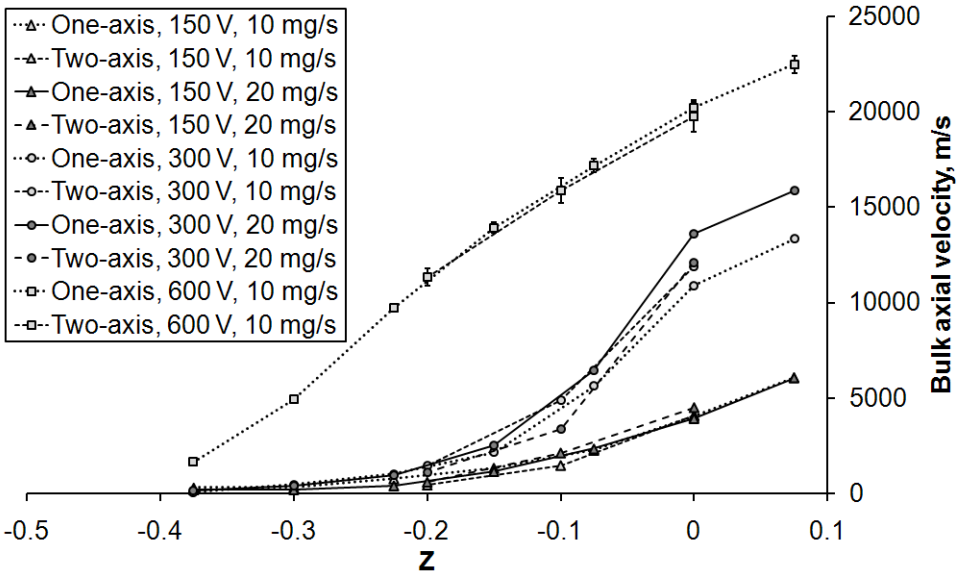

Figure 5-56. Comparison of the bulk axial velocity between the one-axis and two-axis LIF studies along the centerline of the 6-kW Hall thruster's discharge channel. *Representative error bars drawn for the 600 V, 10 mg/s data sets.*

From Figure 5-56, we can see that the variations in the bulk axial velocity between the two LIF studies are small. Furthermore, a careful review of the two-axis data set reveals that the ion speed is no more than 1% different from the bulk axial ion velocity along the centerline, or, in other words, the radial component makes negligible contribution at these locations. Thus, the bulk axial velocity of the ions from the one-axis data set is also roughly equal to the speed of the ions. Operating conditions with anode

mass flow rate of 30 mg/s are not shown because the two-axis data set for these conditions are not complete enough for meaningful comparison.

Figures 5-57 and 5-58 show comparisons of the singly-charged-ion speed along

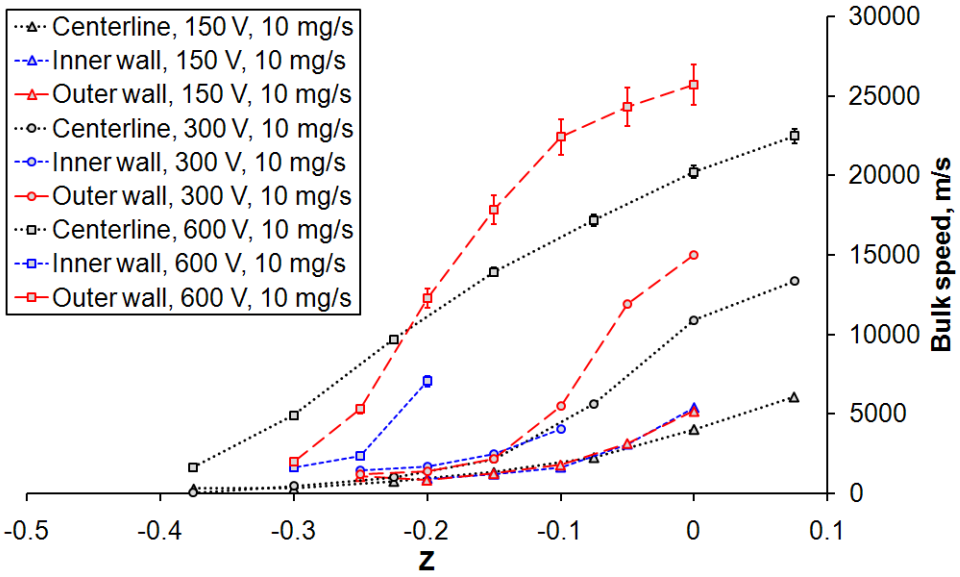

Figure 5-57. Comparison of the singly-charged-ion speed along the channel centerline versus along the channel walls for operating conditions with 10 mg/s anode mass flow rate.
Representative error bars drawn for the 600 V, 10 mg/s data sets.

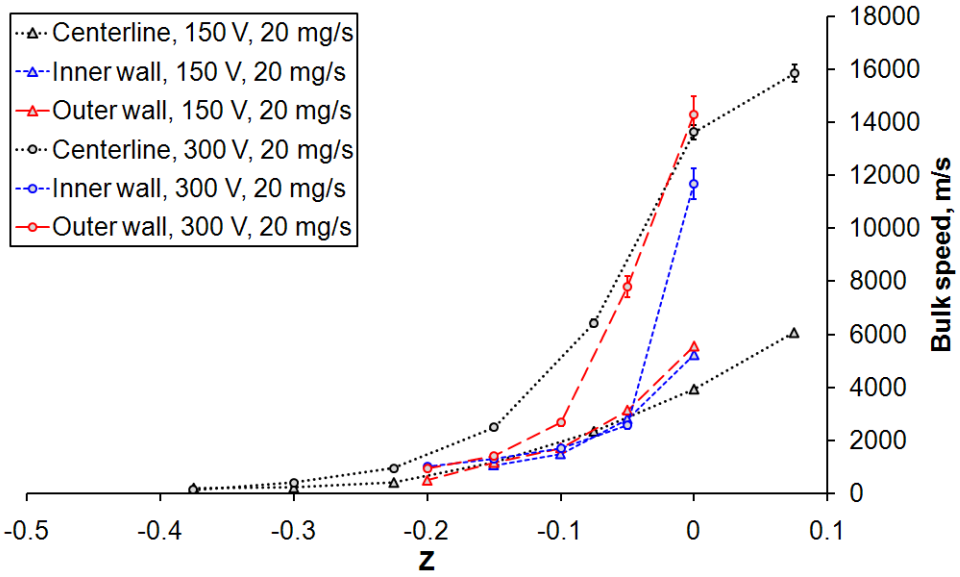

Figure 5-58. Comparison of the singly-charged-ion speed along the channel centerline versus along the channel walls for operating conditions with 20 mg/s anode mass flow rate.
Representative error bars drawn for the 300 V, 20 mg/s data sets.

the channel centerline versus the channel walls for various operating conditions. The speed along the centerline is assumed to be equal to the axial velocity found during the one-axis acceleration zone study, while the speed near the walls are calculated using the axial and radial velocities found in the two-axis near-wall study. The ion speed along the inner and outer walls is calculated from the Z = 0.033 and the Z = 0.967 data, respectively. Some data, particularly along the inner wall, are missing due to poor SNR.

Generally speaking, the singly-charged ions appear to start accelerating further downstream along the walls than along the centerline of the discharge channel. However, their acceleration also takes place over a narrower spatial extent in the axial direction. In particular, at 600 V, 10 mg/s, the ions do not start accelerating along the outer wall until Z ≈ -0.3, but the bulk speed quickly catches up and overtake the values found along centerline by Z ≈ -0.2. The ions near the outer wall continue to maintain a higher speed than what is found along the centerline all the way downstream to the exit plane. Furthermore, with the exception of the 150 V conditions, the ions along the outer wall appear to undergo acceleration further upstream than the ions along the inner wall. It would appear that the acceleration zones near the walls are both shorter and slightly more downstream than the acceleration zone along the channel centerline. This behavior fits the idea of magnetic field lines being equipotentials [45] pretty well.

For a magnetic lens topology [44], the magnetic field lines become denser near the wall and are also further downstream than along the centerline. If these field lines approximate the plasma potential contours, then ions traveling along the walls will undergo acceleration over a narrower axial extent. This idea is illustrated graphically in Figure 5-59, where the red and blue dots are the beginnings and ends of the acceleration zones, respectively, and the green dashed line shows the length of the acceleration zone. The length of these lines are set by the fact that all of the red dots should lie on one magnetic field line while all of the blue dots lie on

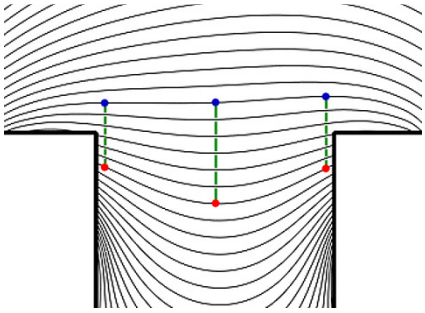

Figure 5-59. An example of how magnetic field lines being electric equipotential can give rise to shorter acceleration zone near the walls when compared to the channel centerline.

another. With this said, we know the magnetic field topology for the 6-kW Hall thruster is symmetric about the channel centerline, yet the inner-wall acceleration zone is axially downstream of the outer-wall acceleration zone, so the magnetic field lines do not perfectly follow the electric equipotentials.

Although the behavior of the singly-charged ions along the channel centerline and along the walls does not match perfectly, they exhibit similar trends. In general, the acceleration zone at all three locations (inner wall, centerline, outer wall) lengthens and recesses into the channel with increasing discharge voltage.

5.4.6 Multi-peaked Velocity Distribution Functions

This section describes some of the ion VDFs that were found to contain multiple peaks. The information in this section is not strictly necessary for understanding the rest of this dissertation but could provide interesting insights into physics in the Hall thruster plasma. Though not absolute certain due to a lack of experimental evidence, some of these physics may play important roles in the erosion problem.

In general, the VDF of the ion exhibit multi-peaked behavior in two types of location, in the acceleration zone and near the channel walls. Some of these multi-peaked VDFs and possible physical explanations have already been briefly described and will be elaborated upon in this section. The multi-peaked behavior associated with the ions found in the acceleration zone of the thruster towards the center of the channel will be discussed first.

Figure 5-60 shows several of the more prominent multi-peaked traces found during the LIF tests. As previously mentioned, this phenomenon is most likely a result of taking time-averaged LIF traces of an oscillatory plasma. The dominant mode of oscillation that can cause ion velocity to fluctuate is the breathing mode. The physics of the breathing mode is explained in great detail in [82]. The detail most important to the multi-peaked phenomenon is that the breathing mode causes an axial oscillation in the

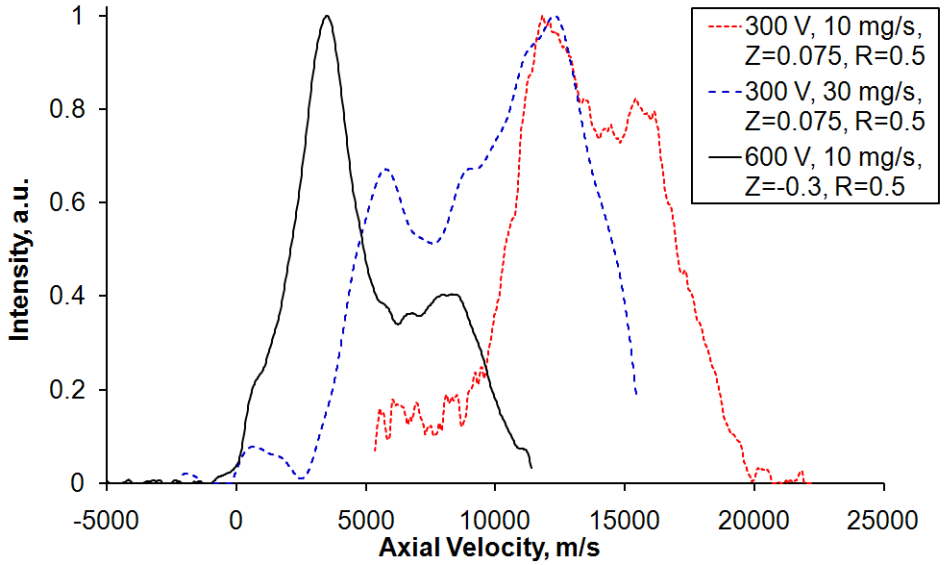

Figure 5-60. Examples of multi-peaked axial ion VDFs found in the acceleration zone along the channel centerline.

location of the acceleration zone. Since the interrogation location is fixed for a given LIF trace, the local ion velocity will appear to fluctuate as different part of the acceleration zone is interrogated at different instance in time.

Figure 5-61 illustrates how a sinusoidal oscillation in velocity can give rise to the appearance of two distinct peaks in a time-averaged VDF. Sub-figure a shows how the axial velocity fluctuates with respect to time. Sub-figure b shows the corresponding probability distribution function (PDF) that describes how likely the velocity is to take on a particular value at a given instance in time. Sub-figure c shows a typical ion VDF that one would measure if one can freeze time. The time-averaged LIF trace can be treated mathematically as a convolution of the PDF (sub-figure b) and the instantaneous VDF (sub-figure c). The result is a bi-modal distribution function. Note that bi-modal behavior induced by time-averaging breathing mode motion should only be observed for axial LIF

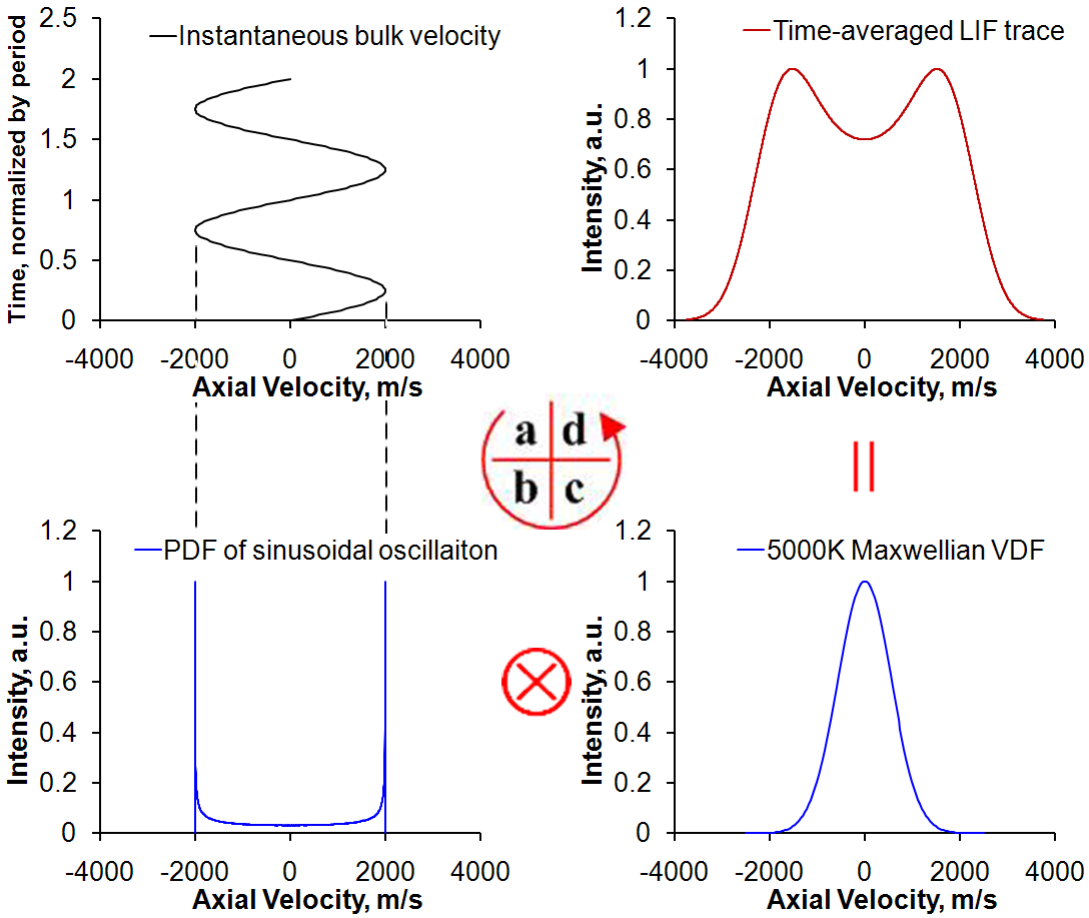

Figure 5-61. An illustration of how sinusoidal velocity oscillation give rise to a bi-modal time-averaged VDF.

traces because the breathing mode should only cause axial displacement of the acceleration zone.

Next, we look at bi-modal scans found near the channel walls. Figure 5-62 shows one of the cleanest set of LIF data that captures bi-modal VDF behavior near the channel walls. These VDFs all originated from the 150 V, 10 mg/s condition because the SNR of the data is in general higher than data from other conditions. Near-wall bi-modal behavior was also observed for other operating conditions though the peaks generally harder to distinguish. Note that all of the VDFs shown in this figure are from the Z = -0.05, R = -0.033 and Z = -0.05, R = 1.033 locations. The interrogation zones for these two data locations were previously identified as intersecting the sheath and the channel wall surface. Note that for the axial VDF at Z = -0.05 and R = -0.033, there are a lot of noise spikes. These spikes can be distinguished from true signal by the fluctuation amplitude of the noise, which is roughly constant throughout the trace. In other words, estimate the noise amplitude from the left end of the trace where there should be little signal and any spikes lower than that amplitude found elsewhere in the trace is most likely just noise. It is readily apparent that each of the four traces shown in Figure 5-62 contains at least two peaks. While the bi-modal behavior of the axial traces could have been due to the breathing mode, that oscillation mode should have little effect on the off-radial traces. Even more telling is the fact that each of the four traces has a peak hovering over the

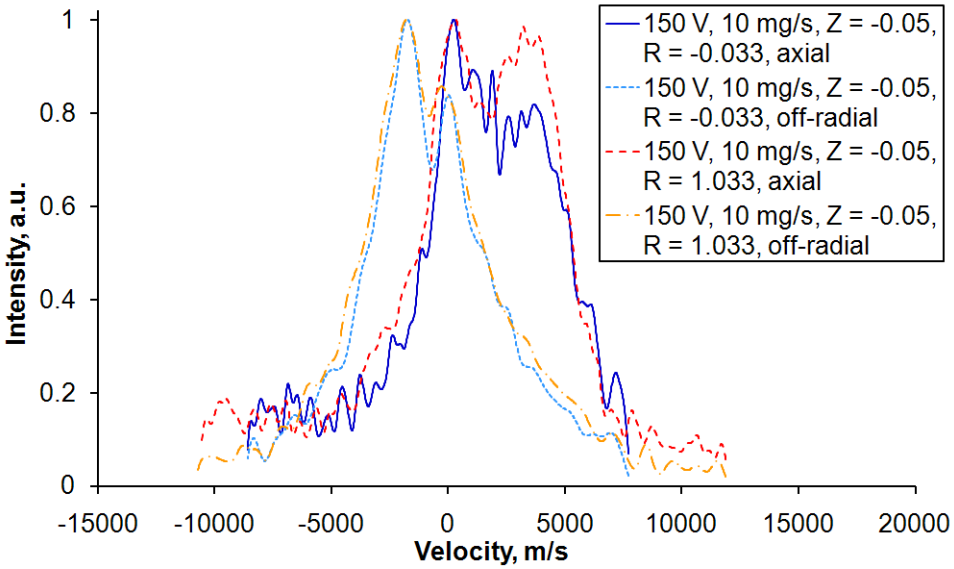

Figure 5-62. Examples of near-wall multi-peaked VDFs.

zero-velocity marker.

Figure 5-63 shows a comparison of the 150 V, 10 mg/s, Z = -0.05 axial VDFs at R = 0.967 and R = 1.033. Although the two data points are immediate neighbors on the data location matrix (Figure 5-1) and the share the same axial coordinate, one displays bi-modal behavior while the other does not. This fact combined with the bi-modal behavior in the off-radial traces suggests the cause of the two-peaked appearance in these near-wall VDFs is probably not plasma oscillations.

The simplest explanation for the observed behavior of the VDFs in Figure 5-62 is that there is a zero-velocity population on top of the beam ions in the VDF. There are two possible explanations for the presence of such ions. One, charge-exchange events can cause neutrals with relatively low speed to become singly-charged ions. It is not clear what role, if any, the sheath and the wall plays in these events. Charge-exchange events may be happening near the walls and not the center of the channel for the simple reason that the neutral density is the highest there [83]. Two, the sheath potential is typical tens of volts and electrons accelerating through the sheath can gain enough energy to impact ionize the neutral xenon (ionization energy of ~12 eV). By scanning close to the walls, one can detect these ions, born in the sheath, before they start accelerating. Note that, in theory, most of these ions should fall toward the walls due to the direction of the local sheath electric field.

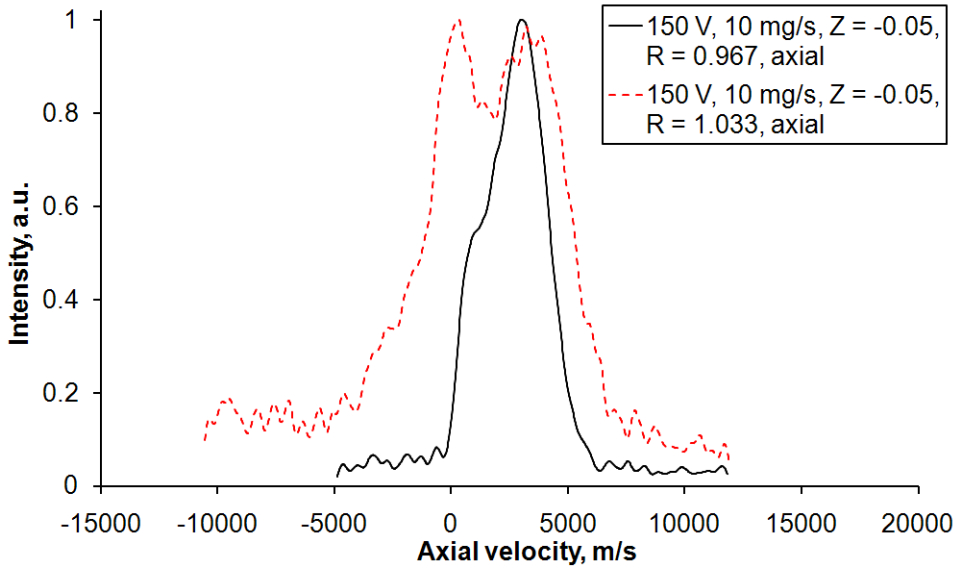

Figure 5-63. Comparison of near-wall axial VDFs from two neighboring data locations that share the same axial coordinate.

5.5 LIF Results for Neutral Xenon

5.5.1 Two-Axis Neutral Velocity Vector Map

Figures 5-64 to 5-70 show the two-axis neutral xenon velocity vector maps for the 6-kW Hall thruster at various operating conditions.

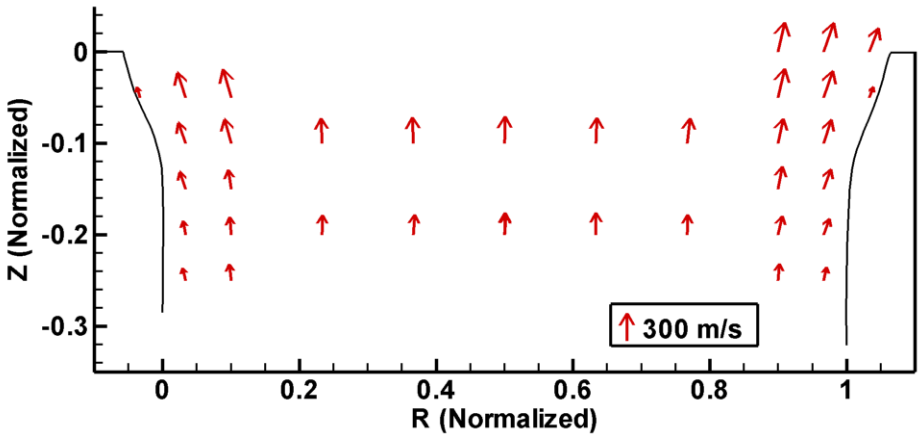

Figure 5-64. Neutral bulk velocity vector map of the 6-kW Hall thruster operating at 150 V, 10 mg/s.

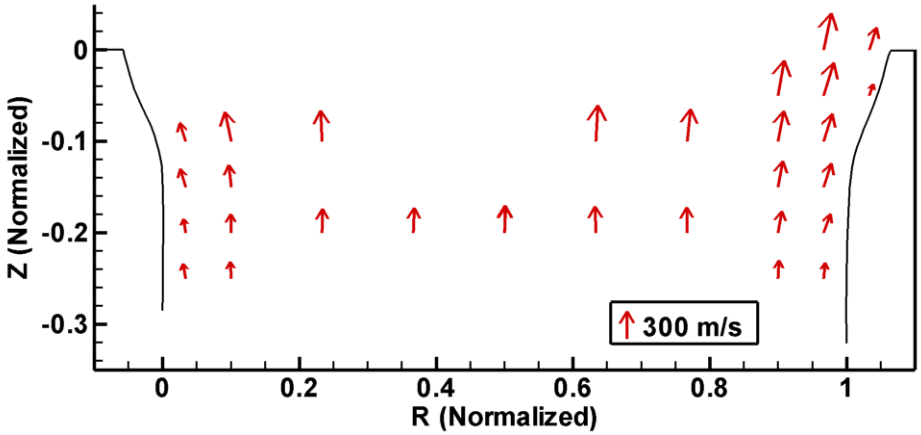

Figure 5-65. Neutral bulk velocity vector map of the 6-kW Hall thruster operating at 150 V, 20 mg/s.

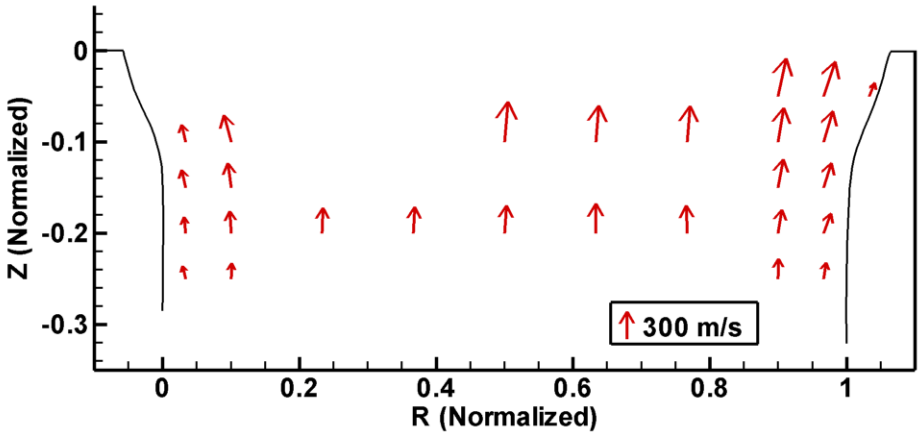

Figure 5-66. Neutral bulk velocity vector map of the 6-kW Hall thruster operating at 150 V, 30 mg/s.

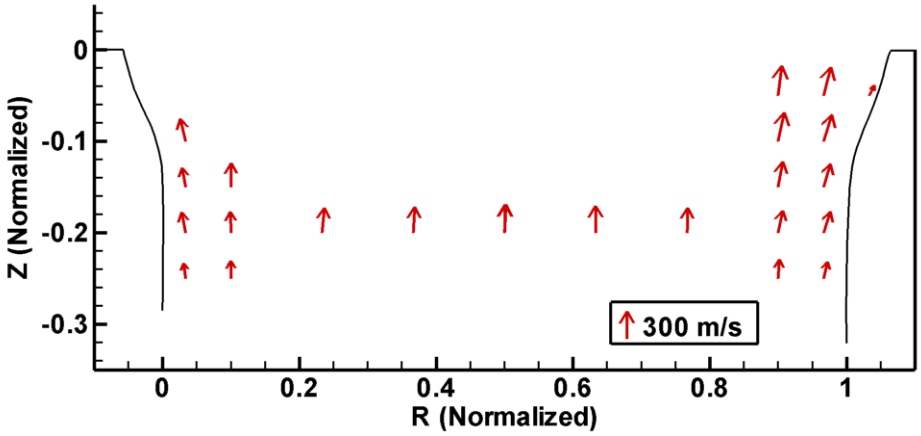

Figure 5-67. Neutral bulk velocity vector map of the 6-kW Hall thruster operating at 300 V, 10 mg/s.

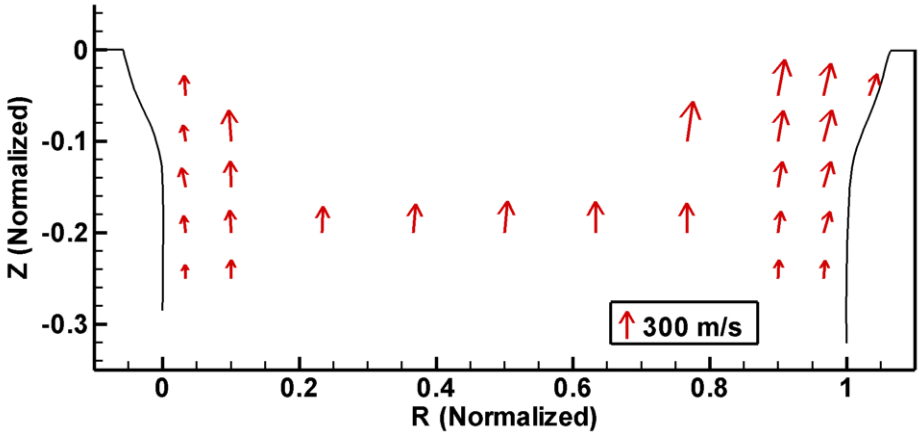

Figure 5-68. Neutral bulk velocity vector map of the 6-kW Hall thruster operating at 300 V, 20 mg/s.

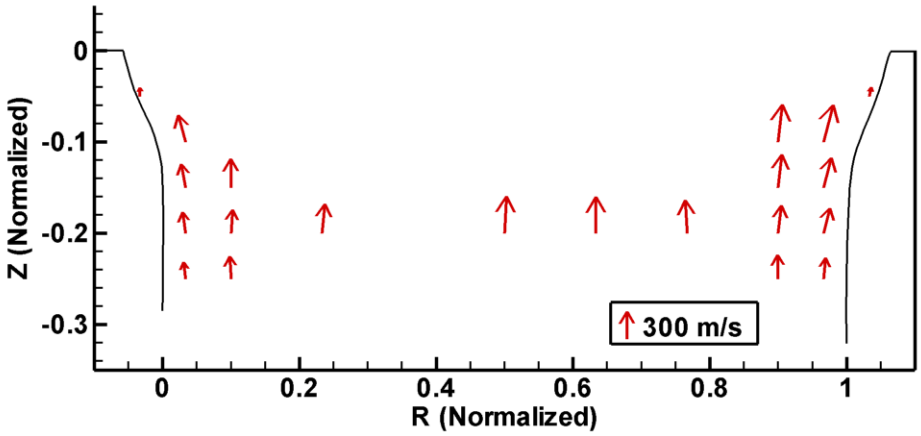

Figure 5-69. Neutral bulk velocity vector map of the 6-kW Hall thruster operating at 300 V, 30 mg/s.

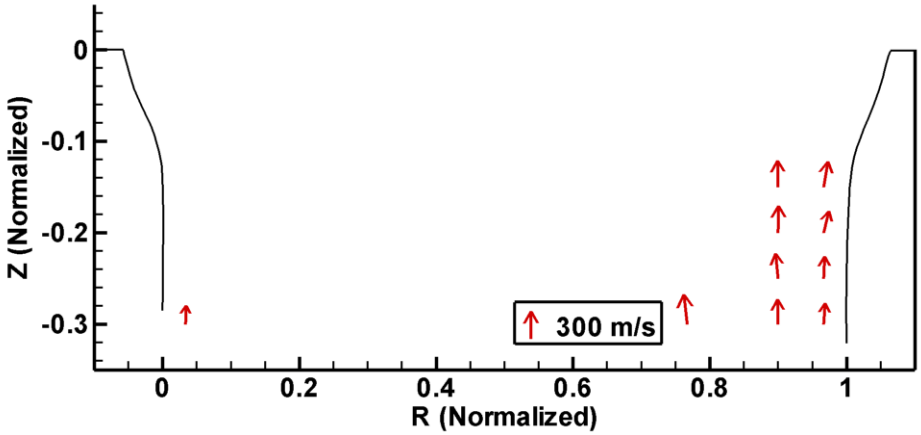

Figure 5-70. Neutral bulk velocity vector map of the 6-kW Hall thruster operating at 600 V, 10 mg/s.

A comparison between the size of the neutral velocity vectors and the reference vectors reveal that the neutral bulk speed is at most a few hundred m/s, far lower than the energy threshold needed to cause channel wall erosion. So it would seem that either bombardment by neutral xenon is not a major source of channel wall erosion or that higher energy neutrals have simply not been detected. A closer look at the neutral xenon VDF is needed and will be done next.

5.5.2 Raw Neutral Xenon LIF Traces

Figure 5-71 shows the raw axial neutral LIF traces along the inner wall of the 6-kW Hall thruster operating at 300 V, 20 mg/s. Figure 5-72 shows the same for the traces along the outer wall. Note that hyperfine structures and Zeeman effect have not been removed from these traces.

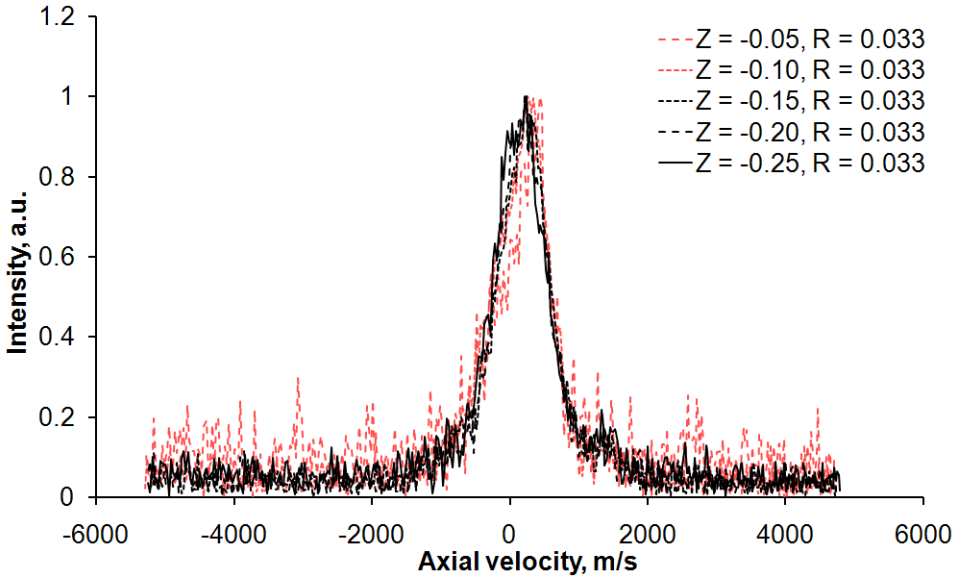

Figure 5-71. Raw axial neutral velocity distribution functions along the inner wall of the 6-kW Hall thruster at 300 V, 20 mg/s.

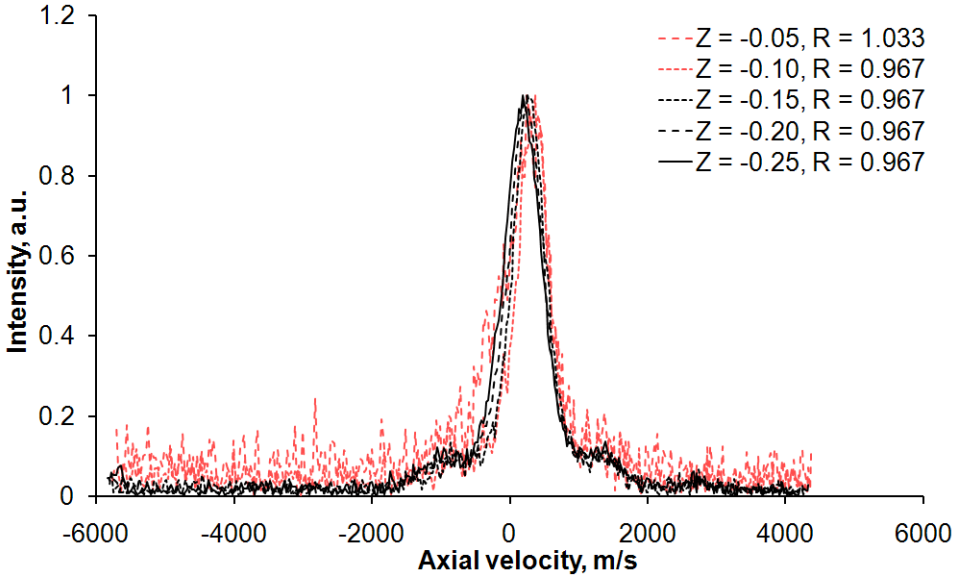

Figure 5-72. Raw axial neutral velocity distribution functions along the outer wall of the 6-kW Hall thruster at 300 V, 20 mg/s.

Figure 5-73 shows the raw off-radial neutral LIF traces along the inner wall of the 6-kW Hall thruster operating at 300 V, 20 mg/s. Figure 5-74 shows the same for the traces along the outer wall.

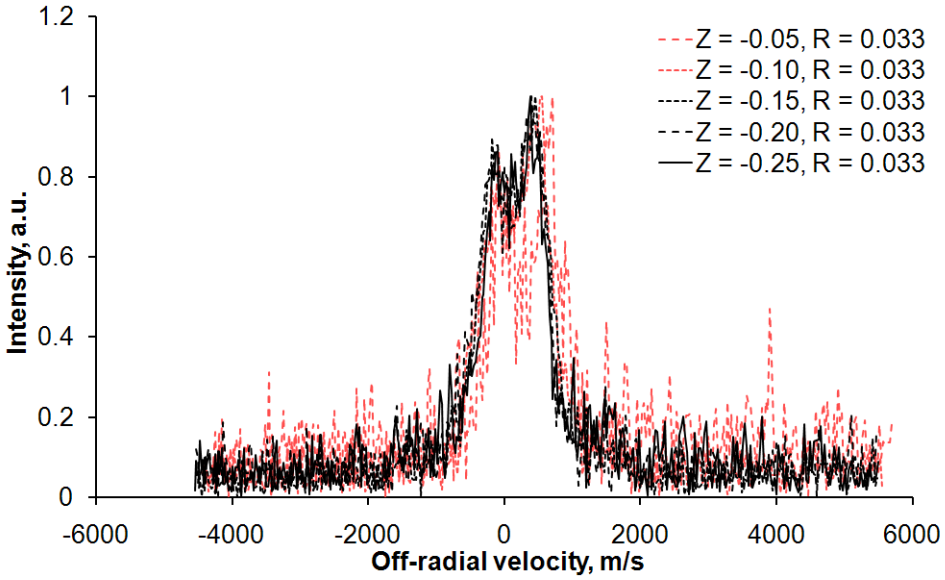

Figure 5-73. Raw off-radial neutral velocity distribution functions along the inner wall of the 6-kW Hall thruster at 300 V, 20 mg/s.

Note that the double-peak feature is an artifact of the Zeeman effect, which is largest near the inner wall where the local magnetic field strength is the highest.

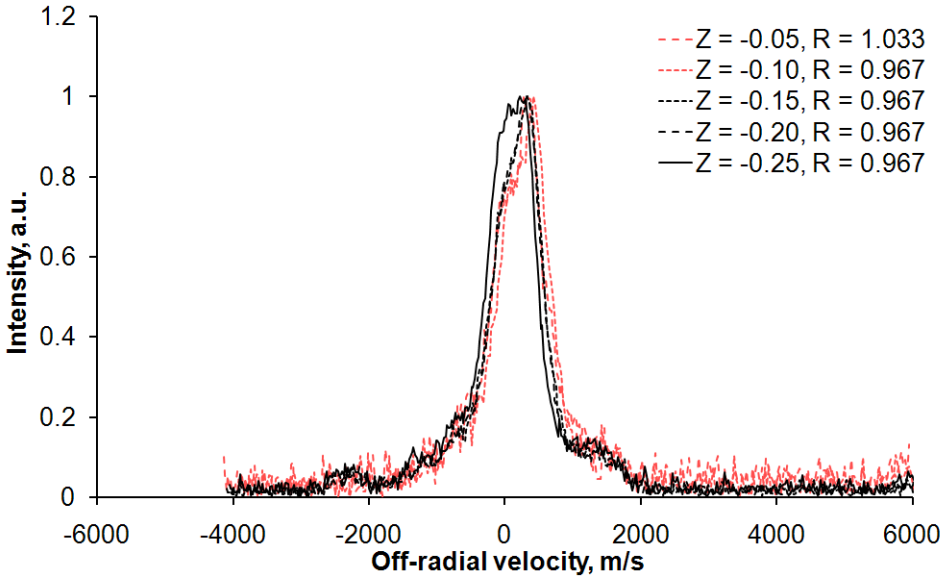

Figure 5-74. Raw off-radial neutral velocity distribution functions along the outer wall of the 6-kW Hall thruster at 300 V, 20 mg/s.

As the various raw LIF traces show, no high-speed neutrals were detected during the two-axis near-wall study. There is a lack of detected high-speed neutrals despite the fact that near-wall physics suggest some level of charge exchange effect should be taking place and slow ions that could possibly be products of charge exchange have been found in the ion LIF traces. One possible explanation is that these high speed neutrals are simply outside of the velocity range that was scanned. Another possibility is that given the high neutral density and low ion density in the near-wall regions, the amount of fast neutrals born from charge exchanging is simply not large enough compared to the slow neutrals that already exist and do not show up very well on the LIF scans.

Although high-speed neutrals were not found in this dissertation, they should not be discounted as being unimportant for future works. It is speculated that aside from direct bombardment of the channel walls, erosion can also occur from indirect bombardment. Here, indirect bombardment refers to particles that undergo direction changing collisions before arriving at the walls. As work by Mikellides, et al., [11] have shown, Hall thruster channel walls can potentially be eroded to a point where the erosion rate becomes very low relative to beginning-of-life erosion rate. Furthermore, profilometry data from older generation thrusters like the SPT-100 [5] have shown that the erosion physics generally have two distinct regime, a high varying erosion rate regime near the beginning-of-life and a low near-constant erosion rate regime towards the end-of-life. Figure 5-75 is a reproduction of the erosion rate plot from [5]. One hypothesis currently being considered is the possibility that the beginning-of-life erosion is dominated by direct bombardment while the end-of-life erosion is dominated by indirect bombardment. That is to say, the channel profile eventually reaches a point where the contribution to the erosion rate from the direct impingement of ions on the channel walls becomes negligible. The study of near-wall neutrals is important for

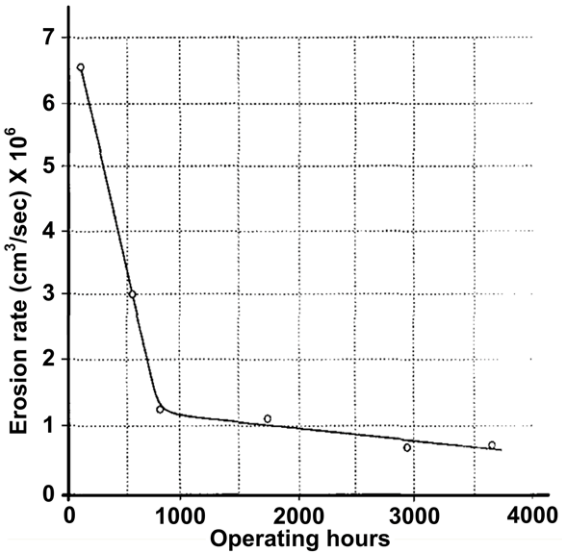

Figure 5-75. Erosion rate versus operating hours for the SPT-100 [5].

proving or disproving this hypothesis because high-speed neutrals re-directed via collisions towards the walls are much easier to distinguish from the background low-speed neutrals than re-directed high-speed ions from background high-speed ions.

With all of this said, since the data does not confirm nor rule out the presence of high-speed neutrals, their characterization is left for future studies. For the purpose of calculations made in this dissertation, neutral xenon will be assumed to make negligible contributions to the erosion of the channel walls.

Having obtained data on the singly-charged xenon ions that bombard the channel walls of the 6-kW Hall thruster, we now move on to studying the sputtered boron product.

Chapter 6
Physics of Boron Spectroscopy

6.1 Basics of Boron Ground State Transition Lines

As a Hall thruster diagnostic, boron spectroscopy was first applied in the form of emission spectroscopy to study erosion rate. A typical experimental setup involves monitoring the thruster plume emission at one of the ultraviolet wavelengths associated with the transition of excited-state boron to ground-state boron [22, 84]. Unfortunately, the results of emission spectroscopy are generally not accurate without a good collisional-radiative model to correlate the measured emission to relative state densities. Such a model does not currently exist for the Hall thruster environment owing to the complex nature of the interactions in the Hall thruster plasma and the difficulties in obtaining data for developing such a model.

To the author's best knowledge, the first published use of laser-based boron spectroscopy in a Hall thruster environment is the experiment that will be described in Chapter 7. Prior to this experiment, Prof. Azer Yalin's group, our collaborator, had successfully demonstrated the use of the cavity ring-down spectroscopy (CRDS) technique to study the erosion of an anode-layer-type Hall thruster using lines for various metal atoms [24, 85, 86].

Table 6-1 shows the two boron ground state transition lines that are of interest to this dissertation. Since CRDS is based on absorption spectroscopy, we are primarily

Table 6-1. Boron ground state absorption lines of interest to this dissertation.

Species	Ground state	Ground state energy, eV	Excitation wavelength (vac), nm	Excited state	Excited state energy, eV
B I	$2p[1/2]_2$	0	249.752	$3s[1/2]_2$	4.964289
B I	$2p[3/2]_4$	0.001895	249.848	$3s[1/2]_2$	4.964289

interested in the absorption characteristics of the boron lines. The B I 249.752 nm and B I 249.848 nm transitions are doublets. However, it is difficult to scan both lines with laser-based spectroscopy in actual testing, so one or the other is chosen. For the CRDS experiment described in this dissertation, the 249.848 nm transition is chosen because it is typically twice as strong as the neighboring 249.752 nm transition. The relative strength of these two transitions, which have almost identical energy, is based on the ratio of degeneracy between the two lines. Since the degeneracy of the $2p[3/2]_4$ ground state, which is 4, is twice as large as the degeneracy of the $2p[1/2]_2$ ground state, which is 2, the former houses roughly double the population density when compared to the latter if the populations are in thermal equilibrium.

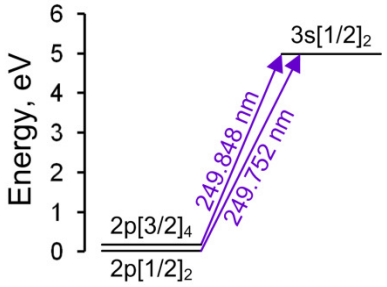

Figure 6-1. Transition diagram for the boron transition used in the CRDS experiment.

Figure 6-1 shows the transition diagram for the transitions shown in Table 6-1. The energy level for the two ground states had to be drawn with an out-of-scale separation between them in order to denote that they are two separate states.

6.2 Line Broadening Mechanism for Ground State Boron

6.2.1 Doppler Shift and Broadening

The velocity distribution function of sputtered particles from bombardment of single-component target is a well-studied subject. This distribution function is described by the Sigmund-Thompson distribution [87, 88], shown in equation (6-1),

$$f(u) \sim \frac{u^3}{\left(u^2 + v_b^{\,2}\right)^{3-2m}} \cos\phi \tag{6-1}$$

$$v_b = \sqrt{\frac{2E_b}{M}} \tag{6-2}$$

where u is the velocity, v_b can be calculated from the binding energy, E_b, using equation (6-2), ϕ is the polar angle of the emitted particle with respect to the surface normal, m is an energy-dependent parameter and is close to 0 for low inbound energy (<1keV), and M

is the mass of the sputtered particle. The binding energy is a constant that depends on the material and the surface condition. This distribution is insensitive to bombardment angle and not very sensitive to bombardment energy at low bombardment energy (<1keV).

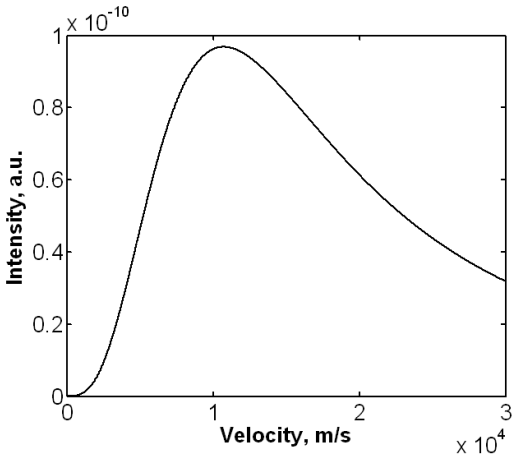

Figure 6-2. Example Sigmund-Thompson distribution with m = 0.17, E_b = 5 eV.

For a typical magnetic-layer Hall thruster, the channel walls are made of boron nitride, which is a multi-component material. Active research is being carried out at Colorado State University to determine how sputtered boron behaves when boron nitride is bombarded by low energy xenon [89]. Preliminary work shows that for xenon bombardment energy of 1200 eV, a Sigmund-Thompson distribution with m = 0.17±0.05 and E_b = 5.0±0.5 eV fits the VDFs obtained in a controlled environment via LIF velocimetry very well [89]. However, it is also pointed out in the same work that the value of the fitted binding energy is very sensitive to the value of m, and that an m of 0-0.3 all produce reasonable fit to the data. Furthermore, the average velocity of the Sigmund-Thompson distribution is also very sensitive to m. If we use the limit at which the Sigmund-Thompson distribution is supposed to be applicable as the limit of integration, the numerically integrated average velocity of the m = 0.17, E_b = 5 eV distribution is ~23 km/s. The aforementioned limit is described by equation (6-3) [88],

$$T_{max} = \frac{4M_1M_2}{(M_1 + M_2)^2} E_{ion} \tag{6-3}$$

where T_{max}, the limit, is also the maximum energy transferred in an elastic head-on collision, M_1 is the mass of the bombarding particle, M_2 is the mass of the sputtered particle, and E_{ion} is the energy of the bombarding particle. The last variable has the subscript "ion" because sputtering experiments are typically performed with mono-energetic ion beams.

Figure 6-2 shows an example Sigmund-Thompson distribution with m = 0.17, E_b = 5 eV. In the Hall thruster environment, the energy of the bombarding particles range

from 0 to a few hundred eV. Whether the sputtered particle distribution function still follows the Sigmund-Thompson distribution is an open area of research. The current hypothesis is that they do based on theories and observed behaviors from the bombardment of single-component metallic targets, which is summarized in [88].

In theory, it is possible to combine the sputtered boron velocity distribution function with angular energy distribution function of the bombarding particles, the geometry of the setup, and differential sputter yield data to calculate the average axial velocity of the boron leaving the Hall thruster channel.

Since some of the boron that exits the channel will also carry a radial component, the CRDS result, which is obtained from shooting a laser beam across the channel exit of the thruster, will be Doppler broadened. The exact amount of Doppler broadening is an open topic of research, and is expected to be no greater than the spread in the Sigmund-Thompson distribution. This value can potentially range from tens of GHz to hundreds of GHz depending on the exact angle at which the sputtered boron is traveling with respect to the CRDS laser beam.

6.2.2 Hyperfine Structures

Table 6-2 lists the nucleus-electron interaction constants and the isotopic shift for the 249.848 nm transition. The author was not able to locate published values for the magnetic dipole interaction constant, A, for the upper state of the 249.848 nm transition. However, a survey of similar B I lines shows that values of this constant are at most a few hundred MHz [91]. Assuming the upper state of the 249.848 nm transition has A values in a similar range, the total broadening due to hyperfine structure, with Zeeman effect at a few hundred Gauss, should be no more than 6 GHz. This broadening is dominated by the isotopic shift. Since the Doppler broadening is expected to be at least a few tens of GHz, the hyperfine structure broadening will most likely be not observable.

Table 6-2. Hyperfine structure constants for the 249.848 nm transition.
The values of these constants are from the references [90-92].

Isotopes	A of lower state, MHz	B of lower state, MHz	J_L	J_U	Isotopic shift, GHz
B-10	23.3	-	3/2	1/2	-5.2
B-11	73.3496	2.6927			0

Note that to calculate boron density for CRDS data one simply integrates over the lineshape. The exact amount of broadening in the lineshape is not important as long as the majority of the portion is included in the integration domain.

6.3 The Excited States and the Singly-Charged State

Neutral boron has a number of excited states which are all ultimately connected to the ground state, with one exception, which will be mentioned later. The longest life time of these excited states is 23 ns of the $2s2p^2\ ^2D$ level [93]. At an average particle speed of 23 km/s, any neutral boron that sputters off of the channel wall in an excited state will have a characteristic decay distance of ~0.5 mm. It is currently unclear how much of the sputtered neutral boron are ejected in the excited state but they will all have decayed to the ground state a few mm from their ejection site assuming no other process is causing further excitation.

There is one neutral boron level that is not connected to the boron ground level, and it is the $2s2p^2\ ^4P$ level. The peculiar thing about this level is that it is only connected to a few boron ion levels and that this level plus those ion levels have not been experimentally observed to connect to the rest of the boron system [91]. In short, it is unclear whether or not sputtered boron can be born into this state. For the time being, it is simply assumed that this state is not accessible by sputtered boron from Hall thruster channel walls.

The ionization energy of neutral boron is 8.2980 eV [93]. There is very little data on whether a significant amount of the sputtered boron comes off as boron ions. A 1997 work by Zhang, et al., [94] is the only study the author is aware of that tackles this subject. In this study, it was found that the amount of ion produced strongly depended on the neutralization of the BN sample being bombarded by a mono-energetic ion beam. The approach of Zhang, et al., was to tweak the electron current until the B^+ signal is maximized in their secondary ion mass spectrometry reading. The most pertinent result of the said study to this dissertation is Figure 3 from [94], which is reproduced as Figure 6-3 in this dissertation for reference. This figure shows how rapidly the B^+ ion signal drops with decreasing ion beam energy. At a bombardment energy of 100 eV, the secondary ion mass spectrometry signal is over 3 orders of magnitude lower than when the

bombardment energy is 1000 eV and about 1 order of magnitude lower than when the bombardment energy is 300 eV.

From the LIF velocimetry data described in Chapter 5, adding in ~30 eV sheath energy [61], we get about ~130 eV xenon ions bombarding the channel walls at the channel exit for 300 V operating conditions. At 600 V, 10 mg/s, the ion energy goes up to 300-400 eV. Since the 1997 work by Zhang, et al., does not compare the amount of boron ion to boron neutral, there are two likely scenarios for what is happening when BN channel walls are bombarded by Hall thruster xenon ions. One, the amount of boron ion production is negligible across the board, and that even at 600 V operation, with an order of magnitude increase in ion production when compared to 300 V operation, the amount is still negligible. Or two, the amount of boron ion production is negligible for 300 V operation but not for 600 V operation. If in fact boron ion production can become significant at higher bombardment energy, then CRDS will increasingly under-predict the erosion rate as the bombardment energy is increased. There is the third scenario where boron ion production is important across all bombardment energy, however, this scenario does not fit with known sputtering theories [88] and experimental data [95]. For the purpose of data analysis, scenario one is assumed to be true.

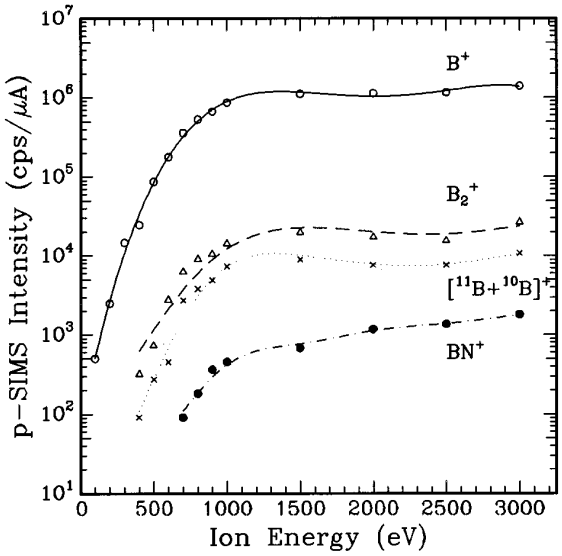

Figure 6-3. Figure 3 reproduced from [94] showing how the number B+ ion drops with decreasing bombardment energy.

Thus all of the boron that comes off of the Hall thruster channel walls is assumed to be in the ground state when they are detected by the CRDS sensor. Whether the presence of the Hall thruster plasma could change the ratio of population between the excited and ionized states to the ground state will be discussed in Section 7.4.5.

Chapter 7
Cavity Ring-Down Spectroscopy

7.1 Principles of Cavity Ring-Down Spectroscopy

Cavity ring-down spectroscopy is a type of absorption spectroscopy. A great review of the topic and recent development of CRDS as a diagnostic can be found in [96]. CRDS is based on the principle that by sending a beam of photons through an absorber and measuring the change in intensity, one can obtain the absolute line-integrated density of the absorber. CRDS differs from traditional absorption spectroscopy in that a cavity, formed from a pair of highly reflective mirrors, is used to trap light. Instead of one pass, the trapped photons make many passes through the absorbing medium. Each time the trapped light strikes one of the mirrors, a small fraction of it leaks out. The intensity of the leaked light is measured by a detector. When the light source is a continuous-wave laser, the beam is extinguished (e.g. with an acousto-optical modulator), and then the light within the cavity will begin to decay. The characteristic time constant of the decaying signal measured by the detector can then be used to calculate the absolute line-integrated density of the absorber inside the cavity. Figure 7-1 illustrates the operational principles of CRDS. The grey arrow pointing from the detector to the acousto-optical modulator (AOM) indicates some form of feedback so the AOM knows when to cut off the laser beam.

Equations (7-1) and (7-2) govern the ring-down behavior and time constant in a

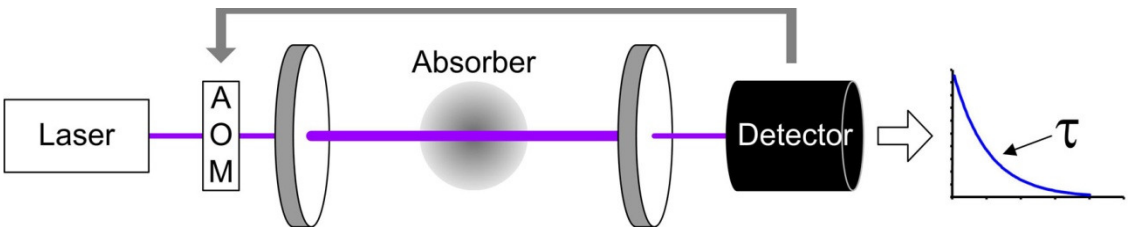

Figure 7-1. Operational principles of cavity ring-down spectroscopy.

typical cavity,

$$S(t,\upsilon) = S_0 \exp\left[-t/\tau(\upsilon)\right] \tag{7-1}$$

$$1/\tau(\upsilon) = \frac{c}{L}\left[\int k(x,\upsilon)dx + (1 - R_M)\right] \tag{7-2}$$

where S is the signal intensity, S_0 is the initial signal intensity, τ is the characteristic time of decay, also called the ring-down time, c is the speed of light, L is the cavity length defined as the distance between the two mirrors, R_M is the mirror reflectivity, and $k(x,\upsilon)$ is the absorption coefficient at a given laser frequency, υ, and position along the cavity optical axis, x, and the integral term represents the line-integrated contribution from the absorber. In practice, the mirror reflectivity changes in time and there is a scattering effect that decreases the effective reflectivity. To extract absolute line-integrated density, the ring-down time is compared to the empty cavity ring-down time, τ_0. This is shown in equation (7-3),

$$Abs(\upsilon) \equiv \int k(x,\upsilon)dx = \frac{L}{c}\left[\frac{1}{\tau(\upsilon)} - \frac{1}{\tau_0}\right] \tag{7-3}$$

where Abs is the sample absorbance and is unitless. By scanning the laser frequency across an entire transition line, one can then integrate for the total absorbance and relate it back to the absolute density of the lower state of the transition being studied. Equation (7-4) shows this relationship,

$$\int n_1 dx = 8\pi \frac{g_1}{g_2} \frac{\upsilon_{21}^{\,2}}{A_{21} c^2} \int Abs(\upsilon)d\upsilon \tag{7-4}$$

where n is absolute density, g is level degeneracy, υ is transition frequency, A is the spontaneous emission rate (called the Einstein A coefficient, not to be confused with the magnetic hyperfine constant), subscript 1 denotes lower state, and subscript 2 denotes upper state. Table 7-1 lists the values of the various constants for the 249.848 nm transition.

Table 7-1. Transition parameters for the B I 249.848 nm transition [93].

Constant	Value
g_1	4
g_2	2
υ_{21}	1.19990e15 Hz
A_{21}	1.67e8 s^{-1}

A typical Hall thruster can usually be assumed to be axisymmetric (there are cases where this is not true, e.g. when an external cathode perturbs the axisymmetry [97]). To obtain erosion product density as a function of

radial position, we begin by taking CRDS density measurements across multiple P positions. Here, the P-axis is orthogonal to the beam axis and parallel to the channel exit plane of the thruster. Greater detail is available in Section 7.3.1. These measurements can then be converted to a function of radial position via Abel inversion. See reference [98] for more information on Abel inversion. The exact inversion algorithm used in this dissertation is called the "onion-peeling" method, which is like Abel inversion but is mathematically much simpler for a discretized data set of relatively small size. A comparison of various inversion methods including Abel and onion-peeling can be found in [99]. The onion-peeling method is described in detail in Section 7.4.2.

The cavity used in CRDS is also called a resonator because only light of the right frequencies can resonate with the cavity and build up in the cavity. In the simplest analysis, the governing equation for the propagating frequencies in a planar cavity can be found by trying to fit a standing wave into the cavity. Since the mirrors are near lossless, the boundary condition at the mirror must correspond to a nodal point in the standing wave. These conditions give rise to equation (7-5) [100],

$$\upsilon_q = q\frac{c}{2L}, \quad q = 1, 2, 3, \ldots \tag{7-5}$$

where υ_q is the frequency of the q-th resonance mode and can only be certain discrete values in a perfectly lossless cavity. In a lossy cavity, transmission can happen over a narrow range of frequencies centered on each resonance frequency. The FWHM value of each transmission peak is described by equations (7-6), (7-7), and (7-8) [100],

$$\delta\upsilon \approx \Delta\upsilon_F / \underline{F} \tag{7-6}$$

$$\underline{F} = \frac{\pi\sqrt{R_M}}{1 - R_M} \tag{7-7}$$

$$\Delta\upsilon_F = \frac{c}{2nL\cos\theta} \approx \frac{c}{2L} \tag{7-8}$$

where $\delta\upsilon$ is the spectral width of the individual resonance mode, $\Delta\upsilon_F$, also called the free spectral range, is the frequency spacing between resonance frequencies, $\underline{F}$ is the cavity finesse, R_M is the mirror reflectivity, n is the refractive index of the medium between the mirrors, and θ is the angle of incidence with respect to the mirror surface normal. Both n and cos θ are typically equal to 1. From these equations we can see that the finesse

increases rapidly as reflectivity approaches 1. Note that $R_M = 1$ represents a lossless mirror. From equation (7-2), we see that we want the mirror to have a reflectivity as close to 1 as possible because then the $(1-R_M)$ term becomes negligible and the ring-down time is sensitive only to the absorber. Furthermore, as the finesse increases, the spectral width decreases and one can obtain higher resolution scans.

Our approach for BN erosion rate measurement is to scan the ground state boron transition at 249.848 nm to measure boron density. Details on this transition are described in Chapter 6. For this wavelength, commercial mirrors can achieve a typical reflectivity of ~0.997, which corresponds to a finesse of ~1000. This reflectivity is poor compared to that available for reflectors in the visible and infrared, which makes sensitive detection challenging compared to standard CRDS. Equation (7-9) [96] shows the two primary attributes of a CRDS sensor that determines its minimum detection threshold.

$$\left[\mathrm{Abs}(\upsilon)\right]_{\min} = \left(1 - R_M\right)\left(\frac{\Delta\tau}{\tau}\right)_{\min} \tag{7-9}$$

where $\Delta\tau$ is the uncertainty in the measured ring-down time. To improve the sensor's sensitivity, one can either increase the mirror reflectivity or decrease the ring-down time measurement uncertainty. For a cavity length of ~1 m, υ_F equals 150 MHz, so $\delta\upsilon$ equals 150 kHz. It is difficult for a commercial laser to achieve a frequency width that narrow (typical diode laser can achieve ~1 MHz), so only a portion of the input laser beam power couples into the cavity. Increasing mirror reflectivity increases measurement sensitivity but can also decrease the beam power that can be coupled into the cavity, which decreases the signal. Furthermore, the mirror reflectivity, the laser power, and the laser linewidth are limited by technology. Alternately, one can improve the CRDS sensor's sensitivity by decreasing the ring-down time measurement uncertainty. This uncertainty is limited by the amount of signal compared to the amount of noise, and the size of the data sample. Typically, the sample size is the easiest to influence and is determined by the ring-down rate, the sampling rate, and the sampling time.

The governing equations for cavity transmission shown in this dissertation are for the TEM_{q00} modes, which are called the fundamental modes. Other non-fundamental modes of transmission are also possible and require wave-optics treatment of the problem, details of which can be found in [100]. For our purpose, we want to couple light

into only the fundamental modes because higher peak intensity can be achieved with those modes than non-fundamental modes. To be more precise, we are coupling into only one of the fundamental modes at any given time because equation (7-5) can only be satisfied for one value of q if the laser is monochromatic. Thus, the TEM_{q00} modes are often collectively referred to as the TEM_{00} mode. The process of coupling light into only the TEM_{00} mode is called mode-matching, the theory of which is described in [101]. Details of how mode-matching is done for the CRDS experiment is described in Section 7.3.4.

7.2 CRDS Sensor Prototype: Challenges and Lessons Learned

This section describes the first CRDS sensor that was tested as a part of this dissertation. This sensor is also known as the version-one sensor or the prototype sensor. Information contained in this section is meant to help others to design their own CRDS sensor from scratch, and to learn from mistakes made with the version-one sensor. Prior to the construction of the prototype sensor, the challenges for the application of CRDS to measuring Hall thruster erosion rate were sorted into two categories. The first category involves thermal management and alignment maintenance during testing. The second category involves proper timing for various parts of the data acquisition train and the how best to maximize the sample size in order to maximize the sensor sensitivity.

7.2.1 Thermal Challenges and Design

The design process started with the assumption that we would be using two mirrors the reflective sides of which are slightly concaved. The radii of concavity for the two mirrors are the same and are slightly larger than the cavity length. Geometric optics calculations show that such a configuration is stable as a resonator [101]. In other words, light is properly trapped in the cavity. The laser light would be injected directly into the cavity through a window port in the side of the vacuum chamber. This “direct beam injection” delivery method was chosen because as of the time of writing, optical fibers that have small enough core diameter to couple light into the cavity properly only last a few tens of hours due to UV-induced damage when transmitting 250-nm light. Using direct beam injection does increase certain alignment challenges. To understand the

thermal-related challenges, we begin by deriving the equations governing the maximum allowable drift speed in cavity length.

Starting from equation (7-5), we can derive the condition for the transmission of light through a cavity of length L in terms of the laser wavelength. This is shown in equation (7-10),

$$L = q\frac{\lambda}{2}, \quad q = 1, 2, 3, \ldots \tag{7-10}$$

where λ is the laser wavelength and is approximately fixed and q is the same mode number as in equation (7-5). We can see from this equation that resonance occurs whenever the cavity length is some half-multiple of the wavelength. To a first order approximation, which is accurate if $\delta\upsilon << \upsilon_F$, we can combine equations (7-6), (7-8), and (7-10) to form equation (7-11),

$$\delta L = \frac{\lambda}{2}\frac{1}{F} \tag{7-11}$$

where δL is the allowable variation in the cavity length such that the transmitted light intensity is equal to or greater than half of the peak transmission. If we assume the laser beam is roughly Gaussian in frequency space (a good approximation for diode laser beam profiles), we can approximate the maximum allowable variation in the cavity length while maintaining resonance condition as measured by percentage of light transmission. In other words, we want to find the condition for avoiding the loss of light from the cavity or distortion to the ring-down due to the cavity going out of resonance during the ring-down. For a 1-m cavity, R_M ~0.997, maintaining light transmission at 98% (±1% variation in cavity transmission) requires the cavity length to stay within $4.1*10^{-11}$ m over the course of the ring down. Assuming five times the ring-down time yield sufficient data for curve-fitting, the maximum allowable drift speed in cavity length is $8.2*10^{-6}$ m/s. This calculation gives the design criterion for the short-term stability of the cavity structure.

Additionally, geometric optics was used to calculate the rough allowable deviation in injection angle to maintain cavity resonance. A detailed explanation of how to use geometric optics can be found in reference [102]. The maximum allowable deviation in injection angle is estimated to be ~0.3° for the 1-m cavity with 1" diameter R_M ~0.997 mirrors. This angle drives the design criterion for the long-term stability of the

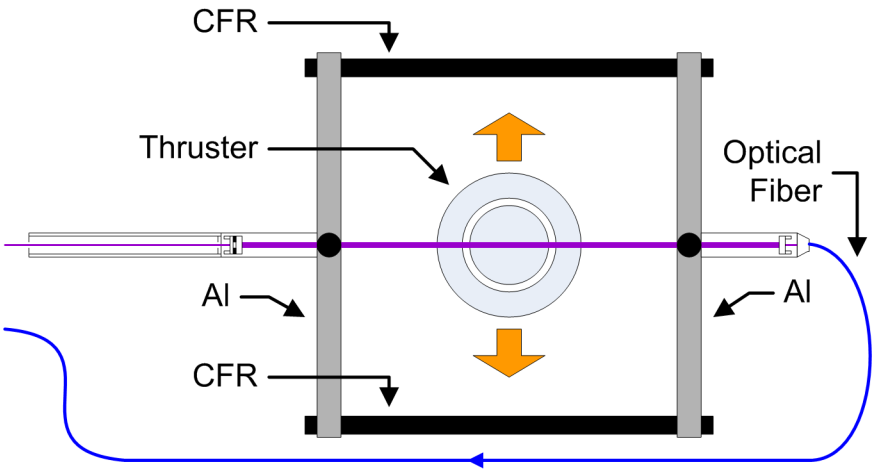

Figure 7-2. Diagram of an initial design for the CRDS sensor.

cavity structure as well as the dimensions of the steering mirrors outside of the chamber needed to compensate for drifts.

One of the main challenges in performing CRDS measurement on a medium-to-high power Hall thruster is to keep the thermal expansion of the cavity low enough to maintain optical alignment. The solution for the management of the thermal expansion on the prototype sensor was the use of a combination of carbon fiber rod (CFR) and aluminum members as illustrated in Figure 7-2. The four members shown in Figure 7-2 (labeled as CFR and Al) formed the physical cavity (as opposed to the optical cavity formed by just the two mirrors). Carbon fiber rod can be manufactured to have a very low coefficient of thermal expansion (~0.1 ppm/K). They are light weight and have about ten times the ultimate yield strength as aluminum. However, their thermal conductivity is only $1/20^{th}$ that of aluminum and they are not as rigid. CFR was used as the bars that keep the cavity mirrors apart at a fixed distance. The joints and connecting pieces were made of aluminum to help dissipate heat. There were two pivot joints located at the midpoint of each aluminum connection in Figure 7-2. Each pivot is connected to a combination CFR/Al structure to keep the height of the cavity fixed relative to the chamber ground. Laser light is injected through a vacuum chamber window directly into the cavity. The cavity's optical axis passes right in front of the exit plane of the thruster. Signal is collected via a multi-mode optical fiber, which, due to having a large core diameter, does not suffer from life time issues. The thruster is mounted on a vertical motion stage and is moved up and down.

In the setup shown in Figure 7-2, as the thruster is moved, one side of the cavity will become more heated than the other. This effect will create a temperature differential between the upper CFR and the lower CFR. From the maximum allowable deviation in injection angle, we calculate that light can still be coupled properly into the cavity for temperature differentials of up to 500 K.

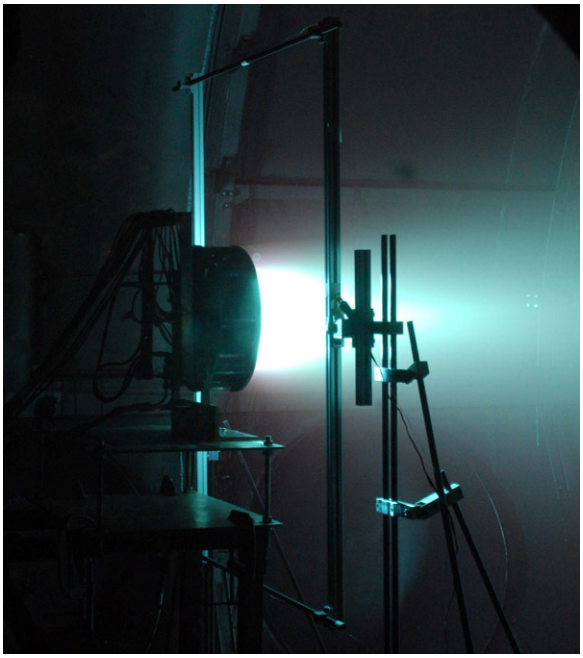

Figure 7-3. Photograph of the 6-kW Hall thruster operating during the CRDS prototype sensor thermal stress test.

To figure out whether or not thermal management will become an issue, a thermal stress test was performed. The cavity structure, without any optics attached, and the 6-kW thruster was put into the vacuum chamber in the same position as they would be during a real CRDS test. The thruster was then operated at nominal condition for two hours and the temperatures at the center of each of the four members that make up the physical cavity are measured via thermocouples. Figure 7-3 shows a photograph of the 6-kW Hall thruster operating during this thermal stress test. The result of this short test showed that the temperature rose at most 6 K in the first half hour and 20 K over the full two hours of thruster operation. The equivalent thermal expansion speed of the upper and lower CFRs is $3*10^{-10}$ m/s, which is four orders of magnitude lower than the maximum allowable cavity drift speed. As it will turn out, aluminum members can be used in place of the CFR in the physical cavity. Using aluminum members and assuming the same temperature rise gives a thermal expansion speed of $7*10^{-8}$ m/s, which is still two orders of magnitude lower than the maximum allowable drift speed.

On the other hand, the CFR members that form the two legs supporting the physical cavity should be kept as CFR if possible. Thermal expansion calculations show that for the temperature rise exhibited during the thermal stress test, the physical cavity would have moved upward 1-2 mm if the legs were made of aluminum. While this can be

compensated by the steering mirrors outside the vacuum chamber, it is not preferable to have to tweak those mirrors after every scan.

Lastly, though not a thermal issue, it should be noted that the shape of the vacuum chamber changes ever so slightly when it is pumped down. If the shift does change the relative position of the two legs for the cavity, then it can be compensated for by tweaking the steering mirrors outside of the vacuum chamber. However, if the shift does change the relative position of the two legs, it is possible that no compensation can be done. If that happens, motors need to be installed behind the cavity mirrors to allow remote tuning of the cavity. For the prototype, it was decided that we would first make an attempt without motorizing the mirrors as they add extra complexity. The motors would be added later if the cavity goes out of alignment in a way that the external steering mirrors cannot compensate for.

7.2.2 Temporal and Statistical Challenges

This section identifies key parameters in the electronics and data-acquisition (DAQ) system design. The empty cavity ring-down time for a pair of R ~0.997 mirrors is ~1 μs. The ring-down time with absorber will therefore be less than 1 μs. This time constant drives a number of important timing issues in the design of a CRDS diagnostic.

First, in order to properly record ring-down traces the DAQ sampling rate needs to be at least an order of magnitude faster than the inverse of the empty cavity ring-down time, and faster still if the absorbers drive the ring-down time down. A 1-μs ring-down time corresponds to a sampling rate requirement of greater than 10 MS/s. An ATS460, 2-channel, 14-bit, 125-MS/s digitizer card was selected to give some margin.

Second, a ring-down starts when the laser beam input to the cavity is cut off. The cutoff time has to be short compared to the ring-down time. If it is not, there will be a portion of the data where the laser beam is still being cut off and has to be removed during data reduction. This portion also corresponds to the part of the data with the highest signal-to-noise ratio. In order to perform fast cut-off, we employ a Neos 35110-2-244-BR acousto-optical modulator (AOM), which can cut off a 1-mm diameter beam in ~100 ns.

Third, the speed of the electronic circuit that detects whether the cavity is filled and signals the AOM to begin ring-down is important. If the circuit takes too long, the wavelength and the cavity length will have moved out of resonance by the time the ring-down begins. As such, the voltage comparator and pulse generator used to determine when ring-down begins must have a response time that is very short compared to the ring-down time constant. We chose an LM361 voltage comparator with a rise time of 20 ns in a Schmitt trigger configuration coupled to a BK Precision model 4030 analog pulse generator with an insertion time of 60 ns to accomplish this task. The choice of the Schmitt trigger configuration adds hysteresis to the circuit response to stabilize that response against noise.

Fourth, the ring-down time sets a limit on how fast one can sweep the cavity. In order to fill the cavity with light, one can either sweep the laser wavelength or sweep the cavity length. However, if the sweep is too fast, the laser light may not couple into the cavity properly and the wavelength and the cavity length may move out of resonance before the ring-down is complete, causing premature ring-down termination [103]. For the prototype sensor, open-loop laser sweeping was used. The maximum wavelength sweep speed is not driven by the same requirement as the maximum cavity length drift speed. This is because assuming the cavity is filled at the moment and that ring-down has begun (AOM has finished cutting off the beam), it no longer matters how fast the wavelength is drifting, that light is not reaching the cavity. As long as the cavity length is stable during the ring down, there will be no excess light loss due to cavity going out of resonance. The real criterion that drives the maximum wavelength sweep speed is the speed with which the cavity can be filled before the laser wavelength moves out of resonance. As it turns out, the characteristic time of cavity filling is the same as the cavity ring-down time. Assuming one wishes to spend ten times the cavity ring-down time within the spectral width, $\delta\upsilon$, of the transmission peak the corresponding maximum data rate, when triggering only once per transmission peak, is ~100 Hz. However, in practice this number can be much greater if multiple triggering happens during one transmission peak and/or the cavity length is also oscillating, e.g. due to vibration in the chamber.

Going back to an ideal situation, there is a lot of dead time between ring-down when the wavelength and the cavity length is out of resonance. We can reduce the dead

time between resonances by switching to a close-loop scheme. The idea is to use the detector response as a feedback signal to keep the cavity length and the wavelength in lock. The close-loop scheme is more complicated as locking of the laser and cavity is usually accomplished with piezoelectric element mounted to one of the cavity mirrors. However, the data rate can be improved by orders of magnitude since one can almost completely eliminate the dead time between transmission peaks. Further detail on various CRDS scanning schemes can be found in the review paper by Berden, et al. [96]. Close-loop scanning was not implemented in this dissertation due to the added complexity. It was decided that attention should be focused on first demonstrating the viability of using CRDS to measure boron density and that further improvements can be left for future work.

Having now reviewed the temporal challenges, we will look briefly at the related statistical challenges. To put it bluntly, we did not know whether the SNR would be high enough to allow statistically meaningful measurements of boron density prior to testing. Initial rough calculations predicted that an SNR of 75 is achievable with an integration time of 10 s (i.e. take 10 s worth of data) [95]. However, we did not know how practical it is to spend 10 s on a single data point when a hundred is typically needed for a single trace. Furthermore, it was not known a priori how much and how fast the mirrors would degrade during testing.

The strength of the signal in the CRDS data is driven by the amount of light coupled into the cavity and the detector sensitivity. Transmission loss through the collection optical fiber was determined to be negligible and the core diameter of the fiber was big enough that all light exiting the cavity is properly captured. The amount of laser light available is limited to ~10 mW by the laser system, which is a Toptica TA-FHG-110 frequency quadrupling diode laser system. The actual amount of light coupled into the cavity is ~30 μW because the same high reflectivity mirrors that form the optical cavity also reflect the majority of the laser light away. The choice of detector is the Hamamatsu R9110 photomultiplier, which combines an excellent quantum efficiency of ~29% at 250 nm with an anode rise time of 2.2 ns.

The noise in the CRDS data comes from electronic and shot noise, background fluorescence, scattering, blackbody radiation, and possibly laser phase noise. Phase noise

is described in much greater detail in reference [103] and generally limits the uncertainty to ~1%. Blackbody radiation is typically very low in intensity in the UV range and easy to remove by installing irises along the beam axis as well as shielding the signal return line from stray light. Scattering effects can be removed by taking ring-down scans with the laser detuned from the transition wavelength. Any scattering can then be rolled into the effective mirror loss term. It is unclear how much electronic and shot noise there is in the system but care is taken to choose low-noise equipment powered by clean power supplies. More details about the equipment, wiring, and power supplies can be found in Section 7.3.2. Background fluorescence could potentially be one of the main sources of noise because the plasma is very energetic and radiates a large amount of energy as light. Luckily, most of the xenon ion and neutral transition wavelengths are pretty far away from the target boron lines. The nearest xenon transition is an ion line at 249.253 nm, but it is a fairly weak line and is located 2.4 THz away from the Boron lines of interest. The nearest nitrogen transition is an ion line at 249.758 nm, but its lower state is a ~21 eV excited state and we do not expect such a state to be well-populated. Nevertheless, the 249.848 nm boron transition may give a more accurate measurement than the 249.752 nm boron transition. Related CRDS work to measure manganese has shown that a simple interference filter is sufficient to block unwanted fluorescence [24]. A band-pass interference filter with a 10-nm bandwidth centered at 250 nm is chosen for filtering out fluorescence.

7.2.3 Lessons Learned

Figure 7-4 shows a photograph of the prototype CRDS sensor with various components pointed out. The injection and collection optics train are supported by four CFR each. However, the weight of the optics and the long moment arm of the injection train required the addition of a triangular support attachment to the cavity, shown on the left side of the photograph. The mirrors are behind the attachment points between the physical cavity and the two legs so it cannot be seen on the photograph. Each leg sits on three Urethane vibration isolators, which are also not shown. There is a T-shape spring attachment located bellow the thruster platform that is designed to catch the thruster if the vertical motion stage loses power and goes into free fall. Should this happen, the encoder on the motion stage will keep track of the thruster position.

A number of valuable lessons were learned while operating the prototype sensor. The first discovery was that the mirrors degrade very rapidly when exposed to the Hall thruster plasma even with narrow beam irises installed in front of the mirrors. The ring-down time constant dropped from ~0.5 μs to ~0.23 μs during initial thruster bakeout, which was done at 300 V, 20 mg/s for 1 hour. The thruster was immediately turned off afterwards but no recovery in the ring-down time was observed. Attempts were then made to take data at the operating conditions 300 V, 20 mg/s and 600 V, 10 mg/s. An

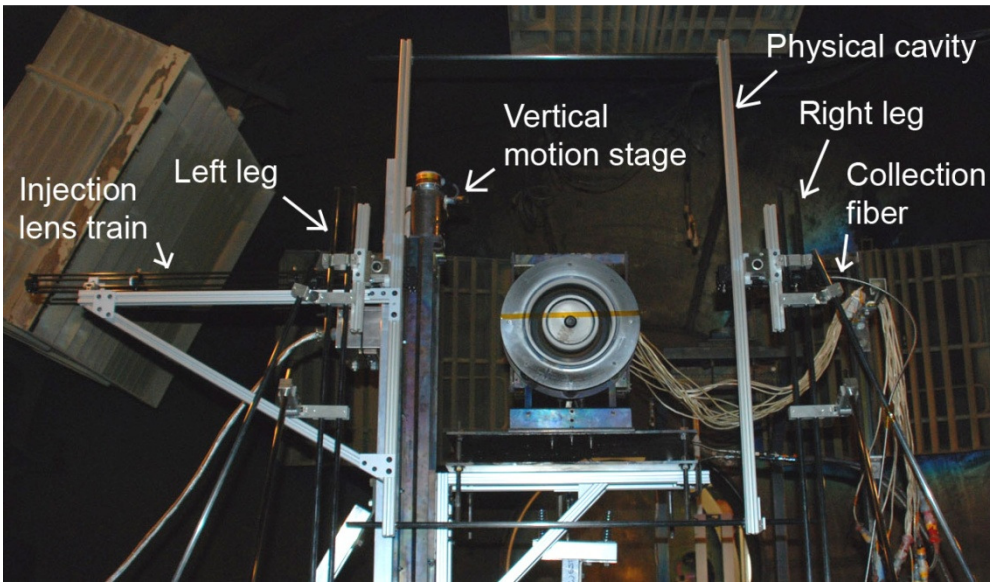

Figure 7-4. Photograph of the prototype CRDS sensor.

SNR of ~2 was found in a few of the 600 V, 10 mg/s traces but by this time the ring-down time constant had dropped to ~0.15 μs. The rapid degradation of the mirrors prompted an addition of two protective covers for the entrances to the two mirrors in the version-two sensor.

The second problem encountered was that the mirrors we used did not meet the requested specification. At an R_M of 0.994, the sensor was not sensitive enough for boron density measurement. Furthermore, the coating and/or substrate of these mirrors did not transmit UV light as well as it should. In the version-two setup, these mirrors were replaced with better ones that actually met the specification of R_M ~0.997 and transmit UV light well.

The third discovery was that although the alignment of the cavity could be maintained during the first live test, it could not be maintained through subsequent chamber pump downs. The change in behavior is attributed to small differences in the chamber structure between tests. The physical cavity could not withstand the geometric changes in the chamber during pump-down since the CFRs it is made from are rather flimsy. This discovery combined with the results of the thermal stress test led to a full re-designing of the physical cavity into a welded aluminum structure, which is much more rigid against warping. It was also confirmed that motorization of the mirrors is highly desirable for maintaining alignment during pump-down, and very useful for maintaining mode-match. Repeated incidence of the physical cavity going out of alignment during chamber pump-down is what prompted the end of the prototype testing.

The fourth problem encountered was that it was sometimes difficult to tell whether the mode-matching into the TEM_{00} mode was achieved. During the prototype testing, the placement of various optics were first calculated using basic mode-matching equations [101] followed by fine tuning the position of the mode-matching lens. The mode-match was found to be not that sensitive to the position of the lens over a 10 mm span but shows signs of change beyond that range. Normally, when mode-matching, one studies the mode pattern and looks for the regular transmission peaks separated by the free spectral range of the cavity. This is because when the light is coupled to only the fundamental mode, transmission is only possible at frequencies separated by the free spectral range. However, vibrations that passed through the vibration isolator and

variations in the laser frequency made it difficult to judge what the spacing between transmission peaks was. Pre-pump-down mode-matching was considered achieved when the coupled power is maximized and the mode pattern captured by instantaneous oscilloscope snapshots looks correct. Since the cavity mirrors were not motorized, it was sometimes difficult to get proper alignment and other non-fundamental modes may have been excited as well. This problem prompted a careful review of the setup for the version-two sensor and led to fine tuning of software filtering, motorization of the mirror, and methods to identify when non-fundamental mode ring-downs were occurring during live test.

7.3 Experimental Setup and Procedure for the CRDS Experiments

In this section, the version-two cavity is described along with detailed procedural information on how the CRDS tests are carried out. The CRDS study was carried out in three tests, also called rounds for the purpose of this dissertation, because mirror degradation prevents all seven operating conditions in the test matrix from being examined in one round of testing. There are no differences in the experimental setup between the three rounds.

7.3.1 Facility, Thruster Settings, and Test Matrix

The facility used for the CRDS testing is the same as the facility used for the LIF testing. Thruster settings and thruster power circuits are all identical as well. More information can be found in Sections 5.2.1 and 5.2.2. However, the cathode was changed between the LIF and CRDS tests. No difference in the performance of the thruster was observed with the switch to the new cathode. More information about the two cathodes can be found in Section 2.4.

The operating conditions tested for the CRDS study is the same as that for the LIF studies and the operating condition table is reproduced in Table 7-2 for convenience.

Table 7-2. Thruster operating conditions for the CRDS experiments.

Discharge voltage, V	Anode mass flow rate, mg/s	Discharge current, A
150	10	9.08
150	20	21.55
150	30	35.35
300	10	8.94
300	20	20.11
300	30	33.25
600	10	9.36

The data locations for the CRDS test will be reported using the non-dimensionalized coordinate P. The P-axis is the axis along which the vertical stage travels. At the 12 o'clock position when viewing the thruster face-on, the values of P are the same as the values of the non-dimensional radial coordinate R. In other words, P = 0 when the CRDS laser beam is directly in front of the nominal channel inner radius at the 12 o'clock position and P = 1 when the beam is directly in front of the nominal channel outer radius at the 12 o'clock position. Figure 7-5 illustrates this concept. The locations at which data were taken for all operating conditions range from P = -0.2 to P = 1.2 at intervals of 0.1 for a total of 15 data locations. Additional traces are taken at P = 2, which is sufficiently far away from the thruster plume along the channel exit plane that there should be no boron signal. Taking these traces help check the "health" of the cavity. The P = 2 traces are taken at least once before taking data at the other data locations and once afterward for a given operating condition. If both traces come out looking like flat lines, then the cavity was assumed to be in good health during the acquisition of the entire data set. One of the main problems the P = 2 traces can detect is poor mode-matching, the details of which can be found in Section 7.3.4.

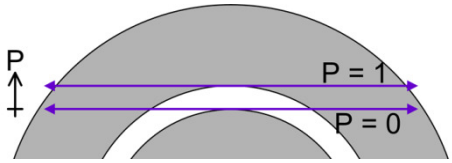

Figure 7-5. Illustration for the P coordinate.

At each data location, 101 wavelengths are scanned to construct the shape of the B I 249.848 nm transition. The total scan range is 70 GHz centered on the transition. The ring-down time obtained at each wavelength is an average of ~150 ring-downs, though the exact number can vary from 100 to 200 depending on how many ring-downs were deemed necessary to obtain good statistics. In total, the data presented in this dissertation

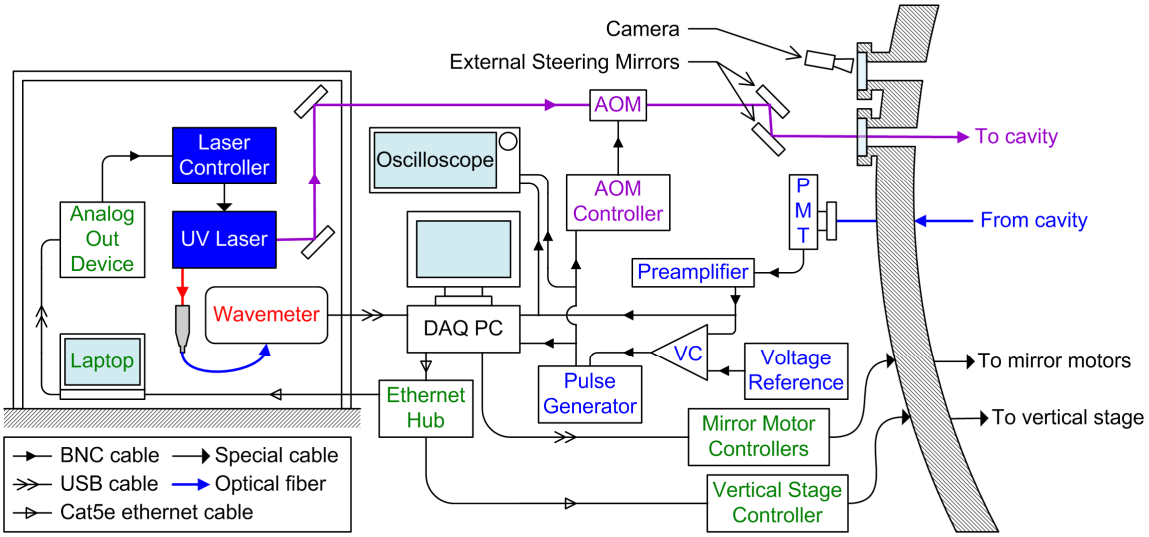

Figure 7-6. Air-side experimental setup for the CRDS study.

are the result of over 1.8 million ring-downs. This count does not include ring-downs recorded during troubleshooting scans, preparatory scans, and scans made during prototype testing.

7.3.2 Air-Side Setup

Figure 7-6 shows the air-side experimental setup for the CRDS study. The DAQ computer, labeled as DAQ PC, uses a custom Labview v9 program along with some manufacturer software to control various components. The DAQ PC communicates through a USB hub (not shown in figure) with four TST001 stepper motor controllers that control the mirror motors on the cavity. Through an Ethernet hub, the DAQ PC communicates with the Aerotech Ensemble MP10 vertical stage controller via the software Motion Composer v3 from Aerotech. Through the same Ethernet hub, the DAQ PC communicates with a laptop computer inside of the laser room, which is the enclosure on the left of Figure 7-6. The communication between the two computers is handled by standard Ethernet protocols programmed into custom Labview programs. Housing the UV laser in the laser room is necessary to protect the laser from dust and noise in the rest of the facility.

The laptop computer inside of the laser room controls a USB-3101 analog output device from Measurement Computing, which controls the scanning voltage of the piezoelectric element in the laser. There is an additional RC low-pass filter (not shown)

between the analog output device and the UV laser controller. This filter is designed to filter out as much high frequency noise as possible while maintaining ~0.1 s rise time. The reason such a complex setup was used to control the UV laser is because the distance between the DAQ PC to the UV laser controller is long enough that the line suffers from electromagnetic interference. The amount of noise is actually very small, but even small jitter in this signal is greatly amplified by the piezo-controller, the piezo-element, and the PID locking of the second and fourth harmonic oscillator inside the laser. When the laser is also performing wavelength sweeps, the total speed with which the wavelength changes is high enough that the fourth harmonic PID controller cannot keep up and the laser power drops to near zero. In other words, if the electronic noise in the scan control signal is too high, the laser power randomly drops to zero. More details about the PID controllers are described later. The solution to this problem is to use the laptop computer as communication relay. Since the laptop computer uses an in-line regulated power source (i.e. wall power goes to battery, then battery power goes to laptop), its USB output exhibits very low noise level. Combined with the RC low-pass filter, the noise in the scan control signal is small enough to be barely noticeable in the final output power of the UV laser.

The UV laser is a Toptica TA-FHG-110 frequency-quadrupled diode laser. The laser is composed of three major components. The first component is a taper-amplified diode laser that typically outputs ~500 mW of 1000-nm light. The second component is a frequency doubling cavity in a bow-tie configuration that takes the 1000-nm light and frequency-doubles it to ~170 mW of 500-nm light. The third component, shown in Figure 7-7, is another frequency doubling cavity in a bow-tie configuration that frequency doubles the 500-nm light to ~10 mW of 250-nm light. The linewidth of the 250-nm laser output is ~4 MHz. The bow-tie cavity is formed from four high reflective mirrors. One of these mirrors in each cavity sits on a piezo-element. Much like the cavity we are using for boron

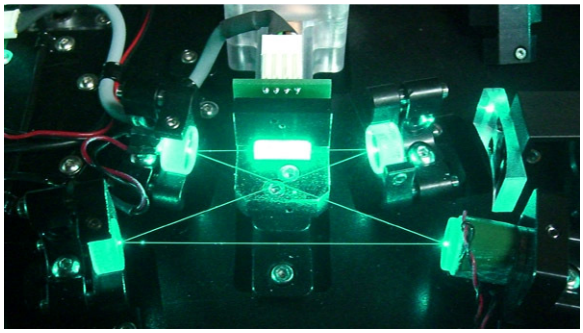

Figure 7-7. Second frequency doubling bow-tie cavity in the UV laser for CRDS.
Alternate title: Motivation #2: it also looks prettier in person.

detection, the bow-tie cavity only goes into resonance when cavity length matches certain quantized criteria. To keep the laser output high, the cavity length needs to be monitored and maintained using the mirrors that sit on piezo-elements. Stabilization of the bow-tie cavity is executed using the Pound-Drever-Hall technique. Control of the piezo-elements is performed by PID controllers. When the PID controller is successfully placing the piezo-actuated mirror at the proper position for laser light to propagate through the bow-tie cavity, the cavity is said to be in "lock". Since the PID controller of the second frequency doubler is trying to lock to the output of the PID controller of the first frequency doubler, the output of the second frequency doubler is extremely sensitive to electronic noise and vibrations. Loud sounds, like the noise from cryopump operation, can cause the second bow-tie cavity to lose lock and is part of the reason why the UV Laser must be housed in the laser room. The sensitivity of the second bow-tie cavity to noise is also the reason why a separate laptop computer is used to control laser scanning.

A small amount of the ~1000 nm laser light generated by the taper-amplified diode laser is split off in the UV laser and outputted in a separate port from the main output. This infrared beam is coupled into a High Finesse WS-7 Fizeau-type wavemeter, which has an accuracy of ±60 MHz at 1000 nm, or ±0.2 pm. The wavelength of the frequency-quadrupled light is calculated from this reading by simply dividing by 4, which gives an accuracy of ±0.05 pm at 250 nm, or ±240 MHz in frequency space. The wavemeter also has a resolution of 10 MHz and updates at up to 150 Hz. It communicates with the DAQ PC via a powered USB hub due to the physical distance to the DAQ PC. During the CRDS test, to excite resonance in the CRDS cavity, the UV laser is continuously executing triangular sweeps at a rate of ~1 pm/s over a range of ~1 pm. This wavelength sweeping normally generates an oscillatory reading in the wavemeter, but the wavemeter software is set to perform floating average of 50 measurements so the actual value recorded into the raw CRDS data is a time-averaged value of the laser wavelength.

The UV laser light is directed by high reflectivity mirrors to a laser platform that sits next to the vacuum chamber. The UV beam passes through a Neos 35110-2-244-BR acousto-optical modulator (AOM), which is powered by a 31110-4DS AOM controller. Roughly 70% of the input beam power is coupled into the first order mode of the AOM output when the AOM is on. This first order beam propagates at a slightly different angle

than the original laser beam that comes into the AOM. An iris is used to block off all of the light coming out of the AOM except the first order beam. This beam is then injected into the vacuum chamber via two external steering mirrors with fine-tuning knobs. The rough alignment of this UV beam is monitored using a video camera and a pair of irises built into the injection lens train of the physical cavity.

The output from the CRDS cavity is returned via optical fiber to a Hamamatsu R9110 photomultiplier (PMT) through a band-pass interference filter with a 10-nm bandwidth centered at 250 nm. This PMT is powered by an Oriel 70705 PMT power supply. The output of the PMT is converted from current to voltage by a FEMTO DHPCA-100 preamplifier operating at a gain of 10^4 in low noise mode, which has a bandwidth of 14 MHz. Preparatory ring-down tests show that this is the highest gain setting that does not affect the value of the measured ring-down time constant (i.e. higher gain settings have smaller bandwidth, which starts to change the measured ring-down time).

The output from the preamplifier is fed to one of the channels on the AlazarTech ATS460 125-MHz DAQ card in the DAQ PC. The DAQ card was operated at only 100 MHz as a precaution against over stressing the card. The preamplifier output is also fed to a LeCroy 314 oscilloscope for live monitoring and a custom voltage comparator circuit built into a Schmitt trigger configuration. The voltage comparator is an LM361 with a switching time of 20 ns. It is powered by an Instek GPS-3303 laboratory power supply. Switching and stabilizing capacitors were added to improve the performance of the Schmitt trigger against random noise. The power lines between the power supply and the circuit were kept short for the same reason. The reference voltage was provided by a potentiometer based voltage divider with a 9-V battery as the voltage source. The non-inverting output of the circuit was used to complete the Schmitt trigger configuration while the inverting output was used to drive the triggering of a pulse generator.

The BK Precision 4030 pulse generator has an insertion delay of 60 ns (i.e. it takes 60 ns from the time when the triggering condition is met to when the pulse generation starts). The pulse width is set to 12 μs, which is more than ten times the typical ring-down time. Output of the pulse generator is fed into the AOM controller to initiate ring-down when the voltage comparator decides the signal in the cavity is high

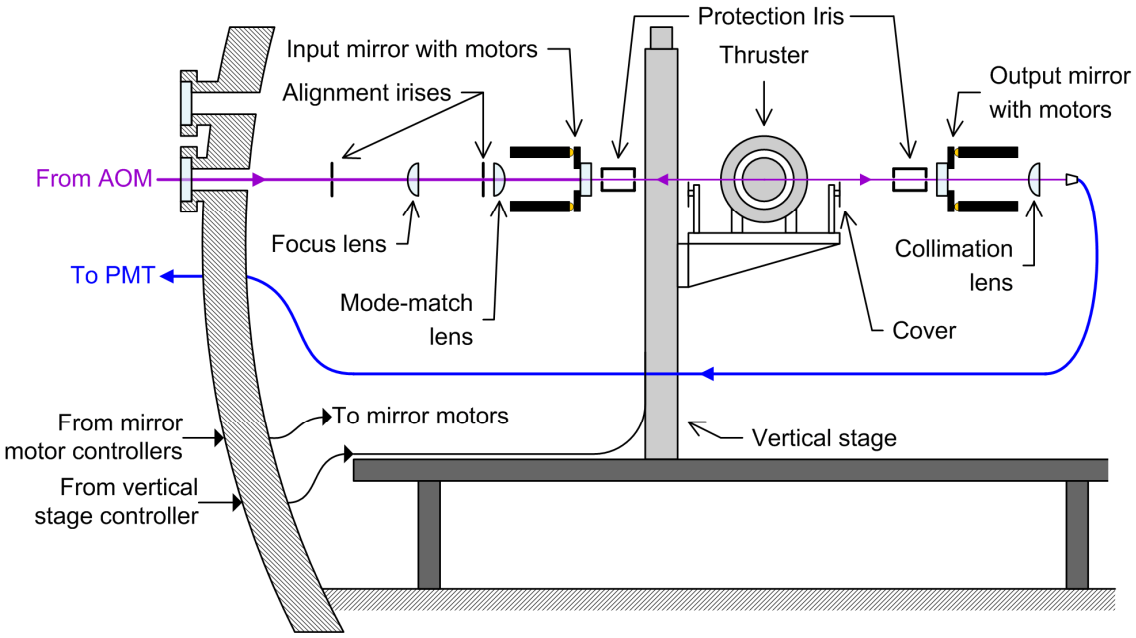

Figure 7-8. Vacuum-side experimental setup for the CRDS study.

enough. This output signal is also fed into the DAQ PC to trigger the start of data acquisition and the oscilloscope for live monitoring.

7.3.3 Vacuum-Side Setup

Figure 7-8 shows the vacuum-side experimental setup for the CRDS study. Figure 7-9 shows a composite photograph of the injection lens train being rough aligned. It is a composite photo because the camera cannot focus on the near and the far part of the lens

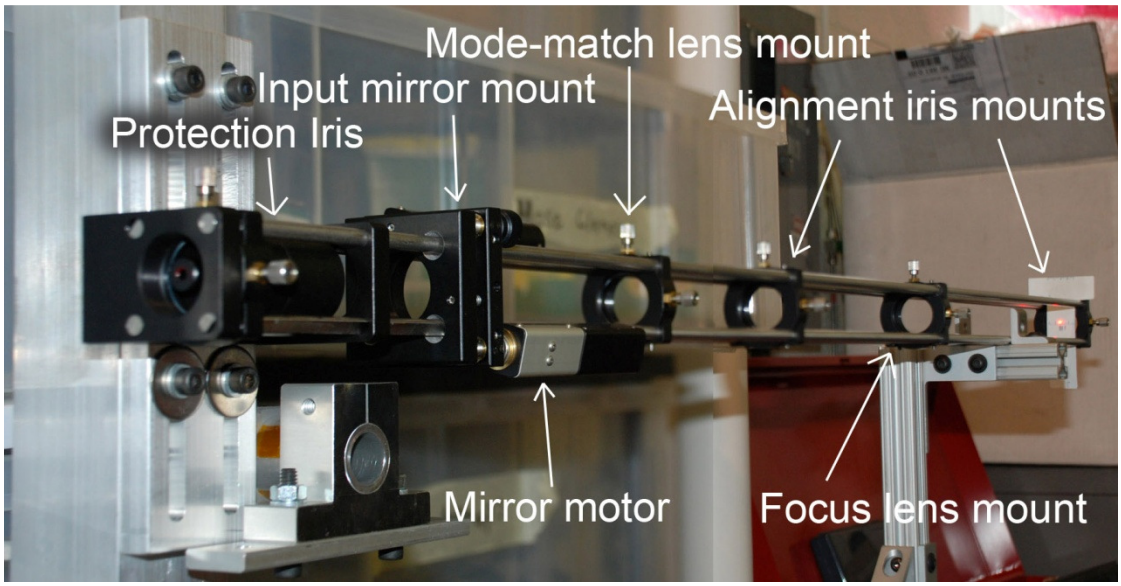

Figure 7-9. Photograph of the injection lens train being rough aligned.

train at the same time. There are no optics shown in Figure 7-9. They are installed once the cavity has been properly positioned inside the vacuum chamber.

The first order beam from the output of the AOM is deflected by the external steering mirrors though a pair of alignment irises. The alignment irises are made of fluorescing discs that turn red when interacting with UV light. A precision drilled hole of 2±0.05 mm diameter is drilled into the center of each disc. The discs are mounted on two-axis translating mounts to help with the positioning of the alignment irises. The UV laser beam is usually slightly bigger than 2 mm in diameter by the time it reaches the first alignment iris. Parts of the beam beyond 2 mm fluoresce the iris and helps alignment. Fluorescence from the irises is monitored by an external camera through a viewport.

At this point, the beam going into the focusing lens is of known diameter and is approximately collimated. The focusing lens has a nominal focal length of 300 mm and is made of uncoated UV-grade fused silica. Its purpose is to bring the UV laser light down to a point of known diameter and location to aid in mode-matching. The original plan also included the addition of a beam cleaning iris at the point where the beam is narrowest to make the beam more Gaussian-like. However, bench-top testing (in the laser room) with the same optics and geometry show that beam cleaning did not noticeably improve mode-matching results. Since such an iris can also make alignment more difficult in vacuum, it was not used in the final design.

The mode-matching lens, which comes next, matches the known position and diameter of the post-focusing-lens UV beam with the position and diameter of the TEM_{00} mode of the CRDS cavity. This lens has a nominal focal length of 125 mm and is made of uncoated UV-grade fused silica. Details on mode-matching are described in Section 7.3.4.

Next in the injection lens train is the input cavity mirror. For the version-two CRDS sensor, this and the output cavity mirror are custom fabricated by MLD Technologies. These mirrors are made from 1” diameter, 1/4” thick super-polished fused silica substrate. One side of each mirror is concaved with a 1-m radius-of-curvature and coated for high reflectivity at ~250 nm. The other side is flat and uncoated. The CRDS cavity beam resides at 4.0±0.5 mm downstream of the thruster channel exit plane and has a beam waist of 0.2-0.25 mm, roughly matching the value from theoretical calculations.

The mirrors sit on kinematic stages that have been modified to have motorized tilt and yaw axes. The input mirror mount is actuated by a pair of Thorlabs ZST6 micro stepper motors, while the output mirror mount is actuated by a pair of ZST13 motors. The difference between the two pairs of motors is the maximum travel, which is 6 mm for the ZST6 and 13 mm for the ZST13. A travel of 6 mm is more than sufficient for cavity mirrors and the space constraint on the injection lens train prevents the use of a longer stepper motor. The space constraint on the collection lens train was not as bad and the ZST13's were already available prior to the designing of the lens train.

There are also protection iris tubes between the Hall thruster plasma and each cavity mirror. These tubes have openings on each end that are 3.6 mm in diameter and are ~80 mm long. The open space between the iris and the mirror mounts as well as along the rest of the lens train is sealed with transparent tape to protect against sputtered particles that are floating around the chamber during testing. In addition, there are two protection covers that extend axially downstream from the platform that holds the Hall thruster. Since we are assuming the thruster is axisymmetric, CRDS results can be obtained by scanning at most half of the thruster in the vertical direction. When the CRDS sensor is not in used, the entire thruster platform can be raised so that the protective covers sit right over the holes of the two protection iris tubes. Although not shown this way in Figure 7-8, the covers sit at only 5 mm in front of their respective protection iris openings when in place. These covers also act as an additional position reference since we will know a priori the P-position at which the covers cut off the cavity beam. All materials that sit near the thruster are made of non-magnetic material and do not perturb the thruster's magnetic field.

Vertical position of the thruster platform is controlled by an Aerotech ATS62150 motion stage with a travel of 1.5 m and is powered by a VS-U33 stepper motor. The resolution of this motion stage is 0.01 mm. The accuracy of this stage is determined by physical calibration from the position of the protective cover and the size of the CRDS cavity beam and is accurate to ±0.1 mm. The position of this vertical stage is monitored by an HS25 1024-line relative encoder from BEI Sensors. The thruster sits on an adjustable platform the lower half of which is supported by the vertical stage. The tilt of

this platform is adjusted so that the channel exit plane of the thruster is parallel to the travel axis of the vertical stage to within ~1°.

The light that comes out of the output cavity mirror is collected by the collimation lens and focused into the collection optical fiber. The collimation lens has a nominal focal length of 50 mm and is made of uncoated UV-grade fused silica. The collection optical fiber has a 550-μm core diameter, is 9 m in length, and has an attenuation of ~25% at 250 nm. This fiber was in use for about 40-50 hours during the version-two sensor testing as well some additional hours during the prototype testing. No noticeable signs of degradation were observed. Output from this optical fiber is fed into the PMT on the air-side for signal collection.

With the optical parts of the CRDS cavity described, we now turn to the structure that holds the optics in place. Figure 7-10 shows a photograph of the actual vacuum-side experimental setup. The physical cavity is made of four aluminum members welded together. Each member has a 2" X 2" hollow square cross section. The walls of the members are 1/4" thick. This design provides significantly greater rigidity than the prototype design and the thermal expansion is still two orders of magnitude lower than

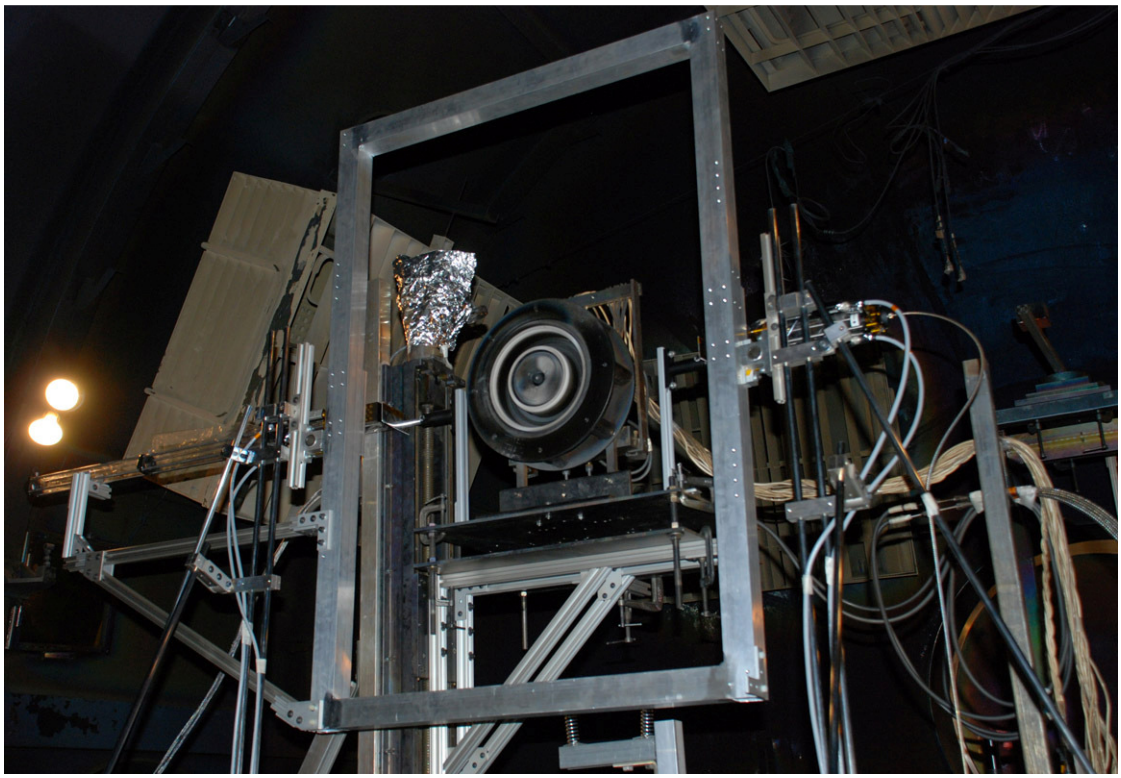

Figure 7-10. Photograph of the vacuum-side experimental setup.

the original design limit. The two legs used to support the physical cavity are the same as the ones used in the prototype design with minor adjustments to the position of the legs to support the shifted center-of-gravity of the new physical cavity. Each leg is, once again, supported by three urethane vibration isolators. In Figure 7-10, the protective covers are sitting right in front of their respective protection iris tube openings. The transparent tape that protects the lens train from sputtered particles is somewhat visible in this photograph. Also shown, to the immediate upper left of the thruster in the photograph, is a sheet of aluminum foil protecting the vertical stage motor and encoder from sputtered particles.

7.3.4 Mode-Matching

Physically, the process of mode-matching can be described as re-focusing laser light from a known position with a known beam waist into the position and beam waist of the TEM_{00} mode of a cavity. Each mode of a laser cavity occupies very specific spatial confines. More details about these mode shapes can be found in [101].

The process of mode-matching begins with theoretical calculations of where to place the mode-match and focus lenses for a given cavity length. These calculations are performed using formulas from [101]. The results are checked using geometric optics calculations. The reason geometric optics calculations are not used from the start is because the formulas from [101] are far more straight-forward and very accurate even though they are supposed to give only approximate solutions.

Next, the obtained positioning scheme is implemented on a laser table inside the laser room and the exact position of the mode-match and focus lenses are fine tuned. The distance between lenses obtained from the fine tuned solution is typically 5-10% smaller than the theoretical value because the true focal lengths of the various optics at 250 nm are shorter than their respective nominal focal lengths. Once the optimal mode-matching positions have been determined, they are replicated on the version-two cavity. For the CRDS study, the distance between the focus lens and the mode-match lens is ~420 mm and the distance from the mode-match lens to the back of the input mirror is ~120 mm. The cavity beam waist is calculated to be ~0.20 mm.

To complete the mode-matching procedure, the mirror motors are actuated iteratively until the mode pattern looks clean and the transmitted signal is maximized.

Figure 7-11 shows an example of the mode pattern of the version-two cavity when mode-matching is completed. Also shown is a dashed line representing the trigger threshold (-0.55 V) of the voltage comparator circuit and the trigger signal from the pulse generator output that controls the AOM. Figure 7-11 represents one of the cleaner mode patterns captured from the version two cavity. Typical patterns look less regular due to the minute amount of vibration that shakes the physical cavity and causes the cavity length drifts speed to fluctuate. It is often necessary to go through several cavity pattern traces to find one that is regular enough to show whether the TEM_{00} mode is properly coupled. Signs of good coupling, shown in Figure 7-11, include large regular spikes that represent transmission at TEM_{00} mode, small (<20% in this case) amplitude non-fundamental mode spikes, and regular triggering of the AOM. It is also possible to tell whether the cavity is close to good mode-match when the amplitude of the transmitted peak is close to when the cavity was still sitting on the laser table in the laser. Of course, the laser beam suffers greater attenuation when the cavity is in the vacuum chamber than when it is in the laser room, but one can still roughly estimate how much signal to expect when mode-match is achieved.

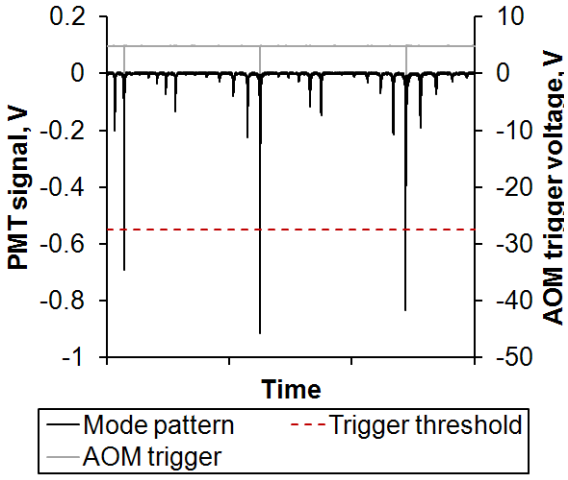

Figure 7-11. Example mode pattern for the version-two cavity.
PMT signal is negative because it is operating in the negatively biased regime.

The high speed of the voltage comparator can also sometimes lead it to misfire, for example when there is a noise spike in the signal. It is possible for the voltage comparator circuit to start ring-down when the laser and the cavity are not in resonance at the fundamental mode. To filter out traces recorded due to these misfire events, a software filter was implemented that rejects any traces with an initial amplitude of less than 0.5 V. Thus, so long as the cavity is in good alignment, the non-fundamental transmission peaks cannot exceed this initial amplitude and will not be recorded.

7.3.5 Curve-Fit Algorithm for the Ring-Downs

Curve-fitting the raw data for the ring-down time constant is accomplished by using an exponential fit sub-routine in Labview v9. This sub-routine uses the bisquare weights method to perform a curve-fit of the data to equation (7-12). The bisquare method is good at rejecting outliers, which may show up as a result of noise spikes.

$$S = S_0 \exp(C_1 t) + C_2 \tag{7-12}$$

where S is the signal intensity, S_0 is the signal amplitude, C_1 is inverted to give the decay time constant, and C_2 is an offset constant. Only negative values of C_1 are allowed in the curve-fit so the raw output for time constant is negative. Lehmann, Berden, and Engeln [104] suggest using ten times the time constant worth of data when there is a possibility of the offset constant varying from trace to trace. This is the driving factor for setting the pulse width to 12 µs (highest ring-down time is 1 µs). All data were taken at 100 MHz so the actual number of data points that comprise each ring-down varies depending on how many points are needed to get approximately 10 times the time constant worth of data. The number of points changes because the ring-down time constant decreases as the mirrors degrade during a test. Also, to ensure that the ring-down has actually begun at the start of data recording, there is a time delay between when the DAQ PC receives the AOM trigger and when it starts recording data. A time delay of 0.55 µs was experimentally found to guarantee that the data recording starts after the ring-down has begun.

Figure 7-12 shows an example of a single ring-down with the associated curve-fit performed by the DAQ PC. This particular trace was obtained prior to the pump-down of the vacuum chamber for the first round of CRDS testing. The data in this figure is representative of data obtained at the start of each round of CRDS testing.

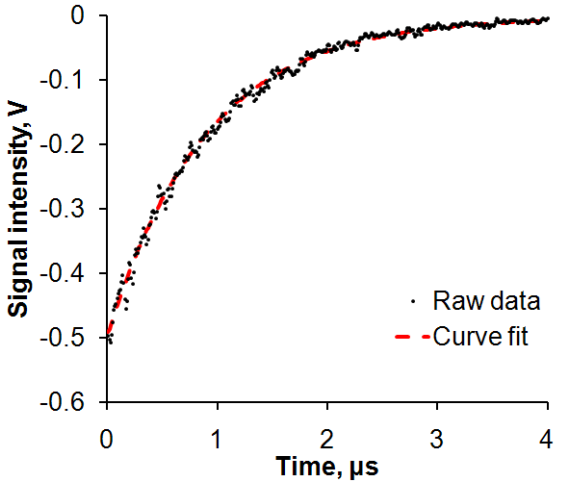

Figure 7-12. An example of curve-fitting to raw data to extract the ring-down time constant.

During an actual test, the custom Labview program that controls the data acquisition only records the time constant extracted from each scan. From the 100-200 time constants recorded this way per wavelength, statistical calculations are carried out on-the-fly and displayed. The health of the cavity can in part be judged by the average ring-down time and the spread in the ring-down time. Benchmark testing of the Labview program showed that it is capable of curve-fitting up to ~700 traces per second and is not a bottleneck for data acquisition. Future CRDS setup using close-looped control of the cavity could potentially require a higher curve-fit rate.

7.3.6 Alignment Maintenance Procedure during Test

During testing, the alignment is tuned before and after each data set. A data set is comprised of all of the data for one operating condition. Sometimes, additional tuning is done if signs of misalignment are detected. The total of time spent to collect data plus alignment and thruster settling time is about one hour per set. The alignment of the cavity will usually last for two to three hours before thermal drifting to the point where the average amplitude of a non-fundamental mode exceeds the trigger threshold, so tuning once an hour is sufficient to keep the cavity in good health.

The first sign of misalignment is usually a drop in transmitted signal amplitude. This is usually accompanied by a drop in trigger rate. Although the transmitted signal amplitude can also drop from mirror degradation, in which case there is an accompanying decrease in ring-down time. Also, the UV laser and the cavity alignment are tuned at the center of the laser's wavelength scan range. When the laser is scanning the edges of this range, the small deflection in the piezo-element that controls the wavelength propagates through the entire system, giving rise to a detectable difference in overall alignment at the CRDS cavity. For this reason, data obtained at the edges of the scan range are more sensitive to misalignment and will show indicative signs before the data obtained at the center of the scan range.

The second sign of misalignment is that the amplitude of a non-fundamental mode exceeds the trigger threshold and the corresponding ring-downs are recorded. At this point the system has to be re-aligned. A cavity in a non-fundamental mode has a slightly different ring-down time than when it is in a fundamental mode. The difference is quite

noticeable when the uncertainty in the ring-down time is ~1% or less, and is the reason why one normally does not want to couple light into non-fundamental modes. When the misalignment of the cavity reaches the point where non-fundamental modes are recorded, three things will happen simultaneously. First, the recorded ring-down time at one or both edges of the scan range will change noticeably. Second, the corresponding uncertainty of the recorded ring-down time will greatly increase, due to the fact that there are now two different ring-down times being averaged together. Three, if one happens to be observing the curve-fitting of the ring-down time live as the scanning proceeds, one will notice two (or possibly more) distinctly different ring-down patterns, one corresponding to the fundamental mode, the other to the non-fundamental mode.

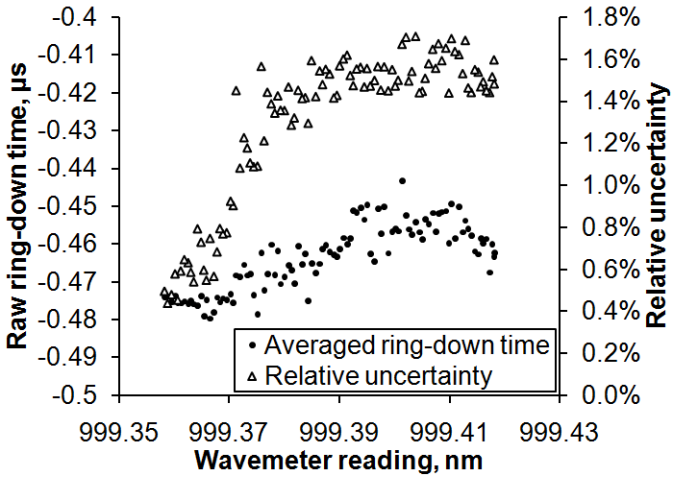

Figure 7-13. An example of raw data recorded when the CRDS cavity is misaligned.

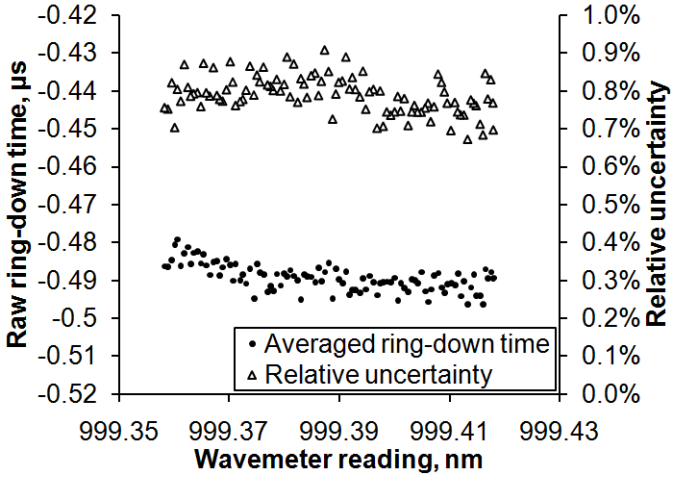

Figure 7-14. An example of raw data recorded when the CRDS cavity is aligned.

Figure 7-13 shows an example of the raw data recorded when the cavity is misaligned. The thruster was operating at 150 V, 20 mg/s at the time and the cavity position was P = 2. The left side of the figure shows very different ring-down time and relative uncertainty than the right side of the figure and the curve-fitting displayed by the Labview program also shows two distinctly different decay patterns. In contrast, Figure 7-14 shows an example of the raw data recorded under the same operating condition and cavity position once the cavity was re-aligned.

Once misalignment is detected, the cavity mirrors are remotely tuned iteratively until the fundamental mode operation is re-established and the signal strength is maximized. If a misalignment occurs during a trace, the trace is discarded, the cavity is retuned, and a new trace is taken. Since the uncertainty is shown to dramatically increase when the cavity is misaligned, the uncertainty of recorded data is also checked in post-processing to make sure no misalignment occurred during the data acquisition.

7.3.7 Post-Test Mirror Cleaning

An interesting discovery that came out of the first round of version two CRDS testing was that the degraded cavity mirrors can be cleaned and reflectivity can recover to pre-degradation levels. Each round of cleaning is done by first applying methanol to lens cleaning tissue and wiping the reflective side of mirror with a mild application of force, followed by the same process with isopropyl alcohol. The mirrors are then reinstalled into the cavity and the ring-down pattern checked.

There were three rounds of testing and the cavity mirrors were cleaned twice. Figures 7-15 and 7-16 show examples of ring-down traces obtained after the first and second cleaning events, respectively. Comparing these figures to Figure 7-12, we can see that the characteristics of the post-clean ring-downs are nearly identical to those obtained prior to the start of the first round of testing.

With that said, it is unclear, how many times the mirrors can be cleaned this way

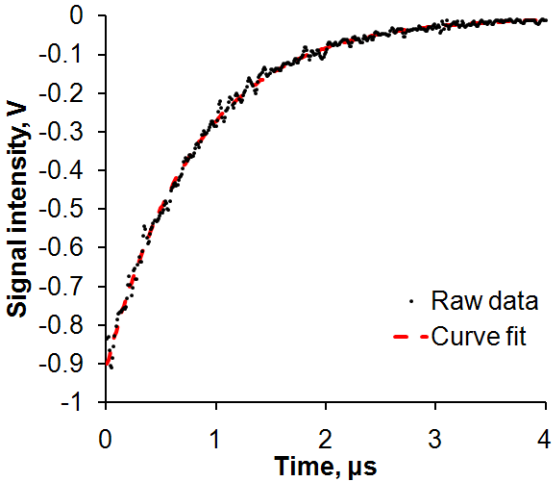

Figure 7-15. An example of ring-down trace obtained after the first mirror cleaning event.

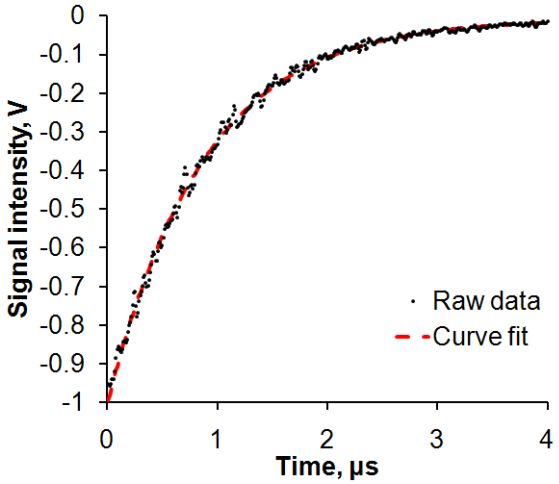

Figure 7-16. An example of ring-down trace obtained after the second mirror cleaning event.

and still maintain functionality. The second cleaning event involved three rounds of cleaning while the first cleaning event involved only one. It is possible that whatever is being cleaned off of the mirrors will become harder and harder to clean off as the mirror surface become more rough with each cleaning.

7.4 Data Reduction for the CRDS Experiments

Having now gone through the experimental setup and procedure, we move on to data analysis and uncertainty analysis.

7.4.1 Calculating the Line-Integrated Boron Density

To calculate the line-integrated boron density, one needs to know the empty cavity ring-down time. Normally, this value remains constant as you perform CRDS scans, but the degradation of mirrors in a Hall thruster environment is fast enough that the value is not a constant. To calculate the empty cavity ring-down time it is assumed that the degradation is due to accumulation of some kind of material on the mirrors' reflective surfaces. This material is treated as an added absorber that grows linearly with time for a given experimental geometry. The exact growth rate may vary with geometry, but is a constant for a given geometry. Mathematically, this constant rate deposition is the same as adding a growth rate times time term into the bracket of equation (7-2). In which case, the empty cavity ring-down can be modeled by equation (7-13),

$$\tau_0 = \frac{1}{at + b} \tag{7-13}$$

where a is a constant that represents the growth rate of the deposition and b is a constant that represents the amount of deposit already on the mirror at the start of the scan. Data points on the edges of the scan range are assumed to be out of resonance and are used to perform the curve-fit for the constants a and b. The data points in-between are assumed to be a part of the 249.848 nm transition lineshape and integrated to find the total line-integrated density. The exact range of integration is the same for all CRDS data analyzed and is 12.4 pm wide centered at 249.848 nm based on the wavemeter measurements. The location of the 249.848 nm marker from the wavemeter measurements corresponds well with the location of peak signal for all obtained data.

Figure 7-17 shows an example of how the curve-fitting to find the empty cavity ring-down time works. This particular trace comes from 300 V, 20 mg/s operating at a cavity position of P = 0.5. As the figure shows, the curve-fit is pretty much shaped like a line due to the small change in ring-down time relative to the amplitude of the ring-down time. Additional data analysis shows that using a linear fit also works fine.

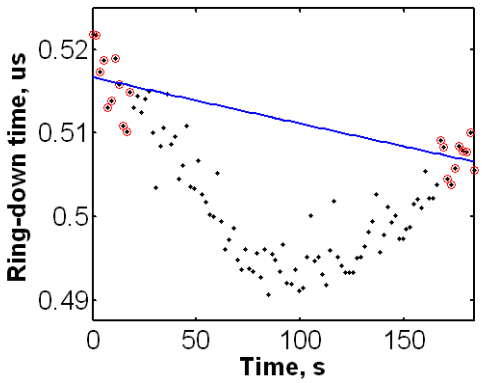

Figure 7-17. An example of how the empty cavity ring-down time is found. *Black dots are data points, red circles near the two edges are data points used to perform curve-fit, and blue line is the resulting fit.*

Once we have the empty cavity ring-down time, we can then use equation (7-3) to calculate the absorbance and equation (7-4) to calculate the line-integrated boron density of the specific ground state we are measuring.

7.4.2 The Onion-Peeling Inversion Method

The onion-peeling method is based on the idea that you can divide the axisymmetric data zone into concentric rings of constant property. The distance that the CRDS cavity beam traverses through each of these concentric rings can be calculated by geometry. A matrix can then be formed that relates the measured line-integrated density to the density within each ring. Figure 7-18 illustrates this concept. This method is intuitively simple and computationally inexpensive when the number of data points is relative low. The matrix formed for this method is an upper triangular matrix and the solution can be obtained by Gaussian elimination. Equation (7-14) shows the matrix representation of the onion-peeling method.

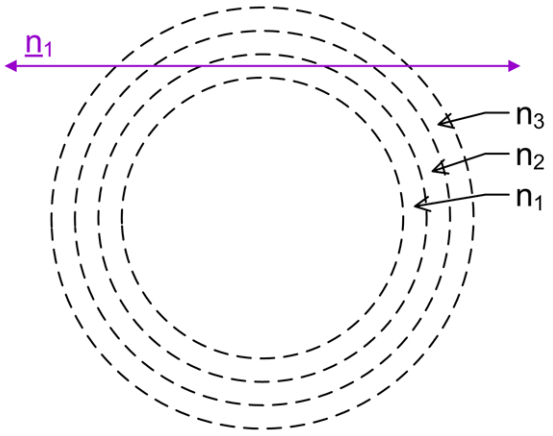

Figure 7-18. Illustration for the onion-peeling concept.

$$\mathbf{K} \cdot \mathbf{n} = \underline{\mathbf{n}} \tag{7-14}$$

where **K** is the geometry matrix and is upper triangular, **n** is the density column vector, and $\underline{\mathbf{n}}$ is the line-integrated density column vector. Equation (7-15) shows formula that governs how **K** is calculated.

$$\mathbf{K}_{ji} = \begin{cases} 0, & i < j \\ 2\sqrt{r_i^{\,2} - p_j^{\,2}}, & i = j \\ 2\sqrt{r_i^{\,2} - p_j^{\,2}} - 2\sqrt{r_{i-1}^{\,2} - p_j^{\,2}}, & i > j \end{cases} \tag{7-15}$$

where the subscript i is both the column index of **K** and the position index for r, the radial coordinate (not in non-dimensional form), and the subscript j is both the row index of **K** and the position index for p, which is the vertical coordinate and the dimensional form of P from Figure 7-5. Note that r and p have a one-to-one correspondence at the 12 o'clock position of the thruster when viewed face-on. Also, the position r_i is associated with the density n_i while the position p_j is associated with the line-integrated density $\underline{n}_j$. More details regarding the mathematical implementation of the onion-peeling method can be found in [99].

Figure 7-19 shows the line-integrated boron density measured for the 6-kW Hall thruster operating at 300 V, 30 mg/s. Figure 7-20 shows the same data after the onion-peeling method is applied. As with the correspondence between r and p, the associated non-dimensional form R and P also have a one-to-one relationship. Each concentric onion ring has a width of 0.1 with the corresponding R value being halfway between the two radii that bound each ring.

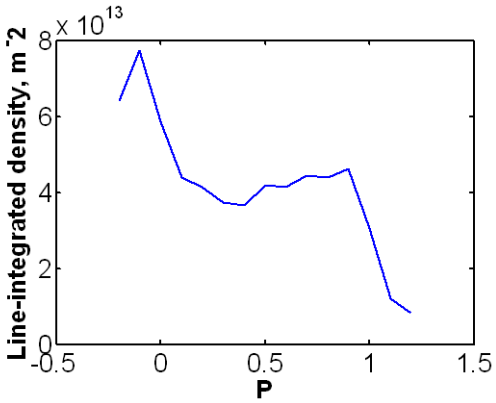

Figure 7-19. Example plot of line-integrated boron density as a function of non-dimensionalized vertical position.

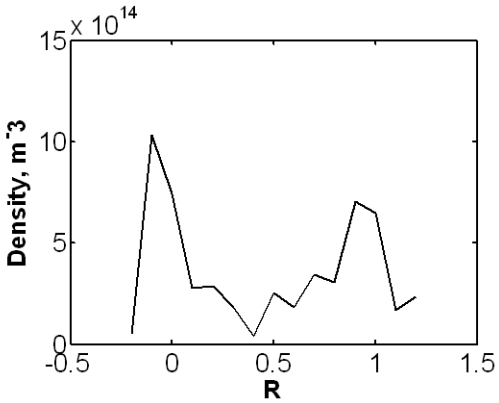

Figure 7-20. Example plot of boron density as a function of non-dimensionalized radial position.

7.4.3 Boron Velocity Modeling

To calculate the boron flux, information about the axial velocity of the boron at each measured location is needed. For the purpose of this dissertation, a very simple model, combined with information on the Sigmund-Thompson distribution measured in [89], will be used to approximate the boron axial velocity.

The model is a simple two point-source model and is formulated in the R-Z plane. Sputtered particles are assumed to radiate isotropically outward in a three-dimensional fashion from two points in the R-Z plane (or two rings in 3D space) at a fixed speed. The speed used is the 23 km/s calculated in Section 6.2.1. Figure 7-21 illustrates the key concepts behind this model. Particles that arrive at the interrogated location from each point source carry an axial velocity that can be calculated from equation (7-16), which is derived from simple geometry,

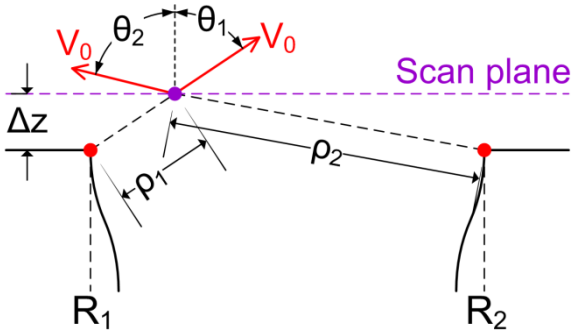

Figure 7-21. Illustration for the two point-source boron velocity model.

$$V_{Z,i} = V_0 \frac{\Delta z}{\sqrt{(r - R_i)^2 + (\Delta z)^2}} \tag{7-16}$$

where V_0 is the speed of the sputtered particles, Δz is the distance from the channel exit plane to the CRDS laser scan plane, and is 4±0.5 mm, r is the radial location of the interrogated point, and R_i is the radial location of the i-th source. Since particles are radiating out in 3D space, the particle density scales as inverse of the square of the distance. The fractional contribution by each source can be computed from the distance between the interrogated point and each source using equation (7-17),

$$f_i = \frac{\rho_i^{-2}}{\rho_1^{-2} + \rho_2^{-2}} \tag{7-17}$$

where ρ_i is the distance from the i-th source to the interrogated location. The average axial velocity of the sputtered boron particles at the interrogated point can then be calculated using equation (7-18).

$$V_Z = f_1 V_{Z,1} + f_2 V_{Z,2} \tag{7-18}$$

The boron flux can now be calculated by combining the density data with the calculated axial velocity. Lastly, the area of each onion ring can be calculated from simple geometry and the total rate of boron sputtering, for the ground state measured, can be derived.

An important assumption in the above simple boron velocity model is that the sources are points in the R-Z plane. In reality, there is contribution to the particles at the interrogated point from neighboring azimuthal locations. The axial velocity of the particles from these locations will be lower than the axial velocity calculated in the model because the arrival angle, θ_i, is always greater when the particles come from out-of-plane. Thus, the calculated velocity is a rough estimate of the axial velocity upper bound.

Additionally, the sputtered boron generally do not radiate isotropically. However, incorporating the anisotropy will not change the answer much because the outer-wall-source contributes very little to the axial velocity near the inner wall, and vice versa. For example, at R = 0, the contribution from the outer-wall-source is 2%, while at R = 0.1, it is 5%. As will be seen in the plotted results later, the majority of the sputtered boron is located very near their respective walls and contribution from the other wall can be approximated as a higher order effect.

7.4.4 Conversion from Boron Flux to Sputter Rate

The conversion from specific-ground-state boron flux to sputter rate can be performed with equation (7-19),

$$SR = \rho_{BN}{}^{-1} m_{BN} \left(\frac{g_{1/2} + g_{3/2}}{g_{3/2}} \right) \Phi_{B,3/2} \tag{7-19}$$

where SR is the sputter rate in m^3/s, ρ_{BN} is the mass density of the BN that the channel walls are made from in kg/m^3, m_{BN} is the mass of a single BN molecule in kg, $g_{1/2}$ is the degeneracy of the $2p[1/2]_2$ ground state of boron, $g_{3/2}$ is the degeneracy of the $2p[3/2]_4$ ground state of boron, and $\Phi_{B,3/2}$ is the flux of boron in the $2p[3/2]_4$ ground state in s^{-1}. The mass density of BN is given by the manufacturer as 2.0 g/cm^3. The above equation is true when the following assumptions apply, one, all sputtered boron are in the ground states. This means they are not in BN molecular complex, excited boron state, or charged boron state. Two, boron ground state population is ejected at thermal equilibrium. It is

unclear at the moment whether this is true but it is the best assumption without a priori knowledge. This assumption means that the ratio of total ground state population to the $2p[3/2]_4$ ground state (lower state of the 249.848 nm transition) is given by ratio of the total degeneracy to the degeneracy of the $2p[3/2]_4$ state. The value of this ratio is 1.5 because the degeneracy is 2 and 4 for the $2p[1/2]_2$ and $2p[3/2]_4$ states, respectively. Three, boron and nitrogen are sputtered in a one-to-one ratio, which is true according to known physics of sputtering and there has been no evidence to the contrary. This means that for every ground-state boron detected, the total mass leaving the thruster is equal to one boron atom plus one nitrogen atom.

7.4.5 Uncertainty and Sources of Error

The uncertainty in the extracted boron density comes from spread in the average ring-down time. One can use equation (7-9) to compute the minimum detectable absorbance for each data trace. Assuming the signal is roughly triangular in nature (i.e. that it has a single peak) one can calculate the signal-to-noise ratio of a data trace. The definition is the same as the one used for LIF data except the noise level is directly computed using equation (7-9). The relative uncertainty in the line-integrated density corresponding to the minimum detectable absorbance is just the inverse of the SNR. Since the uncertainty in the ring-down time is computed with a 95% confidence interval, the relative uncertainty in the line-integrated density derived as previously described also comes with a 95% confidence interval. Further statistical analysis can then be carried out using the nature of the onion-peeling method to compute the uncertainty in each boron density data point. To simplify the analysis for this dissertation, an average relative uncertainty for line-integrated density is computed for each operating condition. The relative uncertainty in the boron density is the same by standard error propagation.

The uncertainty in the relative sputter rate should be dominated only by the uncertainty in the boron density measurement. Uncertainty in the velocity modeling should have a higher order effect on the relative sputter rate as long as it is consistent for all operating conditions. This is because the experimental data from [89] show that the sputtered velocity is not strongly dependent on the energy of the bombarding particle and the LIF data from Chapter 5 show that the variation in the average angle of incidence

near the exit plane for different operating conditions is small (~10° spread). With that said, it is possible that the axial velocity of the sputtered boron is sensitive to even such a small spread in bombardment angle and will require comparison to control experiments and finite-element analysis to check. This is left for future work. For the purpose of this dissertation it is assumed that the uncertainty in the relative sputter rate caused by variation in the axial velocity of the sputtered boron due to differences in operating conditions is negligible compared to the uncertainty of the boron density measurement.

The uncertainty in the absolute sputter rate, on the other hand, is strongly influenced by the velocity modeling. As mentioned earlier, the computed average speed of the Sigmund-Thompson distribution is strongly sensitive to the fitting parameters. For instance, a change in the value of m by 0.1 can cause a 10% change in the computed average speed. Similarly, the exact boundary of integration has a strong influence on the value of the computed average speed. Indeed the best approach to obtaining the axial boron velocity is most likely to measure it directly with boron LIF, but boron LIF in a Hall thruster environment is a completely untested diagnostic. For the purpose of this dissertation, the author will assign an uncertainty of 50% to the sputter rate due to the simplicity of the velocity model used. Refinement of the velocity model through kinetic simulations should greatly reduce this uncertainty, as would boron LIF velocimetry experiment. These are left for future work.

Another source of error contributing to the uncertainty in the absolute sputter rate is the possibility of under-capturing the full lineshape in certain high SNR scans. Some lineshapes do not appear to have completely flattened out by the edge of the scan range. This is often difficult to tell because the lineshapes seem to terminate pretty close to the edges of the scan range. The short scan range is an unfortunate artifact of using direct beam injection to deliver the UV laser light. The slight change in beam angle at the UV laser output leads to variation in the mode-match at the CRDS cavity. If too large of a scan range is used, at the edges of the scan range, the laser light will couple to higher-order modes. Note the laser mode-hop-free range was not a limiting factor in the scan range. So while it is unclear if all of the 249.848 nm lineshape is captured from these select scans, it is assumed that the majority is captured since other scans show proper flattening out at the extreme ends. Under-capturing the lineshape will lead to under-

estimate of the sputter rate, but this under-estimation should be systematic because all of the lineshapes appear to be roughly triangular with similar widths. The similarity in the lineshape is not surprising since the lineshape is driven by similar geometries and the same Sigmund-Thompson-VDF-based Doppler broadening. The maximum systematic error is tentatively listed as 20%.

Other sources of error contributing to the uncertainty in the absolute sputter rate include uncertainty in the mass density of the BN and uncertainty in the correlation ratio between the total sputtered boron density and the boron density of the $2p[3/2]_4$ ground state. The former uncertainty is expected to be a few percent at most, which is negligible compared to the potential error in the axial velocity. The latter uncertainty depends on how much of the boron is in the excited states, the charged states, molecular complex, or the $2p[1/2]_2$ ground state. No information about the amount of sputtered particles that come off as molecular complex is currently published. This dissertation assumes the amount is small based on unpublished microscopy observations by the CSU group (our collaborator on the CRDS project). The ratio of $2p[1/2]_2$ to $2p[3/2]_4$ ground states is assumed to be calculable by assuming thermal equilibrium distribution of states. Confirmation is left for future works. Section 6.3 has already treated excited and ionized boron that may come off of the channel walls. The conclusion was that excited states decay quickly and should be in the ground state when the boron reach the laser beam, and it is assumed that fraction of ionized boron is negligible.

However, it is possible for excited and charged boron to be born from collisions with electrons and ions while the boron is traveling through the plasma. To quantify this effect, we will use cross section data from [105], where Kim and Stone used the binary-encounter-Bethe model to predict the electron impact ionization cross section for boron. From various probe studies inside the channel of the Hall thruster [106], including two for this 6-kW thruster [46, 61], the electron temperature ranges from 20 to 50 eV. We will use 40 eV because it is the typical value for operation at 300-600 V discharge voltage. From [105] at 40 eV, the listed total boron ionization cross section as $4.07*10^{-20}$ m^2. Note that this cross section includes both direct electron-impact ionization and excitation to the $2s2p^2$ levels, which auto-ionizes almost immediately. For electron density, we will assume the plasma in the channel is quasi-neutral and use ion density

measured in probe studies by Reid [46] and Shastry [61]. The highest measured electron density is ~$1*10^{19}$ m^{-3} at the center of the plasma when the thruster was operating at 300 V, 30 mg/s [46]. Near the walls the highest measured electron density is ~$1*10^{18}$ m^{-3} [61]. For near-wall CRDS scans, the fraction of ionized particles is <3%. However, for the distance from the walls to the center of the channel, going through the densest part of the Hall current, the ionization fraction is up to 60%. On the other hand, if the boron does not pass through the bulk of the Hall current, the ionization fraction is only 10%. While Reid's dissertation work showed that the bulk of the Hall current is usually just upstream of the thruster exit plane, there is uncertainty as to the exact location [46]. We can assume that any boron which has a negative axial component (going back into the channel after being ejected) will likely be ionized and accelerated out of the channel. With that said, work by Rubin, et al., [21] have shown that the sputtered boron are preferentially ejected in the downstream direction when the channel walls are bombarded at large angles of incidence as was found in Chapter 5. CRDS results, shown later, also confirm that the majority of the sputtered material is detected very close to the site of sputtering. Thus, while it is not possible at this point to establish exactly how much boron is leaving the Hall thruster channel as ions, the fraction should be small and will tentatively be assigned a maximum systematic error of 20%. That is to say, the absolute sputter rate can be underestimated by up to 20%. If we wish to reduce this uncertainty, kinetic simulations based on the acquired LIF data will need to be carried out in order to properly establish how much of the sputtered boron actually goes into the Hall current, this is left for future work.

Although Kim and Stone [105] did not list all of the excitation cross sections, we can use the excitation cross sections for some of the auto-ionizing levels as representative examples. These cross sections are less than half of the total ionization cross section and the other excited states will generally decay via emission instead of auto-ionization. As mentioned in Section 6.3, the excited states that do not auto-ionize have a typical characteristic decay length of less than 0.5 mm. So, although some excitation will occur based on mean-free-path calculations, these excited boron atoms should quickly decay. Thus, we can also safely neglect the excitation of sputtered boron in the plasma plume.

On a related topic that is not directly addressed by the experiments in this dissertation, the validity of boron CRDS as an erosion measurement diagnostic is still being researched. The latest study on this topic was performed by Tao, et al., in 2009 [23] where they studied the behavior of boron CRDS under a controlled environment. The optical cavity used in that study is nearly identical to the one used in this dissertation. The target is boron nitride being bombarded by a nearly-mono-energetic ion gun. The CRDS signal was found to increase roughly linearly with the gun current and the gun voltage over the test matrix of 30-60 mA and 800-1200 V. This behavior is expected based on prior controlled ion gun experiments involving quartz crystal microbalance (QCM) [21]. So, the general behavior of the CRDS sensor has been validated but direct comparison of absolute sputter rate measurement against another technique, like weight loss, still needs to be made. This task is left for future work.

7.5 CRDS Results

7.5.1 Boron Density Plots

Figure 7-22 shows sample raw CRDS traces for the 300 V, 30 mg/s operating condition. This data set has the lowest uncertainty of all of the data set and nicely

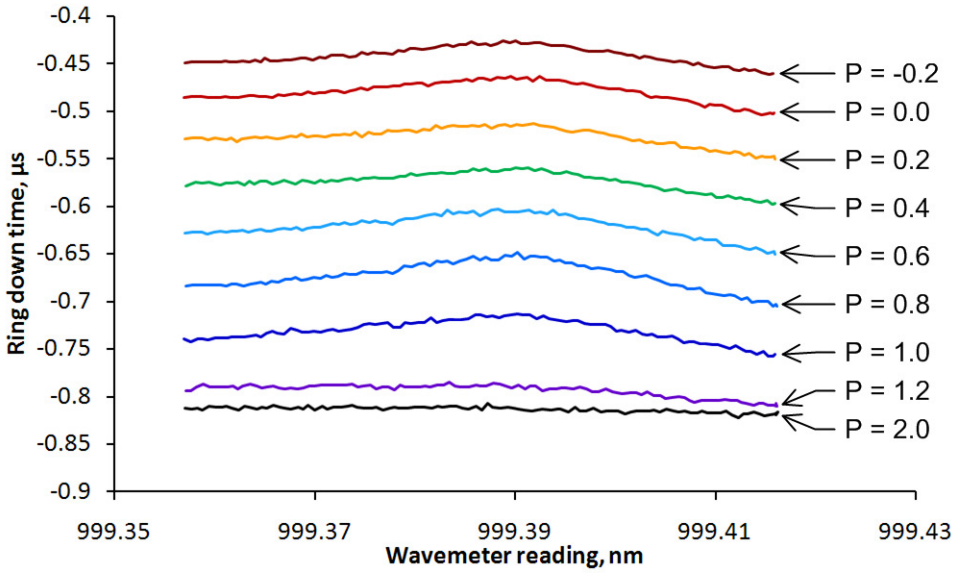

Figure 7-22. Sample raw CRDS traces for the 6-kW Hall thruster operating at 300 V, 30 mg/s.

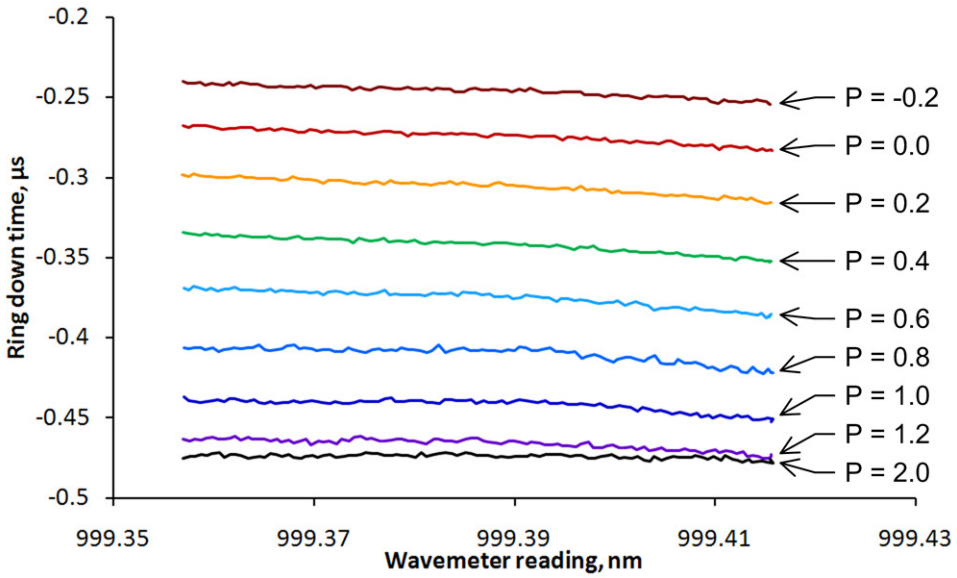

Figure 7-23. Sample raw CRDS traces for the 6-kW Hall thruster operating at 150 V, 30 mg/s.

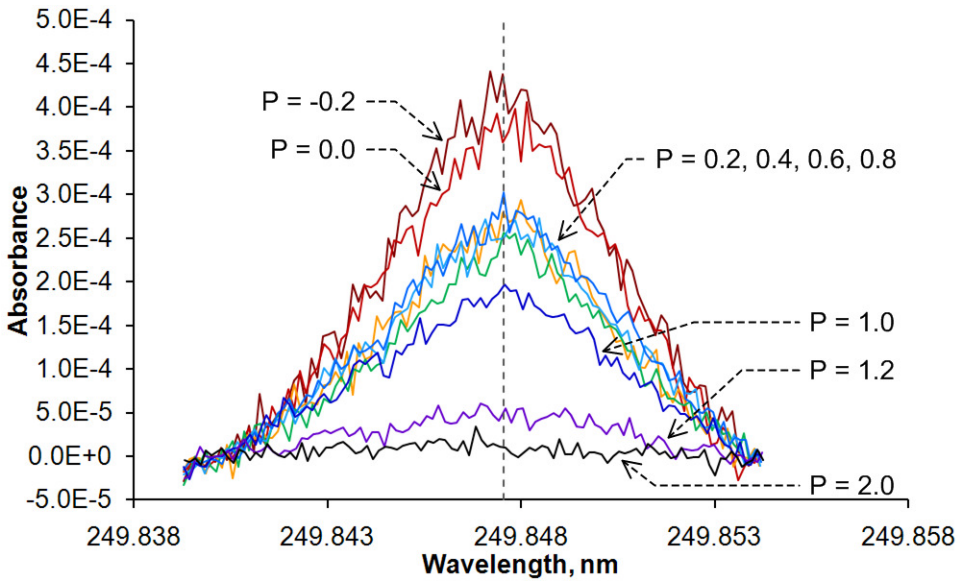

Figure 7-24. Sample processed CRDS traces for the 6-kW Hall thruster operating at 300 V, 30 mg/s.

illustrates the kind of signal to expect when the CRDS diagnostics is working properly. In the figure, P is the non-dimensionalized vertical coordinate. In contrast, Figure 7-23 shows sample raw CRDS traces for the 150 V, 30 mg/s condition, where the measured sputter rate was about the same as the minimum detectable sputter rate. While sputtering information can still be extracted for this condition once the data is put through the data reduction methods previous described, it is clear from the figure that the extracted sputter rate is close to being undetectable.

Figure 7-24 shows the 300 V, 30 mg/s data set after it is processed. The amplitude and shape of the signal as well as the amount of noise in the data is much easier to see in this figure than in Figure 7-22. Figure 7-24 also shows a vertical dashed line indicating the average wavelength of the 249.848 nm transition as reported by the NIST atomic spectra database, ver. 4 (last updated September 2010). The exact value is 249.84755 nm, which is the reason why the line does not line up with the 249.848 nm marker on the wavelength axis. The lineshape shows good symmetry, which is expected for a lineshape dominated by Doppler broadening because the UV beam spends roughly the same amount of time traveling forward and backward. In general, the shape looks Gaussian but could just as easily be described by a triangle for the amount of noise present.

Figures 7-25 to 7-31 show the line-integrated density as a function of vertical position and the corresponding boron density as a function of radial position for the 6-kW

Hall thruster at various operating conditions. In each figure, the left sub-plot is the line-integrated density plot and the right sub-plot is the boron density plot obtained from onion-peeling.

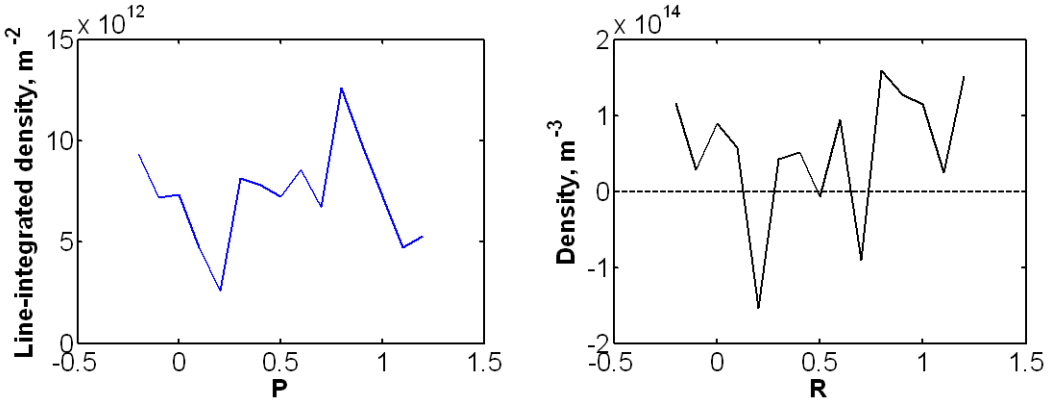

Figure 7-25. The line-integrated density plot and the corresponding boron density plot for the 6-kW Hall thruster operating at 150 V, 10 mg/s.

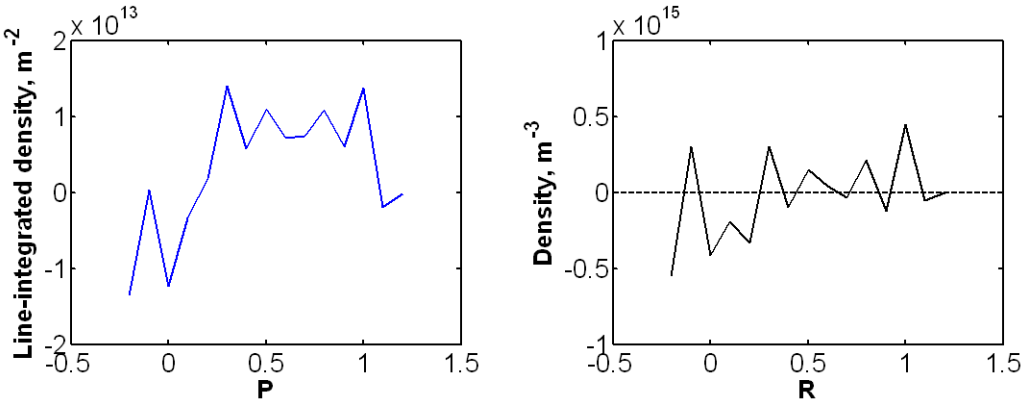

Figure 7-26. The line-integrated density plot and the corresponding boron density plot for the 6-kW Hall thruster operating at 150 V, 20 mg/s.

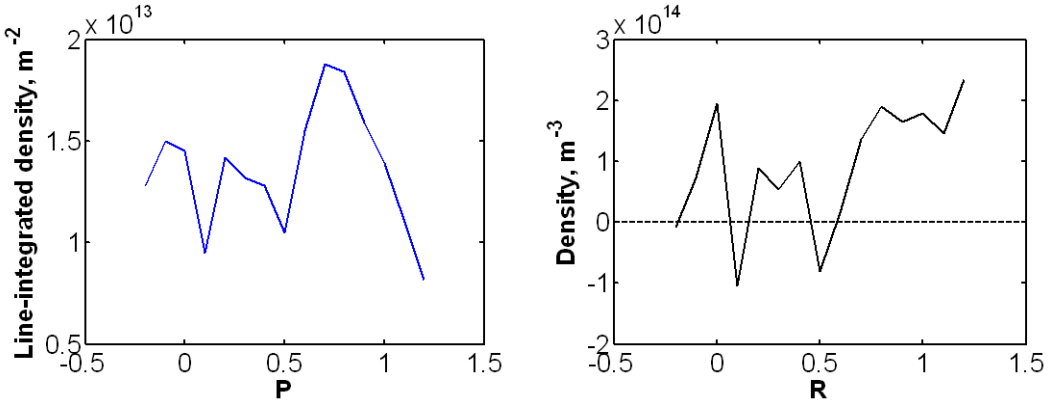

Figure 7-27. The line-integrated density plot and the corresponding boron density plot for the 6-kW Hall thruster operating at 150 V, 30 mg/s.

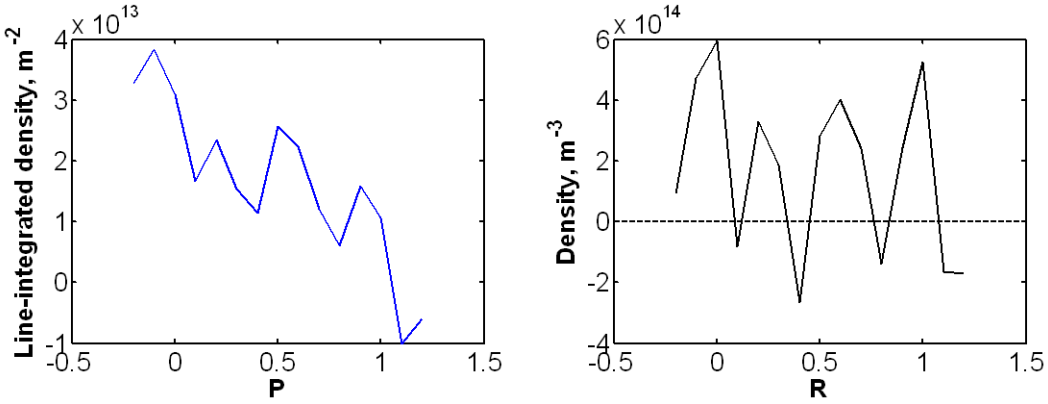

Figure 7-28. The line-integrated density plot and the corresponding boron density plot for the 6-kW Hall thruster operating at 300 V, 10 mg/s.

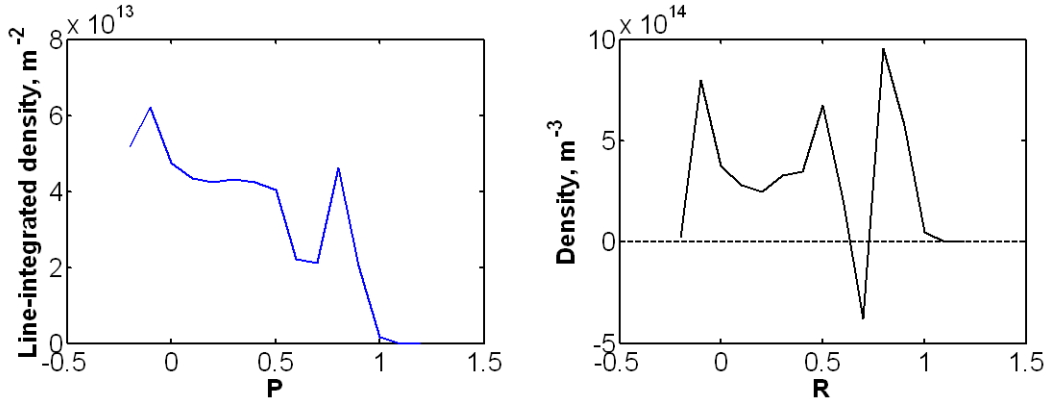

Figure 7-29. The line-integrated density plot and the corresponding boron density plot for the 6-kW Hall thruster operating at 300 V, 20 mg/s.

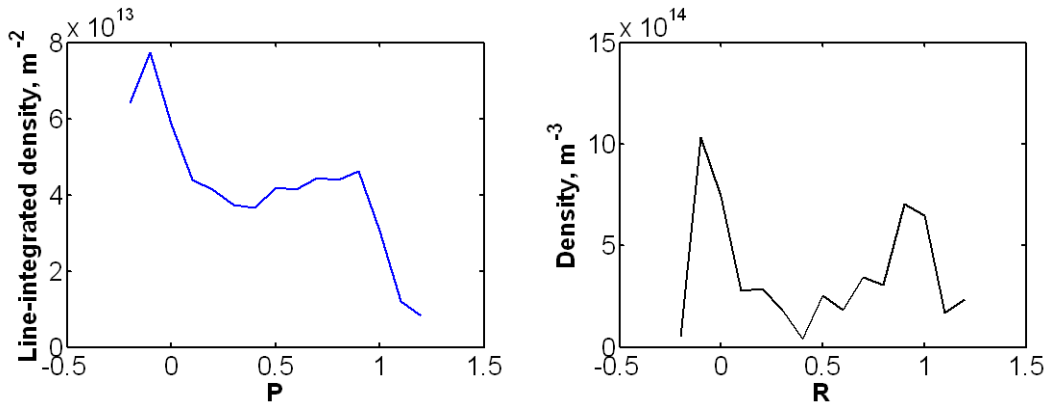

Figure 7-30. The line-integrated density plot and the corresponding boron density plot for the 6-kW Hall thruster operating at 300 V, 30 mg/s.

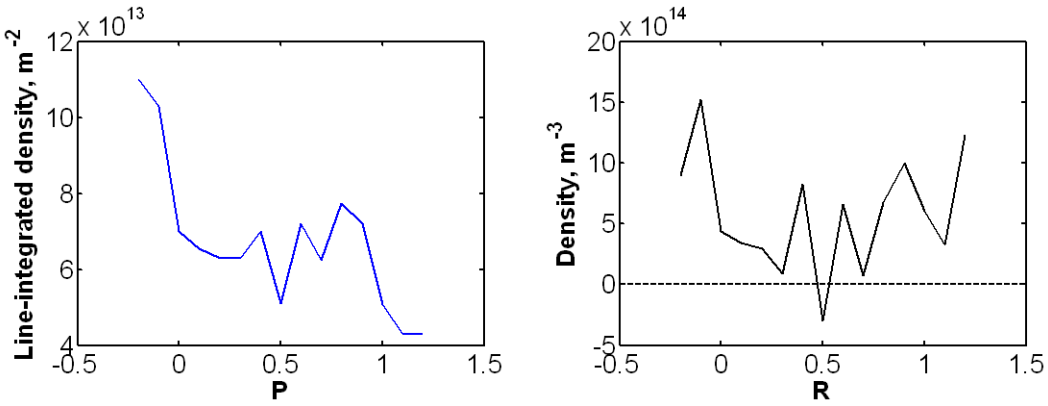

Figure 7-31. The line-integrated density plot and the corresponding boron density plot for the 6-kW Hall thruster operating at 600 V, 10 mg/s.

From the figures of boron density, we can deduce that there was pretty much no detectable boron when the thruster was operating at 150 V, 10 mg/s and 150 V, 20 mg/s. At 150 V, 30 mg/s, the amount of boron is close to the minimum detectable threshold so the uncertainty is most likely high. The amount of uncertainty for all traces will be reported later in the sputter rate calculations. Note that for the 300 V, 20 mg/s condition, there were no usable traces at P = 1.1 and 1.2 due to a misalignment problem that was not caught until post-processing. The boron density is artificially set to 0 at those two data points so the amount of erosion for the 300 V, 20 mg/s condition may be slightly under-estimated (up to 10% based on data at other 300 V condition, which is less than the sensor uncertainty for this set). Also note that the density plot for the 600 V, 10 mg/s condition shown in Figure 7-31 does not quite come down to zero at the edges of the

measured domain. This feature suggests the possibility that boron exists beyond R = -0.2 and R = 1.2 at this operating condition. If this is true, the real sputtered rate is under-estimated in the measurement. However, the error is unlikely to be large because the geometry of the thruster does not really allow boron to show up that far away from the channel. The uncertainty contributed is assumed to be smaller than the sensor uncertainty at this operating condition.

Figure 7-32 shows the measured sputtered boron flux as a function of non-dimensionalized radial position for the 300 and 600 V operating conditions. Flux plot for 150 V operating conditions are not shown since we already know the corresponding signals are close to or below the minimum detectable threshold of the sensor. Figure 7-32 shows a higher boron flux coming out of the inner channel wall than the outer channel wall. However, the respective area of boron emission is not the same due to changing radius. As a comparison, Figure 7-33 shows the sputtered boron particle rate (flux times area) of each onion ring for the same operating conditions. The total amount of boron being sputtered off of the inner channel wall appears to be about the same as the total amount of boron being sputtered off of the outer channel wall. However, these profiles may contain a large amount of error since the simple velocity model is used in the calculations. It is therefore difficult to draw any definitive conclusions regarding the relative amount of boron sputtered from one channel wall versus the other at this time.

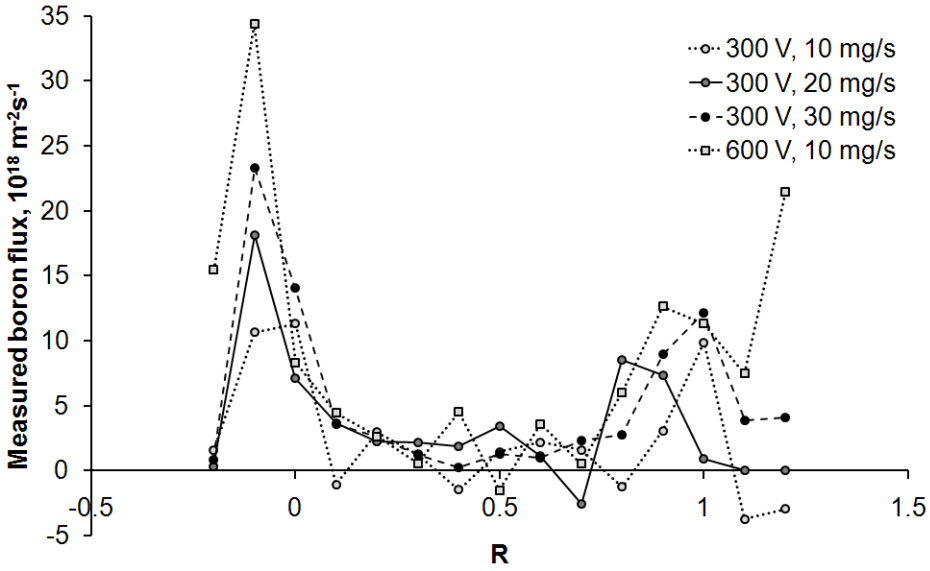

Figure 7-32. Measured boron flux as a function of non-dimensionalized radial position for the 6-kW Hall thruster operating at 300 and 600 V discharge voltage.

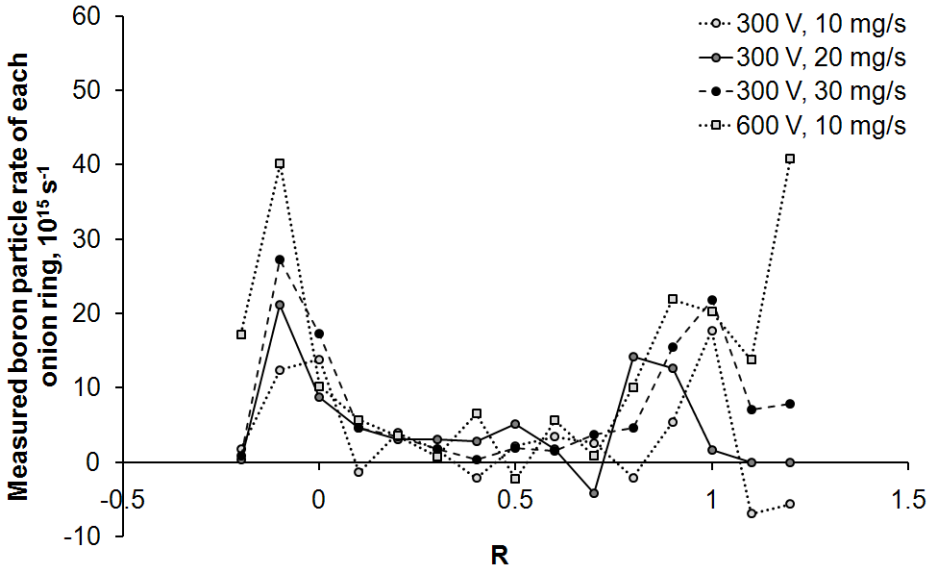

Figure 7-33. Measured boron particle rate of each onion ring as a function of non-dimensionalized radial position for the 6-kW Hall thruster operating at 300 and 600 V discharge voltage.

Figures 7-32 and 7-33 show that the detected boron is largely concentrated near the two walls. Recall the nominal wall radii are at R = 0 and R = 1. Due to enlargement of the channel at the exit plane, the actual radial location of the intersection between the two channel walls and the channel exit plane is at R = -0.09 and R = 1.09. While the boron peak over the inner channel wall appears to be well defined, the boron peak over

the outer channel wall is much more diffused and may be more susceptible to noise. It is unclear whether this difference is real or an artifact of noise.

7.5.2 Sputter Rate Results

Table 7-3 lists the relative and absolute channel wall material sputter rate calculated from the CRDS data for the 6-kW Hall thruster at various operating conditions. This table shows how close the signals for the various 150 V conditions are to the minimum detectable threshold. The uncertainties of 62% and 100% for the 150 V, 10 mg/s and 150 V, 20 mg/s conditions, respectively, show that any signal present is drowned out by the noise. The minimum detectable sputter rate, shown in parentheses, is listed in place of the sputter rate for these two operating conditions. The true sputter rate is around or less than the value in parentheses.

The relative sputter rate is normalized against the sputter rate for the 300 V, 10 mg/s operating condition. The sputter rate appears to rise with both discharge voltage and anode mass flow rate. In particular, doubling the discharge voltage from 300 V to 600 V while maintaining an anode mass flow rate of 10 mg/s roughly quadrupled the sputter rate. On the other hand, increasing the mass flow does not seem to have as dramatic of an effect. The 300 V, 20 mg/s and the 600 V, 10 mg/s operating conditions have almost the same discharge power but have very different relative sputter rates. The sputter rate appears to vary non-linearly with the discharge voltage and almost linearly with the

Table 7-3. Relative and absolute sputter rate calculated from the CRDS data for the 6-kW Hall thruster.

The high uncertainty in the absolute sputter rate is primarily due to uncertainties in the modeling of boron velocity.

Discharge voltage, V	Anode mass flow rate, mg/s	Discharge power, kW	Relative sputter rate	Associated relative uncertainty	Sputter rate, mm^3/s	Associated relative uncertainty
150	10	1.36	0.4	62%	(3.7e-4)	-
150	20	3.23	-0.1	100%	(5.7e-4)	-
150	30	5.30	0.7	39%	1.1e-3	63%
300	10	2.68	1.0	39%	1.5e-3	63%
300	20	6.03	1.6	17%	2.3e-3	53%
300	30	9.98	2.5	7%	3.7e-3	51%
600	10	5.62	4.2	21%	6.0e-3	54%

anode mass flow rate. The exact nature of the relations between the discharge voltage and anode mass flow rate with the sputter rate will be explored in Chapter 8.

7.6 Additional Applications of CRDS

This section is included in this dissertation to describe other ways to utilize CRDS density measurement besides parametric scientific study. It is mainly written to share ideas and is not strictly relevant to the dissertation.

The ultimate goal of studying the channel wall erosion problem is to develop predictive tools that can aid in the design of long life Hall thrusters. The purpose of performing CRDS measurement is to obtain data that can help such development. However, it may be possible to use CRDS results itself as a predictive tool. Two applications of CRDS data for flight qualification come to mind.

The first application is to use the CRDS data to project the life of the channel walls from one operating condition to another. By performing boron density measurements in near-real-time, a thruster that has been qualified can potentially be re-qualified for another operating condition depending on how much difference there is between the conditions. At the very least, one can establish upper engineering limits for how much erosion to expect at the new operating condition.

The second application is to use the CRDS data with LIF measurements to perform accelerated life test. For example, we obtain data on the bombarding particles and the erosion rate for thruster X, then use that data to make a short time step (~500 hours) prediction of what the channel walls will become. We now manufacture a new set of channel walls with the predicted profile and repeat the measurements. Depending how rapidly the walls are eroding, we can change the size of the time step. 500 hours of thruster operation would normally take 3 weeks. Since the CRDS and LIF data can be taken in 1-2 days once set up, the turnaround time for accelerated life test would only depend on how quickly new channels can be manufactured. As such, this approach may only work well if there is a channel wall shaping capability on-site. However, even if it takes 1 week to shape a new ring, that will still be 3x time saving. If the channel wall erosion rate drops over time, the time step can be made longer and the time saving will increase. Of course, this approach is not a replacement for full-flight qualification, which

tests other components besides the channel walls, but it is a more accurate way to re-qualify the channel walls or just to study how the erosion of the channel walls progress.

Chapter 8
Sputter Yield Estimates and General Trends

8.1 Sputter Yield Models for Hall Thruster

Having now described the optical diagnostics used to study Hall thruster erosion and the associated results, we now turn to different ways to interpret and compare the results. We begin with a general description of sputter yield models and end by discussing the models that will be used to predict sputter yield from the LIF data.

In general sputter yield models for xenon bombardment of boron nitride are comprised of an angular component and an energy component, though some do not have an angular component. The energy component generally has an energy threshold. The existence of an energy threshold is due to the fact that sufficient energy must be transferred into the sputtered particle to overcome the binding energy of the matrix. Since xenon is much more massive than boron, the maximum energy transferred in a head-on collision is only a small fraction of the xenon's energy. The equation for energy transfer in head-on collisions was shown in equation (6-3) and is replicated in equation (8-1) [88] for convenience,

$$T_{max} = \frac{4M_{Xe}M_B}{(M_{Xe}+M_B)^2} E_{ion} \tag{8-1}$$

where M_{Xe} and M_B are the atomic masses of xenon and boron, respectively, and E_{ion} is the energy of the incoming xenon particle. The energy transferred is labeled with the subscript "max" because head-on collision provides the maximum energy transfer. Substituting in the values for the masses, we get a maximum energy transfer of 28%. The binding energy of boron in a BN matrix is ~5 eV [89], so the xenon particle needs at least ~20 eV to knock out a boron particle. In reality, the xenon particle requires even higher energy because the majority of collisions are not head-on and a large amount of energy is

dissipated into the BN matrix instead of transferring to a sputtered particle. Nitrogen, being similar in mass and bonded in a similar fashion to the BN matrix, will exhibit similar sputtering behavior. Thus, the energy threshold for the sputtering of BN by bombarding xenon is on the order of a few tens of eV.

One of the current models used to describe BN sputtering by xenon bombardment was developed by Yamamura and Tawara [28]. To obtain the constants in the model, Gamero fitted the model to Garnier's data [35, 36]. The result is shown in equation (8-2),

$$Y_V(E) = \left(0.0099 + 6.04 \cdot 10^{-6}\alpha^2 - 4.75 \cdot 10^{-8}\alpha^3\right)\sqrt{E}\left(1 - \sqrt{\frac{58.6}{E}}\right)^{2.5} \tag{8-2}$$

where Y_V is the sputter yield in mm^3/C, α is the angle-of-incidence (AOI) and is in degrees, and E is the bombarding energy in eV. This model is used in HPHall-2 to predict erosion.

Another model for BN sputtering by xenon bombardment, used by Rubin, et al., [21] involves a combination of the Bohdansky fit and a semi-empirical equation for modeling angular dependence from Yamamura. The equation is reproduced in equations (8-3) to (8-5),

$$Y = k\left[1 - \left(\frac{E_{th}^E}{E}\right)^{2/3}\right]\left[1 - \frac{E_{th}^E}{E}\right]^2 (\cos\alpha)^{-f} \exp\left[-Z\left(\frac{1}{\cos\alpha} - 1\right)\right] \tag{8-3}$$

$$f = 1.71\left(1 + 2.5\frac{\sqrt{E_{th}^{\alpha}/E}}{1 - \sqrt{E_{th}^{\alpha}/E}}\right) \tag{8-4}$$

$$Z = f\cos\left(90 - 286\left[\frac{0.09145}{\sqrt{E}}\right]^{0.45}\right) \tag{8-5}$$

where k, E^E_{th}, and E^{α}_{th} are constants obtained from QCM experimental data. The two models shown above will be used to try to correlate LIF data to the channel wall sputter rate.

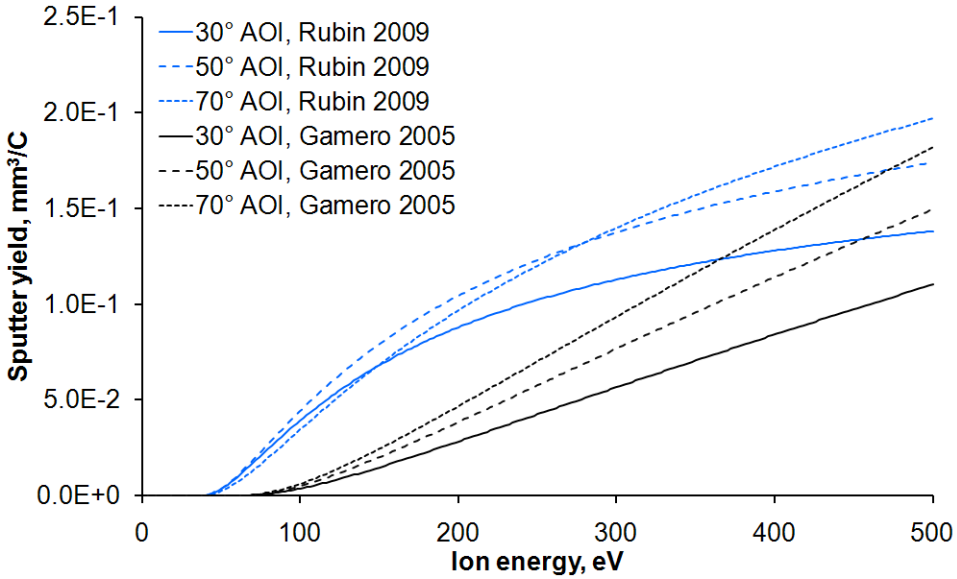

Figure 8-1. Sputter yield versus ion energy for two sputtering models at different angles of incidence.

Rubin 2009 refers to the model described in [21] *and Gamero 2005 refers to the model described in* [35].

Figure 8-1 shows the sputter yield as a function of energy at different AOIs for the two models described.

8.2 Combining LIF and Wall Probe Data with Sputtering Model

To obtain the sputter yield, one needs to know the current density of the ions hitting the channel walls and the local sheath energy in addition to the LIF data shown in this dissertation. The current density and sheath energy will be drawn from Rohit Shastry's dissertation work where he flush-mounted Langmuir probes into modified channel walls of the 6-kW Hall thruster [61]. Since Shastry determined that the Hobbs and Wesson model should be used to calculate the sheath energy, and this calculation scales with the measured electron temperature, the accuracy in the measurement is quoted to be ~15%. Given the spread of the sheath energy shown in the appendix of [61] generally does not vary by more than ±15%, it is decided to simply take one value of the sheath energy per channel wall per condition by averaging all data for each wall and condition. For the ion current, Shastry's probe data does not completely cover our domain of interest, so a parabolic fit versus axial position is performed for each wall and condition and the fit is used to calculate ion current information.

To calculate the average bombardment energy and angle, the LIF data along the R = 0.033 and R = 0.967 radial locations are used as the pre-sheath velocities for the xenon ions arriving at the inner and outer walls, respectively. The average velocity data from LIF is broken into components parallel and normal to the local channel surface at the same axial location. An amount of velocity equivalent to the sheath energy is added to the normal component in the direction of the wall. The post-sheath energy and angle-of-incidence (AOI) is then calculated from the parallel and the post-sheath normal velocity components. Next, a sputter yield model is applied based on the average post-sheath energy and AOI to obtain a local sputter yield in volume per incident current. The local area represented by each data point is then calculated based on the wall location and surface angle, from the first profilometry measurement. Each data point is set to represent an axial extent of ±0.025 in non-dimensional coordinates. For example, the data points at Z = -0.1 would be associated with the channel wall area from Z = -0.075 to Z = -0.125. The data points at Z = 0 would be associated with the channel wall area from only Z = -0.025 to Z = 0 since no wall exist beyond Z = 0. The local current density from curve-fits performed on Shastry's data [61] is then multiplied by the calculated local channel wall area and the sputter yield to obtain the local sputter rate. The total sputter rate for a given condition is the sum of the local sputter rate. For the 150 V, 30 mg/s condition, insufficient LIF data were obtained to make a sputter rate calculation. For the 300 V, 30 mg/s condition, there is no probe data available. For the 600 V, 10 mg/s condition, there is also no probe data, but probe data at 500 V, 10 mg/s condition are deemed sufficient for a rough calculation. For operating conditions where insufficient data are available on the inner wall but not the outer wall, the total sputter rate is calculated as two times the outer wall sputter rate.

Table 8-1 shows a summary of the pre- and post-sheath ion energy and angles as well as the averaged sheath energy used. Data is only shown for Z = 0 to Z = -0.15 because there are no high energy ions upstream of Z = -0.15, except in the case of 600 V, 10 mg/s where more data are shown.

Table 8-1. Pre- and post-sheath ion energy and angle-of-incidence from combining LIF and wall probe data.

Wall probe data is from [61]. *The sheath energy for the 600 V, 10 mg/s condition is from probe data measured at the 500 V, 10 mg/s condition.*

			Pre-sheath			Post-sheath	
Operating condition, (V, mg/s)	Z	R	Ion energy, eV	AOI, deg.	Sheath energy, eV	Ion energy, eV	AOI, deg.
(150, 10)	0.00	0.033	20	72	40.6	60	33
	-0.05	0.033	6	78	40.6	47	21
	-0.10	0.033	2	67	40.6	42	11
	-0.15	0.033	1	54	40.6	42	7
	0.00	0.967	18	79	43.4	62	32
	-0.05	0.967	7	77	43.4	50	21
	-0.10	0.967	2	76	43.4	46	12
	-0.15	0.967	1	61	43.4	44	8
(150, 20)	0.00	0.033	19	77	41.6	60	33
	-0.05	0.033	5	80	41.6	47	19
	-0.10	0.033	1	70	41.6	43	10
	-0.15	0.033	1	54	41.6	42	6
	0.00	0.967	21	80	42.2	63	34
	-0.05	0.967	7	79	42.2	49	21
	-0.10	0.967	2	79	42.2	44	12
	-0.15	0.967	1	65	42.2	43	8
(150, 30)	0.00	0.967	13	84	36.8	50	31
	-0.05	0.967	4	79	36.8	41	19
	-0.10	0.967	1	75	36.8	38	11
(300, 10)	0.00	0.967	153	86	27.6	181	67
	-0.05	0.967	96	86	27.6	124	62
	-0.10	0.967	20	83	27.6	48	40
	-0.15	0.967	3	63	27.6	31	17
(300, 20)	0.00	0.033	93	81	32.6	125	58
	-0.05	0.033	4	89	32.6	37	20
	-0.10	0.033	2	73	32.6	35	13
	-0.15	0.033	1	59	32.6	34	9
	0.00	0.967	139	88	32.2	171	64
	-0.05	0.967	41	85	32.2	74	48
	-0.10	0.967	5	79	32.2	37	21
	-0.15	0.967	1	61	32.2	34	10
(600, 10)	0.00	0.967	449	83	38.8	488	72
	-0.05	0.967	403	85	38.8	442	72
	-0.10	0.967	342	86	38.8	381	71
	-0.15	0.967	217	72	38.8	256	61
	-0.20	0.967	102	68	38.8	141	52
	-0.25	0.967	19	64	38.8	58	31
	-0.30	0.967	3	60	38.8	42	13

Table 8-2 shows a summary of the results from combining LIF and wall probe data with a sputtering model. Generally speaking, the results show the same trend as the

Table 8-2. Sputter rate results from combining LIF and wall probe data with sputtering model.

Wall probe data is from [61]. *IW = inner wall and OW = outer wall.*
**Calculated as two times the outer wall sputter rate.*

Discharge voltage, V	Anode mass flow rate, mg/s	Gamero 2005 [35]			Rubin 2009 [21]		
		IW sputter rate, mm^3/s	OW sputter rate, mm^3/s	Total sputter rate, mm^3/s	IW sputter rate, mm^3/s	OW sputter rate, mm^3/s	Total sputter rate, mm^3/s
150	10	2.7e-7	1.8e-6	2.0e-6	2.5e-3	4.8e-3	7.4e-3
150	20	3.2e-7	8.9e-6	9.2e-6	5.3e-3	8.2e-3	1.3e-2
150	30	-	-	-	-	-	-
300	10	-	7.3e-3	1.5e-2*	-	2.4e-2	4.8e-2*
300	20	1.8e-3	4.1e-3	5.9e-3	8.8e-3	2.3e-2	3.1e-2
300	30	-	-	-	-	-	-
600	10	-	7.5e-2	1.5e-1*	-	1.5e-1	3.0e-1*

CRDS measurements in that the erosion rate increases greatly with the discharge voltage and does not change as much with the anode mass flow rate. The uncertainties in these results are very difficult to calculate as there are many sources that can contribute large amounts of error. Possible sources of error include the fact that the LIF data used is not right next to the channel walls, particularly at the exit plane where LIF data tends to have low SNR. The probe measurements and the use of curve-fitting also contribute their own share of uncertainties which can be read about in Shastry's dissertation [61]. Furthermore, it is still unclear how accurate the existing sputtering models are. In particular, very small variation in the bombardment energy near the threshold energy of the model leads to large changes in the predicted sputtering yield making the 150 V results highly inaccurate. Lastly, the results in Table 8-2 are calculated from average bombardment energy and AOI. In reality, the bombarding particles have a spread in energy and angle that can lead to very different results when taken into account. This, along with improvement of LIF diagnostics, are left for future works.

Table 8-3. Comparison of sputter rate results from different diagnostics on a per-condition basis.

Numbers in parenthesis denote the real value is around or below the minimum detectable sputter rate, which is shown inside the parenthesis. Numbers in brackets denote projected value based on the scaling that the sputter rate is a linear function of the anode mass flow rate at a given discharge voltage.

Discharge voltage, V	Anode mass flow rate, mg/s	CRDS, sputter rate, mm^3/s	Gamero 2005 [35], sputter rate, mm^3/s	Rubin 2009 [21], sputter rate, mm^3/s	Operation percentage at 543 hours
150	10	(3.7e-4)	2.0e-6	7.4e-3	7.2%
150	20	(5.7e-4)	9.2e-6	1.3e-2	5.8%
150	30	1.1e-3	[1.1e-5]	[2.0e-2]	5.9%
300	10	1.5e-3	1.5e-2	4.8e-2	29.4%
300	20	2.3e-3	5.9e-3	3.1e-2	28.7%
300	30	3.7e-3	[2.1e-2]	[7.9e-2]	9.5%
600	10	6.0e-3	1.5e-1	3.0e-1	7.8%
misc					5.7%

8.3 Comparing LIF, CRDS, and Profilometry Data

The next step to reviewing the data obtained in this dissertation is to compare the amount of sputtered materials predicted by each diagnostic. We begin by first comparing the CRDS results to the LIF and wall probe with sputtering model results obtained in Section 8.2 on a per-condition basis. This comparison is shown in Table 8-3. To obtain modeled values for the 30 mg/s conditions, the sputter rate will be assumed to scale linearly with the anode mass flow rate for a given discharge voltage. This is approximately true based on the 300 V CRDS data. In this case, the sputter rate at 30 mg/s can be projected by simply adding up the sputter rate at 10 and 20 mg/s for a given discharge voltage. The result is shown in brackets. The miscellaneous operating conditions include a large range of voltages and anode mass flow rate and will be assumed to have an effective sputter rate that is equal to the sputter rate of the 300 V, 20 mg/s operating condition.

We can now multiply the measured and modeled sputter rate by the amount of time spent at each operating condition and sum them up to obtain the total volume of BN ejected from the channel walls at 543 hours of thruster operation. We can also compare

Table 8-4. Comparison of total sputtered volume by diagnostics.

**Calculated using standard error propagation from a spatial measurement uncertainty of 0.05 mm.*

***Calculated using standard error propagation from the uncertainty at each operating condition.*

Methodology	Total volume of sputtered BN, mm^3	Relative uncertainty
Profilometry at 543 hours	8437	8%*
CRDS measurement	4249	56%**
LIF, wall probe, and Gamero 2005	3.9e4	-
LIF, wall probe, and Rubin 2009	1.1e5	-

these results to the volume of ejected material measured from profilometry. The volume of ejected material measured from profilometry can be calculated by geometry. The comparison of results is shown in Table 8-4.

The most accurate method to measure the total volume of sputtered BN is, predictably, profilometry. The next most accurate approach is CRDS being in the right order of magnitude but still off by 50%. The least accurate approach is combining LIF, wall probe data, and sputtering model, which produced results that are an order of magnitude too large. However, profilometry took much more operational time and resources to obtain results than CRDS did.

Several important observations can be made from the comparison in Table 8-4. First, the CRDS results appear to under predict the ejected BN volume. Some possible reasons have already been discussed in Section 7.4.5 and include poor velocity modeling, ionization of neutral boron before they are detected, boron that exits the thruster in BN molecular complex, and possible incompleteness in the scanned lineshape. Also, it is possible that the record of operating hours is incomplete as the timing record is only accurate to a few minutes per entry. With that said, the total uncertainty in the operating hours should be no more than 10 hours or ~2%.

Another major possibility is that the erosion rate has changed between when the thruster was fresh and when it achieved 543 hours of operation. Erosion rate data from the SPT-100 [5] replicated in Figure 8-2 (same as Figure 5-75) clearly show a rapid decrease in erosion rate during the first 800 hours of thruster operation. It would therefore not be surprising if a similar phenomenon is happening in the 6-kW thruster thus giving rise to under-prediction when using the CRDS data to calculate the total ejected volume of BN. A more rigorous study involving taking CRDS and profilometry data at short regular intervals is needed to ascertain the influence of the above phenomenon.

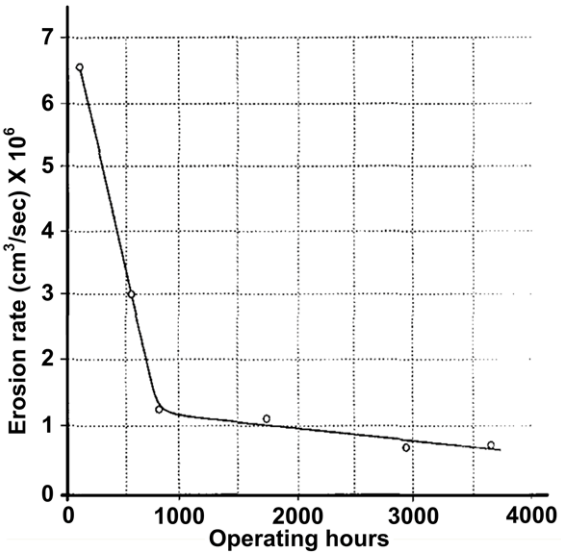

Figure 8-2. Erosion rate versus operating hours for the SPT-100 [5] (same as Figure 5-75).

Second, using LIF and wall probe data with a sputter model is not a good way to calculate the sputter rate as each component introduces a large amount of error and the final results can be very different from reality. Nevertheless, the exercise does show us that there is still much we do not understand about modeling erosion physics or even how to go about studying erosion physics.

With all of the above said, the author believes the results shown in Table 8-4 is a very good sign. The fact that CRDS result is in the same order of magnitude as the profilometry results and the fact that we know many of the shortcomings and physical effects that can potentially account for the difference in results mean that CRDS really does have the potential to become a very unique and highly accurate tool for studying Hall thruster channel wall erosion. Furthermore, while the uncertainty in the absolute sputter rate measured by CRDS is high, the uncertainty in the relative sputter rate is much lower. We will now move on to correlation of the sputter rate to key Hall thruster parameters using the relative sputter rate measurements.

8.4 The Importance of the Acceleration Zone Length for Channel Wall Erosion

At this point we are ready to review the hypothesis and the associated basic assumptions posed at the start of the dissertation. For convenience, the assumptions in the hypothesis are relisted here. One, the sputter yield of channel wall material is proportional to the energy of the bombarding particle. Two, the ion flux to the channel wall is proportional to the discharge current. Three, the acceleration zone does not change in structure and does not move. These assumptions give rise to a very simple scaling relation that says the erosion rate scales with the product of the discharge voltage and the anode mass flow rate.

Assumption one seems roughly accurate based on Garnier's sputter yield data [36] but more recent research [21] suggests the possibility that the relation between sputter yield and ion energy is not linear. Assumption one also does not take into account the threshold energy the value of which can greatly change any calculated sputter yield. This dissertation does not contain data that directly address assumption one and so it will be left for other research projects to address.

Assumption two also appears to be roughly true based on the CRDS data obtained for the three 300 V operating conditions. To a first order approximation, the measured relative sputtered rate appears to scale with the anode mass flow rate.

Assumption three is divided into two parts. The LIF data has shown that the acceleration zone structure does appear to be universal for a given magnetic field topology. However, there is also clear evidence that the acceleration zone moves with the discharge voltage. Any movement correlated to the mass flow rate is much less obvious. Indeed, the LIF data shows very significant changes in the extent and location of the acceleration zone with discharge voltage. When the discharge voltage was increased from 300 V to 600 V while the anode mass flow rate was kept constant at 10 mg/s, the axial length of the channel wall exposed to high energy ions almost doubled, as seen in Figure 5-57. Thus, it has become clear that the length of the acceleration zone at the channel walls is a key parameter that drives the channel wall erosion rate.

Admittedly, the definition of the acceleration zone does not lend itself to describing the wall erosion problem. A better definition for the region of the channel

walls exposed to erosion would be to take the location where the ions have enough energy to cause sputtering as the starting point (upstream point) and the channel exit plane as the ending point (downstream point). This new zone will be called the erosion zone.

We will now attempt to establish the value of the erosion zone length for the 10 mg/s operating condition in order to establish an engineering relationship between the discharge voltage and the erosion zone length. The threshold energy is set to 50 eV, the sheath energy is set to 35 eV, which is roughly the average of the sheath energy shown in Table 8-1. The difference, 15 eV, is the minimum pre-sheath kinetic energy that the ion must have to cause erosion. The location extrapolated to have this average ion kinetic energy is the starting point of the erosion zone. The ending point is the exit plane. For this exercise, we will be ignoring the channel wall shape and use only the axial distance between the starting and ending points as the erosion zone length.

Figure 8-3 shows the same information as Figure 5-57 except the speed axis has been converted to equivalent energy axis and a dotted line representing the 15-eV threshold has been added. The ion energy along the inner and outer walls is calculated from the Z = 0.033 and the Z = 0.967 data, respectively. The energy axis of the plot was purposely set to cut off at 50 eV to highlight the intersection between the lines that connect the data points and the 15 eV threshold line. Table 8-5 shows the result of the

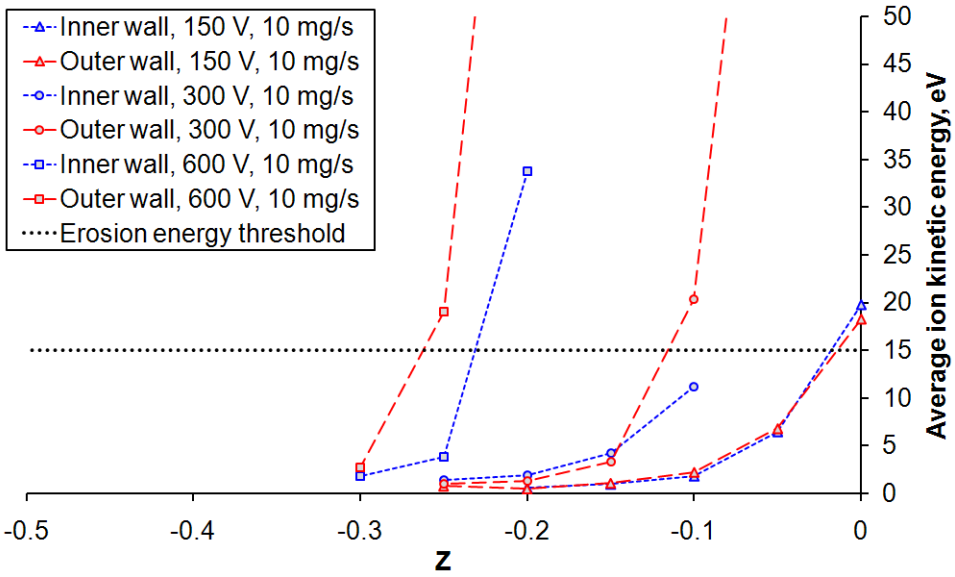

Figure 8-3. Average near-wall ion kinetic energy as a function of axial location.

Table 8-5. Axial erosion zone length calculated from near-wall LIF data.

Discharge voltage, V	Anode mass flow rate, mg/s	Channel wall	Non-dimensional axial erosion zone length
150	10	Inner	0.018
150	10	Outer	0.014
300	10	Outer	0.116
600	10	Inner	0.231
600	10	Outer	0.262

axial erosion zone length calculations.

From Table 8-5, we see that the area exposed to erosion is much smaller for the 150 V, 10 mg/s condition than for other conditions. Indeed, the CRDS sensor was unable to measure a statistically meaningful amount of erosion product at 150 V, 10 mg/s. We will ignore the results of this condition for the purpose of correlating the erosion zone length to the discharge voltage while keeping in mind that the erosion rate for 150 V, 10 mg/s must be very small relative to all other conditions. Since not enough data are available to calculate the erosion zone length for the 300 V, 10 mg/s condition at the inner wall, we will simply compare the values at the outer wall for different discharge voltages. Table 8-5 shows that operation at 600 V gave rise to an erosion zone ~2.3 times as long as the operation at 300 V for an anode mass flow rate of 10 mg/s. Unfortunately, the exact value of the threshold energy is a parameter that is still being actively researched and the LIF data near the walls are not high enough in resolution and fidelity for the above calculation to be very accurate. Furthermore, it is unclear whether the relationship between the erosion zone length and the discharge voltage holds for all anode mass flow rates. For the time being, we have established that the erosion zone length changes greatly with discharge voltage and is very important for channel wall erosion physics.

8.5 A Simple Engineering Model for Relative Erosion Rate

Having studied the importance of the erosion zone length and the weaknesses of the hypothesis, we will now attempt to create a new engineering relationship for correlating the erosion rate to the discharge voltage and the anode mass flow rate. The goal is a simple first-order relationship that can be used to predict erosion rate across

multiple operating conditions for a given Hall thruster with a given magnetic field topology.

The simple engineering erosion model on which the new correlation equation will be based is illustrated in Figure 8-4. The figure shows a close-up cut-away view of the inner channel wall near the channel exit of a Hall thruster. Various colored circles denoting bombarding xenon ion and sputtered products are shown to illustrate the basic mechanism. The Z-R axis is shown for directional reference. The main concept and assumption of the simple engineering erosion model is that the sputter rate profile along channel wall can be described by only two parameters, E_0 and Z_T. E_0 is the maximum energy of the bombarding particle reached at the channel exit plane, and Z_T is the axial start location of the erosion zone and its magnitude is the axial erosion zone length. This concept holds true if the near-wall acceleration zone profiles are roughly universal, which they have been shown to be from LIF data. Next, it is assumed that the sputter yield scales linearly with the ion energy for ion energies that exceed the threshold. If this is true, the red curve in Figure 8-4 also represents the sputter yield as a function of axial location. Lastly, the ion current density to the wall will be assumed to be roughly constant for a given mass flow rate in the erosion zone. While not completely true, Shastry's wall probe data [61] shows that this is an adequate first order approximation.

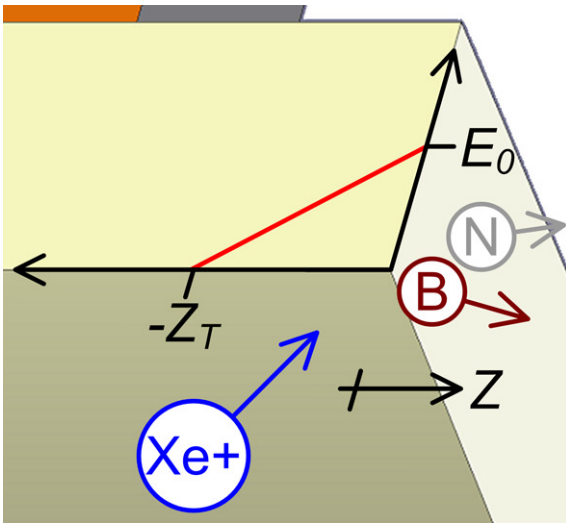

Figure 8-4. Illustration of the concepts behind the simple engineering erosion model.

Taking all the above assumptions together, one can form a scaling relationship for the total sputter rate in terms of the erosion zone length, the maximum particle energy, and the ion current density to the wall. This is shown in equation (8-6),

$$SR \sim Z_T E_0 j_W \tag{8-6}$$

where SR is the sputter rat and j_W is the ion current density to the wall. From previously shown empirical data for 10 mg/s operating conditions, doubling the discharge voltage roughly doubles the erosion zone length so we will assume that Z_T scales with the

discharge voltage. The maximum bombardment energy is also assumed to scale with the discharge voltage. Next, since the LIF data cannot establish definitively how much the erosion zone moves with changing anode mass flow rate, Z_T is assumed to be independent of the anode mass flow rate. On the other hand, the ion current density to the wall is assumed to scale linearly with the anode mass flow rate. The final scaling relationship for the channel wall sputter rate for a given magnetic field topology is shown in equation (8-7),

$$SR \sim V^2 \dot{m}_A \tag{8-7}$$

where V is the discharge voltage and $\dot{m}_A$ is the anode mass flow rate. Note that equation (8-7) uses the anode mass flow rate as one of the driving parameter as oppose to equation (1-1), which use the discharge current. To a first-order approximation, the discharge current scales with the anode mass flow rate and either could be used as the correlation parameter to scale the ion current density to the walls. Nevertheless, the anode mass flow rate is picked for engineering modeling in this dissertation because, one, it is one of the control parameters used in the dissertation throttling table, and two, the discharge current includes electron current, which does not carry enough momentum to cause sputtering. To check the validity of the scaling equation, the sputter rate measured by CRDS is normalized by the right-hand-side of equation (8-7). The result is shown in Table 8-6. In this table, the relative uncertainty and the deviation of the normalized sputter rate from the average across all seven operating conditions are shown.

The simple engineering erosion model appears to work remarkably well but one

Table 8-6. Validation of the simple engineering model for relative erosion rate.

Discharge voltage, V	Anode mass flow rate, mg/s	CRDS, sputter rate, mm^3/s	$SR/(V^2 \dot{m}_A)$, $mm^3/(V^2\text{-mg})$	Relative uncertainty	Deviation from average
150	10	(3.7e-4)	1.64e-9	100%	9%
150	20	(5.7e-4)	1.27e-9	100%	-16%
150	30	1.1e-3	1.63e-9	39%	8%
300	10	1.5e-3	1.67e-9	39%	11%
300	20	2.3e-3	1.28e-9	17%	-15%
300	30	3.7e-3	1.37e-9	7%	-9%
600	10	6.0e-3	1.67e-9	21%	11%

must be cautious because there is a great amount of uncertainty in the CRDS results. The fact that the results fit well with a discharge-voltage-squared dependence could be a coincidence and many more tests will need to be performed to check the results. It is also difficult to generalize this result to any other situation. For example, the length of the channel walls protruding downstream of the magnetic circuit may be an important design parameter that will change the amount of wall area exposed to the Hall thruster plasma, and therefore the scaling relationship. Thus, the most important result we should take away from the above model is not the specific scaling law that fits the data but how the scaling law was generated. In particular, it is the scaling of the erosion zone length with the discharge voltage and the anode mass flow rate that is the key to creating the scaling equation. One should keep in mind that the dependence of the erosion zone length on the discharge voltage and the anode mass flow may vary greatly from thruster to thruster and through different stages in the life of a thruster. Indeed, the next important problem to study in Hall thruster channel wall erosion physics is why does the erosion zone length increases with discharge voltage and how do we predict the value of that length. For future works, many more CRDS studies of the type described in this dissertation should be carried out on many different thrusters across many different operating conditions and stages of life to gain a better understanding of how the erosion zone length changes.

Having arrived at a correlation equation from a half-empirical, half-physical standpoint, we will now attempt to write a similar but purely phenomenological equation. This time, we will propose a more general form for the engineering sputter rate equation. This form is shown in equation (8-8),

$$SR \sim V^{x} \dot{m}_{A}^{\ y} \tag{8-8}$$

where the discharge voltage is raised to the power of x and the anode mass flow rate is rated to the power of y. To find the value of x, we perform power-law curve-fit on data from pairs of operating conditions having the same anode mass flow rate (i.e. 150 and 300 V at 30 mg/s and 300 and 600 V at 10 mg/s). The value of x calculated this way is then averaged. The data from 150 V, 10 and 20 mg/s conditions are not used because they do not represent positive boron detection. The value of y is found by performing power-law curve-fit on the data from the 300 V operating conditions. Figure 8-5 shows the

curve-fit for the coefficient x, and Figure 8-6 shows the curve-fit for the coefficient y. The final result is shown in equation (8-9).

$$SR \sim V^{1.9} \dot{m}_A^{0.8} \tag{8-9}$$

Interestingly, the curve-fit produced a value for the exponent of the anode mass flow rate at 0.8. Since this value is less than 1, sputter rate increases a bit less slowly than linear with the anode mass flow rate. This may provide evidence that the erosion zone is indeed being pushed downstream as the anode mass flow rate increases. There is a lot of uncertainty in the values of x and y and, due to the limited quantity of data available, it is not meaningful to try to establish the value of the uncertainty. Nevertheless, the general result is very similar to equation (8-7).

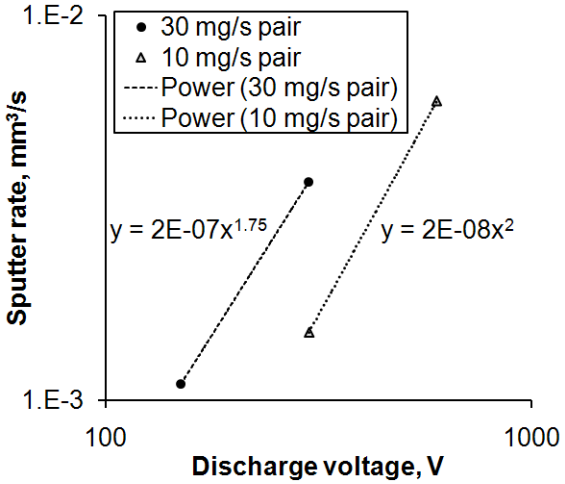

Figure 8-5. Curve-fit for the exponent of the discharge voltage in the general relative erosion rate equation.

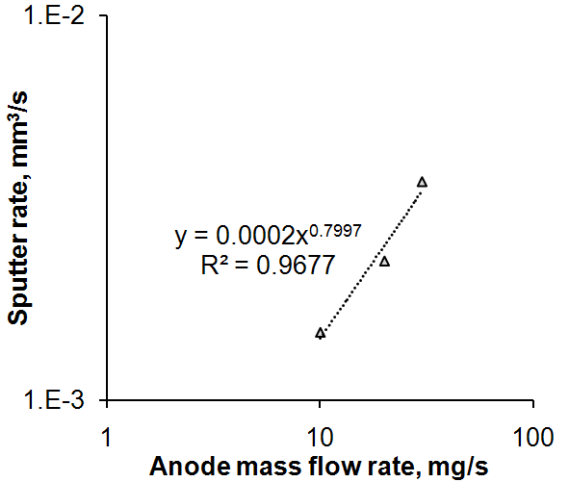

Figure 8-6. Curve-fit for the exponent of the anode mass flow rate in the general relative erosion rate equation.

8.6 Additional Erosion Physics not included in the Simple Model

The simple engineering erosion model, being a first-order model, does not account for many complex phenomena that could affect the erosion rate. Three of the phenomena known to have the potential to greatly alter erosion physics will be discussed in this section. They are the magnetic field topology, the divergence of the plume, and the presence of multiply-charged species.

The magnetic field topology plays a vital role in the determination of the erosion zone length. The importance of the magnetic field was not studied in this dissertation because the addition of the magnetic field topology as a control parameter would have made the scope of this work too broad for the resources available. Note that while the topology was not varied the field strength did change between operating conditions to maximize anode efficiency. It is likely that much of the results found in this dissertation could change should the magnetic field topology be a variable. For example, if the magnetic field topology is designed such that the bulk plasma is completely downstream of the exit plane for 300 V operation (assuming that this is possible) but that the plasma extends back into the channel when a thruster with such a topology operates at 600 V, the dependence of the erosion zone length on the discharge voltage would exhibit a threshold-like behavior. The total erosion rate would actually be zero below a certain discharge voltage and thus not be well-described by equation (8-8). The relationship between the erosion zone length and the magnetic field topology is an important topic for future research.

Plume divergence and the associated angle-of-incidence of the xenon particles is another factor that the simple engineering model does not account for. Table 8-1 shows that the angle-of-incidence with which xenon ions approach the channel walls can change greatly and nonlinearly with the discharge voltage. The AOI increases from ~30° to ~60° when the discharge voltage is increased from 150 V to 300 V, but then only increases an additional ~10° to ~70° when the discharge voltage is increased from 300 V to 600 V. Given the complex relationship between the AOI and sputter yield, it is unclear how the erosion rate should change as a result of changing plume divergence.

The presence of multiply-charged species can also greatly influence the scaling of the erosion rate with the discharge voltage and the anode mass flow rate. Simplistically speaking, a doubly-charged ion should acquire twice as much kinetic energy before hitting the channel walls than a singly-charged ion, a triply-charged ion should acquire three times as much energy, and etc. Furthermore, ExB probe data by Reid, et al., shows that the ratio of multiply-charged to singly-charged ion increases with both discharge voltage and anode mass flow rate for the 6-kW Hall thruster [51]. The main results of their work are reproduced in Table 8-7 for reference. Since the fraction of multiply-charged species increase with voltage and multiply-charged species have greater bombardment energy, it would appear that part of the increase in erosion rate with increasing discharge voltage can be attributed to multiply-charged species.

However, reality is likely to be much more complicated. Multiply-charged species are not born as far upstream as singly-charged species. This is due to the fact that the probability of a neutral atom being stripped of X electrons in one impact is generally much lower than the probability of a neutral undergoing X impacts each stripping away one electron (where X = 2+). It is also likely that multiply-charged species will bombard the wall with different AOI than singly-charged species.

Furthermore, while ExB data shows as much as 24% doubly-charged and 7% triply-charged species for the 600 V, 10 mg/s operating condition, this measurement is made of the plume in general and does not give information about the multiply-charged

Table 8-7. Summary of ExB probe data for the 6-kW Hall thruster, reproduced from work by Reid, et al. [51].

All data shown in this table were taken at 8 thruster diameters away from the thruster.

		Species fraction		
Discharge voltage, V	Anode mass flow rate, mg/s	Singly-charged	Doubly-charged	Triply-charged
150	10	0.971	0.072	0.018
150	20	0.907	0.120	0.033
150	30	0.811	0.198	0.052
300	10	0.922	0.105	0.034
300	20	0.795	0.193	0.073
300	30	0.632	0.347	0.081
600	10	0.746	0.242	0.073

species fractions near the channel walls. One argument for the lack of multiply-charged species near the channel walls would be that the density of the Hall current is much lower near the walls than near the center of channel. The reduction in electron density greatly decreases the probability of a neutral undergoing multiple ionizing impacts when compared to the probability of a neutral undergoing a single ionizing impact.

There is also physical evidence suggesting that multiply-charged species may not be an important factor, at least to a first-order approximation. Table 8-7 shows that the doubly-charged species fraction increases (almost linearly) with the anode mass flow rate, reaching 0.347 for the 300 V, 30 mg/s operating condition. Such a large increase in doubly-charge fraction should imply a large increase in the average energy of the particles bombarding the channel walls. The sputter yield should therefore increase with the anode mass flow rate, and the exponent to which the anode mass flow rate is raised in the general scaling equation (labeled as y) should be greater than one. Yet, curve-fit to the CRDS erosion data showed a value for y of 0.8 (Figure 8-6). To further complicate matters, the fact that y is less than 1 may be due to the acceleration zone moving downstream with increasing anode mass flow rate. This effect competes with any increase in the erosion rate due to increases in multiply-charged species fractions. At the present, there is simply not enough data to form any quantitative conclusions about the effect of multiply-charged species on the physics of Hall thruster channel wall erosion.

.

Chapter 9
Conclusion

9.1 Summary

This dissertation described successful demonstrations of the use of optical diagnostics, both laser-induced fluorescence velocimetry and cavity ring-down spectroscopy, to obtain a clearer physical picture of the Hall thruster channel wall erosion problem. The novel use of two-axis LIF velocimetry inside of a 6-kW Hall thruster produced maps of the velocity distribution functions of the ions bombarding the channel walls. This is the first published instance of the use of LIF to study the bombarding particles near the discharge channel walls in a Hall thruster environment. Additionally, the dissertation successfully demonstrated the first use of boron CRDS to study channel wall erosion in a Hall thruster environment to measure the channel wall erosion rate. The combined data from the two optical diagnostics highlighted the importance of the erosion zone length as a key parameter that drives channel wall erosion.

The length of the erosion zone was found to scale roughly with the discharge voltage of the Hall thruster. There are also subtle signs that the erosion zone length decreases slightly with increasing anode mass flow rate. A simple engineering erosion model was formulated to produce a scaling equation that can be used to correlate relative erosion rate across different discharge voltages and anode mass flow rates for the 6-kW Hall thruster with symmetric magnetic field topology towards the beginning of life. The equation shows that, to a first-order approximation, the erosion rate scaled linearly with the anode mass flow rate and with the square of the discharge voltage.

However, the key result of this dissertation is not the determination of the power to which the discharge voltage is raised. Instead, it is the method used to derive the scaling equation. The said method highlights the importance of determining the scaling relationship between the length of the erosion zone and the controlling parameters. This

method has not been previously used because the diagnostics needed to study the aforementioned scaling relationship has not been previously available. Thus, it can be said that the truly original contribution of this dissertation to the study of the channel wall erosion problem is the development of the two optical diagnostics and a demonstration of their potential use.

The optical diagnostics used in this dissertation were also found to have potentials to aid flight qualification in manners beyond just scientific studies. LIF velocimetry was found to be an accurate and non-intrusive tool for determining where the erosion zone starts. This information may become valuable for future effort to design magnetic field topology that pushes the erosion zone downstream of the axial position where the magnetic circuitry lies. CRDS was found to give near-real-time data that can save a vast amount of testing resources if the information of interest is relative erosion rate. Once the diagnostics is more matured, it may even make sense to use CRDS to perform accelerated channel wall erosion testing.

9.2 Future Work

To continue the research performed in this dissertation, future works should be focused on two areas, improving the diagnostics and performing similar studies on more thrusters at different stages in life.

For the LIF diagnostic, one important area of improvement would be increasing the signal-to-noise ratio so that data can be obtained closer to the walls and at operating conditions with oscillatory plasma. The relative lack of data for 30 mg/s operating conditions compared to other conditions illustrates the importance of obtaining data with higher SNR.

Another area of improvement for the LIF diagnostics is to improve the spatial resolution, so that the near-wall acceleration zone can be more finely mapped. This is, by itself, not difficult to accomplish as we can simply use a denser data grid. However, increasing data grid density means increasing data acquisition time so improvement in data rate can be very helpful for obtaining better spatial resolution.

For the CRDS diagnostic, there are many areas of improvements and they will be described in order of decreasing importance. The first area of improvement involves the

sensitivity of the cavity. Currently, we can only obtain an SNR of ~10 for the operating conditions with relatively high sputter rates. This SNR is not high enough to properly spatially resolve the exiting boron or to measure low-sputter-rate conditions like 150-V conditions. There are several ways to improve the SNR of CRDS. We can improve the mirror reflectivity, the strength of the signal, the ring-down and data acquisition rates, and the acquisition time. Currently, the acquisition time is limited by the degradation of the mirrors, so stopping or slowing down the degradation of the mirrors is an important next step. Adding a piezo-element to the cavity to increase the ring-down rate is another improvement to consider.

The second area of improvement for CRDS involves the modeling of the axial velocity along the measurement plane. The best approach would be to use boron LIF velocimetry to measure the axial component of the neutral boron being ejected. Another approach would be to use a more accurate 3D velocity model. The data for this model would, preferably, come from more controlled experiments to measure sputtered boron VDFs at bombardment energies closer to what is typical in Hall thrusters (50-500 eV).

The third area of improvement for CRDS involves the conversion of relative sputter rate to absolute sputter rate. Many key assumptions about quantities like the fraction of boron exiting the channel that is in an ionized form, the fraction of boron that are sputtered in molecular form, and the relative ratio of the two boron ground states need to be checked.

The fourth area of improvement for CRDS involves converting to an optical fiber beam deliver scheme. Using such a scheme would remove the limitation in wavelength scan range currently imposed by the use of direct beam injection. Making this improvement will increase our measurement fidelity because we can then make certain that the whole boron spectral lineshape is properly captured.

The fifth area of improvement for CRDS is more of an engineering improvement. The isolation of the cavity from vibration should be improved in order to make the tuning of the cavity easier. This improvement is not pressing as this dissertation describes how to keep the cavity aligned even with mild vibration present. However, it is a very useful improvement for making CRDS more user-friendly and may be a necessary improvement

if the cavity length is to be driven by a piezo-element (as opposed to mechanical vibrations).

Two types of future works are very helpful to the continuation of the scientific effort described in this dissertation. First, it is imperative that more data be collected at different high-voltage operating conditions (300+ V) because the creation of a better engineering scaling relationship for erosion requires more high-voltage data. This was demonstrated in this dissertation when an attempt was made to create such a scaling relationship.

Second, more studies of the type described in this dissertation should be performed on various thrusters and at various stages of life. By comparing the erosion data across different thrusters we can determine whether the scaling relationship for predicting erosion rate at different operating conditions is universal across thrusters. By comparing the erosion data across different stages of life for the Hall thruster, we can determine whether and how the erosion physics changes as the channel erodes. These data will be very useful for developing models and design tools for Hall thrusters.

In addition, a wide variety of scientific investigations has now been enabled by a developed boron CRDS erosion diagnostic that was previously impractical to perform. For example, we can now measure the channel wall erosion rate as a function of changing magnetic field topology in an attempt to search for a topology that minimizes erosion rate (as opposed to maximize performance). We can study the effect of different wall materials (this may require multiple lasers and mirror sets) to identify materials that are optimal from an erosion standpoint. There are a vast number of possible scientific ventures now available for exploring.

Furthermore, it would be interesting and highly valuable to Hall thruster erosion research to investigate how exactly the Hall thruster magnetic circuit erodes once exposed to plasma. The community has repeated cited the erosion of the channel walls and subsequent exposure of the magnetic circuit to plasma as a primary failure mode in the Hall thruster. However, little is known about the duration of time between the beginning of circuit exposure to plasma and actual failure of the thruster. A simple experiment to study the behavior of this failure mode could be performed by shaping the channel based on existing channel profiles such that the magnetic circuit is exposed.

Ultimately, the study of the Hall thruster channel wall erosion needs to be tied back to the basic underlying physics so that accurate, predictive models can be built to help future design works. In particular, a very important next step forward is to decipher the physical basis by which the erosion zone length increases with the discharge voltage. There is also the important question of how the erosion zone length changes with the magnetic field topology. Understanding these mechanisms are the keys to creating a model that can correlate erosion rate across different operating conditions for a Hall thruster, and eventually will lead to a model for predicting absolute erosion rate.

Bibliography

[1] Hofer, R. R., "Development and Characterization of High-Efficiency, High-Specific Impulse Xenon Hall Thrusters", Ph.D. Dissertation, Aerospace Engineering, University of Michigan, Ann Arbor, MI, 2004.

[2] Brown, D. L., "Investigation of Flow Discharge Voltage Hall Thruster Characteristics and Evaluation of Loss Mechanisms", Ph.D. Dissertation, Aerospace Engineering, University of Michigan, Ann Arbor, MI, 2009.

[3] Oleson, S. R., "Mission Advantages of Constant Power, Variable Isp Electrostatic Thrusters", *36th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-2000-3413, Huntsville, AL, 16-19 Jul., 2000.

[4] Oh, D., "Evaluation of Solar Electric Propulsion Technologies for Discovery Class Missions ", *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-2005-4270, Tucson, AZ, 11-13 Jul., 2005.

[5] Absalamov, S. K., et al., "Measurement of Plasma Parameters in the Stationary Plasma Thruster (SPT-100) Plume and Its Effect on Spacecraft Components", *28th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-1992-3156, Nashville, TN, 6-8 Jul., 1992.

[6] Kamhawi, H., Manzella, D. H., and Peterson, P. Y., "High Voltage Hall Accelerator Wear Test Update", *55th Joint Army-Navy-NASA-Air Force Propulsion Meeting*, Newton, MA, 12-16 May, 2008.

[7] de Grys, K., Mathers, A., Welander, B., and Khayms, V., "Demonstration of 10,400 Hours of Operation on a 4.5 kW Qualification Model Hall Thruster", *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2010-6698, Nashville, TN, 25-28, Jul., 2010.

[8] Manzella, D., Yim, J., and Boyd, I., "Predicting Hall Thruster Operational Lifetime", *40th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2004-3953, Fort Lauderdale, FL, 11-14 Jul., 2004.

[9] Sommier, E., Scharfe, M. K., Gascon, N., Capelli, M. A., and Fernandez, E., "Simulating Plasma-Induced Hall Thruster Wall Erosion With a Two-Dimensional Hybrid Model", *IEEE Transactions on Plasma Science*, Vol. 35, No. 5, doi:10.1109/TPS.2007.905943, Oct., 2007, pp. 1379-1387.

[10] Yim, J. T., "Computational Modeling of Hall Thruster Channel Wall Erosion", Ph.D. Dissertation, Aerospace Engineering, University of Michigan, Ann Arbor, MI, 2008.

[11] Mikellides, I. G., et al., "Magnetic shielding of the channel walls in a Hall plasma accelerator", *Physics of Plasmas*, Vol. 18, No. 3, doi:10.1063/1.3551583, Mar. 8, 2011, pp. 033501.

[12] Cheng, S. Y.-M., "Modeling of Hall Thruster Lifetime and Erosion Mechanisms", Ph.D. Dissertation, Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, 2007.

[13] Goebel, D. M. and Katz, I., *Fundamentals of Electric Propulsion: Ion and Hall Thrusters*, JPL Space Science and Technology Series, ISBN: 9780470429273, Wiley, 2008.

[14] Petrosov, V. A., et al., "A 2000 Hours Lifetime Test Results of 1.3 kW T-100 Electric Thruster", *24th International Electric Propulsion Conference*, IEPC-1995-041, Moscow, Russia, Sep., 1995.

[15] Mason, L. S., Jankovsky, R. S., and Manzella, D. H., "1000 Hours of Testing on a 10 Kilowatt Hall Effect Thruster", *37th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2001-3773, Salt Lake City, UT, 8-11 Jul., 2001.

[16] Marchandise, F. R., et al., "The PPS1350 Qualification Demonstration 7500h on Ground, about 5000h in Flight", *29th International Electric Propulsion Conference*, IEPC 2005-209, Princeton, NJ, 31 Oct. - 4 Nov., 2005.

[17] Yalin, A. P., Surla, Y., Farnell, C., Butweiller, M., and Williams, J. D., "Sputtering Studies of Multi-Component Materials by Weight Loss and Cavity Ring-Down Spectroscopy", *42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2006-4338, Sacramento, CA, 9-12 Jul., 2006.

[18] Rubin, B., Topper, J. L., and Yalin, A. P., "Total and Differential Sputter Yields of Boron Nitride Measured by Quatz Crystal Microbalance and Weight Loss", *30th International Electric Propulsion Conference*, 2007-074, Florence, Italy, 17-20 Sep., 2007.

[19] Peterson, P. Y. and Manzella, D. H., "Investigation of the Erosion Characteristics of a Laboratory Hall Thruster", *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2003-5005, Huntsville, AL, 20-23 Jul., 2003.

[20] Misuri, T., Milani, A., and Andrenucci, M., "Development of a Telemicroscopy Diagnostic Apparatus and Erosion Modelling in Hall Effect Thrusters", *31st International Electric Propulsion Conference*, 2009-036, Ann Arbor, MI, 20-24 Sep., 2009.

[21] Rubin, B., Topper, J. L., and Yalin, A. P., "Total and Differential Sputter Yields of Boron Nitride Measured by Quartz Crystal Microbalance", *31st International Electric Propulsion Conference*, 2009-042, Ann Arbor, MI, 20-24 Sep., 2009.

[22] Pagnon, D., Balika, L., and Pellerin, S., "QCM and OES: two ways used to study simultaneously HET thruster chamber ceramic erosion. First QCM results on PPS100-ML validate previous OES measurements." *31st International Electric Propulsion Conference*, 2009-118, Ann Arbor, MI, 20-24 Sep., 2009.

[23] Tao, L., Lee, B., Yalin, A., Yamamoto, N., and Gallimore, A. D., "Development of a Cavity Ring-Down Spectroscopy Sensor for Boron Nitride Erosion in Hall Thrusters", *31st International Electric Propulsion Conference*, 2009-146, Ann Arbor, MI, 20-24 Sep., 2009.

[24] Yamamoto, N., Tao, L., Rubin, B., Williams, J. D., and Yalin, A. P., "Sputter Erosion Sensor for Anode Layer-Type Hall Thrusters Using Cavity Ring-Down Spectroscopy", *Journal of Propulsion and Power*, Vol. 26, No. 1, Jan-Feb, 2010, pp. 142-148.

[25] Manzella, D. H., "Stationary Plasma Thruster Ion Velocity Distribution", *30th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-1994-3141, Indianapolis, IN, 27-29 Jun., 1994.

[26] Hargus, W. A., Jr. and Cappelli, M. A., "Laser-induced fluorescence measurements of velocity within a Hall discharge", *Applied Physics B*, Vol. 72, No. 8, doi:10.1007/s003400100589, Jun., 2001, pp. 961-969.

[27] Hargus, W. A., Jr. and Nakles, M. R., "Ion Velocity Measurements within the Acceleration Channel of a Lower Power Hall Thruster", *30th International Electric Propulsion Conference*, 2007-172, Florence, Italy, 17-20 Sep., 2007.

[28] Yamamura, Y., "Theory of Sputtering and Comparison to Experimental Data", *Nuclear Instruments and Methods*, Vol. 194, No. 1-3, doi:10.1016/0029-554X(82)90575-4, Nov., 1982, pp. 515-522.

[29] Zhang, Z. L. and Zhang, L., "Anisotropic Angular Distribution of Sputtered Atoms", *Radiation Effects & Defects in Solids*, Vol. 159, No. 5, doi:10.1080/10420150410001724495, May, 2004, pp. 301-307.

[30] Yalin, A. P., Williams, J. D., Surla, V., and Zoerb, K. A., "Differential sputter yield profiles of molybdenum due to bombardment by low energy xenon ions at normal and oblique incidence", *Journal of Physics D: Applied Physics*, Vol. 40, No. 10, doi:10.1088/0022-3727/40/10/025, 4 May, 2007, pp. 3194.

[31] Sigmund, P., "Theory of Sputtering. I. Sputtering Yield of Amorphous and Polycrystalline Targets", *Physical Review*, Vol. 184, No. 2, doi:10.1103/PhysRev.184.383, Aug., 1969, pp. 383-416.

[32] Bohdansky, J., "A universal relation for the sputtering yield of monatomic solids at normal ion incidence", *Nuclear Instruments and Methods in Physics Research Section B*, Vol. 2, No. 1-3, doi:10.1016/0168-583X(84)90271-4, Mar., 1984, pp. 587-591.

[33] Parra, F. I., Ahedo, E., Fife, J. M., and Martinez-Sanchez, M., "A two-dimensional hybrid model of the Hall thruster discharge", *Journal of Applied Physics*, Vol. 100, No. 2, doi:10.1063/1.2219165, 26 Jul., 2006, pp. 023304.

[34] Hofer, R. R., et al., "Efficacy of Electron Mobility Models in Hybrid-PIC Hall Thruster Simulations", *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2008-4924, Hartford, CT, 21-23 Jul., 2008.

[35] Gamero-Castano, M. and Katz, I., "Estimation of Hall Thruster Erosion using HPHall", *29th International Electric Propulsion Conference*, IEPC 2005-303, Princeton, NJ, 31 Oct. - 4 Nov., 2005.

[36] Garnier, Y., Viel, V., Roussel, J. F., and Bernard, J., "Low-energy xenon ion sputtering of ceramics investigated for stationary plasma thrusters", *Journal of Vacuum Science and Technology*, Vol. 17, No. 6, doi:10.1116/1.582050, Nov/Dec, 1999, pp. 3246-3254.

[37] Kamhawi, H., Manzella, D., Pinero, L., Haag, T., and Huang, W., "In-Space Propulsion High Voltage Hall Accelerator Development Project Overview", *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2010-6860, Nashville, TN, 25-28, Jul., 2010.

[38] Cartmell, M. P. and McKenzie, D. J., "A review of space tether research", *Progress in Aerospace Sciences*, Vol. 44, No. 1, doi:10.1016/j.paerosci.2007.08.002, Jan, 2008, pp. 1-21.

[39] Johnson, L., Young, R., Montgomery, E., and Alhorn, D., "Status of solar sail technology within NASA", *Advances in Space Research*, doi:10.1016/j.asr.2010.12.011, Dec, 2010.

[40] Chin, G., et al., "Lunar Reconnaissance Orbiter Overview: The Instrument Suite and Mission", *Space Science Reviews*, Vol. 129, No. 4, doi:10.1007/s11214-007-9153-y, 2007, pp. 391-419.

[41] Brophy, J. R. and Rodgers, D. H., "Ion Propulsion for a Mars Sample Return Mission", *36th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-2000-3412, Huntsville, AL, 16-19 Jul., 2000.

[42] Whiffen, G. J., "An investigation of a Jupiter Galilean Moon Orbiter trajectory", *AAS/AIAA Astrodynamics Specialist Conference*, AIAA-2003-0544, Big Sky, MT, 3-7 Aug., 2003.

[43] Larson, W. J. and Wertz, J. R., *Space Mission Analysis and Design*, 3rd ed., ISBN: 9781881883104, 1999.

[44] 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, doi:10.2514/1.15952, Jul.-Aug., 2006, pp. 721-731.

[45] Morozov, A. I., "Focussing of Cold Quasineutral Beams in Electromagnetic Fields", *Soviet Physics Doklady*, Vol. 10, No. 8, Feb., 1966, pp. 775-777.

[46] Reid, B. M., "The Influence of Neutral Flow Rate in the Operation of Hall Thrusters", Ph.D. Dissertation, Aerospace Engineering, University of Michigan, Ann Arbor, MI, 2008.

[47] Brown, D. L., et al., "Performance Characterization and Design Verification of the H6 Laboratory Model Hall Thruster", *54th Joint Army-Navy-NASA-Air Force Propulsion Meeting*, Denver, CO, 14-17 May, 2007.

[48] 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, 14-17 May, 2007.

[49] Reid, B. M. and Gallimore, A. D., "Langmuir Probe Measurements in the Discharge Channel of a 6-kW Hall Thruster", *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2008-4920, Hartford, CT, 21-23 Jul., 2008.

[50] Reid, B. M. and Gallimore, A. D., "Plasma Potential Measurements in the Discharge Channel of a 6-kW Hall Thruster", *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2008-5185, Hartford, CT, 21-23 Jul., 2008.

[51] Reid, B. M., Shastry, R., Gallimore, A. D., and Hofer, R. R., "Angularly-Resolved ExB Probe Spectra in the Plume of a 6-kW Hall Thruster", *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2008-5287, Hartford, CT, 21-23 Jul., 2008.

[52] Shastry, R., Hofer, R. R., Reid, B. M., and Gallimore, A. D., "Method for analyzing ExB probe spectra from Hall thruster plumes", *Review of Scientific Instruments*, Vol. 80, No. 6, doi:10.1063/1.3152218, 22 Jun., 2009, pp. 063502.

[53] Smith, T. B., Huang, W., Reid, B. M., and Gallimore, A. D., "Near-field Laser-induced Fluorescence Velocimetry of Neutral Xenon in a 6-kW Hall Thruster Plume", *30th International Electric Propulsion Conference*, 2007-252, Florence, Italy, 17-20 Sep., 2007.

[54] Huang, W., Reid, B. M., Smith, T. B., and Gallimore, A. D., "Laser-Induced Fluorescence of Singly-Charged Xenon in a 6-kW Hall Thruster Plume", *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2008-5102, Hartford, CT, 21-23 Jul., 2008.

[55] Huang, W., Drenkow, B., and Gallimore, A. D., "Laser-Induced Fluorescence of Singly-Charged Xenon Inside a 6-kW Hall Thruster", *45th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2009-5355, Denver, CO, 2-5 Aug., 2009.

[56] 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*, 2009-113, Ann Arbor, MI, 20-24 Sep., 2009.

[57] Huang, W. and Gallimore, A. D., "Laser-induced Fluorescence Study of Neutral Xenon Flow Evolution inside a 6-kW Hall Thruster", *31st International Electric Propulsion Conference*, 2009-087, Ann Arbor, MI, 20-24 Sep., 2009.

[58] Huang, W., Ngom, B. B., and Gallimore, A. D., "Using Nonlinear Zeeman Spectroscopy to obtain In-Situ Magnetic Field Measurement in a Hall Thruster", *31st International Electric Propulsion Conference*, 2009-088, Ann Arbor, MI, 20-24 Sep., 2009.

[59] Shastry, R., Gallimore, A. D., and Hofer, R. R., "Near-Wall Plasma Characterization of a 6-kW Hall Thruster", *31st International Electric Propulsion Conference*, 2009-133, Ann Arbor, MI, 20-24 Sep., 2009.

[60] McDonald, M. S., Lobbia, R. B., and Gallimore, A. D., "High Speed Interrogation of the Near-Field Plume of a 6-kW Laboratory Hall Thruster", *31st International Electric Propulsion Conference*, 2009-112, Ann Arbor, MI, 20-24 Sep., 2009.

[61] Shastry, R., "Experimental Characterization of the Near-Wall Region in Hall Thrusters and its Implications on Performance and Lifetime", Ph.D. Dissertation, Aerospace Engineering, University of Michigan, Ann Arbor, MI, 2011.

[62] Haag, T. W., "Thrust stand for high-power electric propulsion devices", *Review of Scientific Instruments*, Vol. 62, No. 5, doi:10.1063/1.1141998, May, 1991, pp. 1186.

[63] Cedolin, R. J., Hargus, W. A., Hanson, R. K., and Cappelli, M. A., "Laser-Induced Fluorescence Diagnostics for Xenon Hall Thrusters", *32nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-1996-2986, Buena Vista, FL, 1-3 Jul., 1996.

[64] Cedolin, R. J., "Laser-Induced Fluorescece Diagnostics of Xenon Plasmas", Ph.D. Dissertation, Mechanical Engineering, Stanford University, Stanford, CA, 1997.

[65] Huang, W., Smith, T. B., and Gallimore, A. D., "Obtaining Velocity Distribution using a Xenon Ion Line with Unknown Hyperfine Constants", *40th AIAA Plasmadynamics and Laser Conference*, AIAA-2009-4226, San Antonio, Texas, 22-25 Jun., 2009.

[66] Candler, C., *Atomic Spectra and the Vector Model*, 2nd ed., Van Nostrand, Princeton, N. J., 1964, pp. 313-315.

[67] Hill, E. L., "Relative Intensities in Nuclear Spin Multiplets", *Proceedings of the National Academy of Sciences*, Vol. 15, No. 10, 15 Oct., 1929, pp. 779-784.

[68] D'Amico, G., Pesce, G., and Sasso, A., "High Resolution Spectroscopy of Stable Xenon Isotopes", *Hyperfine Interactions*, Vol. 127, doi:10.1023/A:1012666811805, Aug., 2000, pp. 121-128.

[69] Suzuki, M., Katoh, K., and Nishimiya, N., "Saturated Absorption Spectroscopy of Xe using a GaAs Semiconductor Laser", *Spectrochimica Acta Part A*, Vol. 58, No. 11, doi:10.1016/S1386-1425(02)00069-0, Aug., 2002, pp. 2519-2531.

[70] Jackson, D. A. and Coulombe, M. C., "Isotope Shifts in the Arc Spectrum of Xenon", *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, Vol. 338, No. 1614, doi:10.1098/rspa.1974.0086, Jun., 1974, pp. 277-298.

[71] Firestone, R. B., *Table of Isotopes*, 8th ed., edited by C.M. Baglin and S.Y.F. Chu, ISBN: 978-0471356332, Wiley-Interscience, 1996.

[72] Svanberg, S., *Atomic and Molecular Spectroscopy*, 4th ed., ISBN: 3-540-20382-6, Springer-Verlag, Berlin, 2004.

[73] Karimov, R. G. and Klimkin, V. M., "Xe I and Xe II Radiative Lifetimes and Transition Probabilities", *Russian Physics Journal*, Vol. 14, No. 3, doi:10.1007/BF00822259, Mar., 1971, pp. 308-311.

[74] Brostrom, L., Mannervik, S., Passian, A., and Sundstrom, G., "Investigation of some transitions and lifetimes in Xe II", *Physical Review A*, Vol. 49, No. 5, doi:10.1103/PhysRevA.49.3333, May, 1994, pp. 3333-3337.

[75] Harkness, H. W. and Heard, J. F., "The Stark Effect for Xenon", *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, Vol. 139, No. 838, doi:10.1098/rspa.1933.0027, Feb., 1933, pp. 416-435.

[76] Manola, S. and Konjevic, N., "Stark broadening of the singly ionized xenon line: Temperature variation", *Physical Review A*, Vol. 38, No. 11, doi:10.1103/PhysRevA.38.5742, Dec. 1, 1988, pp. 5742-5744.

[77] Ngom, B. B., "Magnetic Field Simulation and Mapping Based on Zeeman-Split Laser-Induced Fluorescence Spectra of Xenon in the Discharge Channel of 5-6 kW Coaxial Stationary-Plasma Hall Thrusters", Ph.D. Dissertation, Aerospace Engineering, University of Michigan, Ann Arbor, MI, 2009.

[78] Savitzky, A. and Golay, M. J. E., "Smoothing and Differentiation of Data by Simplified Least Squares Procedures", *Analytical Chemistry*, Vol. 36, No. 8, doi:10.1021/ac60214a047, Jul., 1964, pp. 1627-1639.

[79] Dushman, S., *Scientific Foundations of Vacuum Technique*, 1st ed., Ch. 4, Wiley, New York, 1958.

[80] Smith, T. B., "Deconvolution of Ion Velocity Distributions from Laser-Induced Fluorescence Spectra of Xenon Electrostatic Thruster Plumes", Ph.D. Dissertation, Aerospace Engineering, University of Michigan, Ann Arbor, MI, 2003.

[81] Petrov, G. M., "A simple algorithm for spectral line deconvolution", *Journal of Quantitative Spectroscopy & Radiative Transfer*, Vol. 72, No. 3, doi:10.1016/S0022-4073(01)00125-X, Feb., 2002, pp. 281-287.

[82] Fife, J. M., Martinez-Sanchez, M., and Szabo, J., "A Numerical Study of Low-Frequency Discharge Oscillations in Hall Thrusters", *33rd AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-1997-3052, Seattle, WA, 6-9 Jul., 1997.

[83] Huang, W., Gallimore, A. D., and Hofer, R. R., "Neutral Flow Evolution in a Six-Kilowatt Hall Thruster", *Journal of Propulsion and Power*, Vol. 27, No. 3, doi:10.2514/1.54141, May-Jun, 2011, pp. 553-563.

[84] Hargus, W. A., Jr. and Strafaccia, J., "Optical Boron Nitride Insulator Erosion Characterization of a 200 W Xenon Hall Thruster", *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-2005-3529, Tucson, AZ, 11-13 Jul., 2005.

[85] Surla, V., Wilbur, P. J., Johnson, M., Williams, J. D., and Yalin, A. P., "Sputter erosion measurements of titanium and molybdenum by cavity ring-down spectroscopy", *Review of Scientific Instruments*, Vol. 75, No. 9, doi:10.1063/1.1786354, Sep., 2004, pp. 3025.

[86] Yalin, A. P., Surla, V., Butweiller, M., and Williams, J. D., "Detection of sputtered metals with cavity ring-down spectroscopy", *Applied Optics*, Vol. 44, No. 30, Oct., 2005, pp. 6496-6505.

[87] Sigmund, P. i., in *Sputtering by Particle Bombardment I*, R. Behrisch, Editor, Springer, Berlin, 1981, pp. 9.

[88] Betz, G. and Wien, K., "Energy and angular distributions of sputtered particles", *International Journal of Mass Spectrometry and Ion processes*, Vol. 140, No. 1, doi:10.1016/0168-1176(94)04052-4, Dec., 1994, pp. 1-110.

[89] Tao, L. and Yalin, A., "LIF Velocity Measurement of Sputtered Boron Atoms from Boron Nitride Target", *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2010-6526, Nashville, TN, 25-28 Jul., 2010.

[90] Johansson, S. G., Litzen, U., Kasten, J., and Kock, M., "A 25 mA isotope shift in B I at 2090 A - A possible diagnostic of the cosmic B-11/B-10 ratio?" *Astrophysical Journal*, Vol. 403, No. 1, doi:10.1086/186713, 20 Jan., 1993, pp. L25-L28.

[91] Kramida, A. E. and Ryabtsev, A. N., "A critical compilation of energy levels and spectral lines of neutral boron", *Physica Scripta*, Vol. 76, No. 5, doi:10.1088/0031-8949/76/5/024, 16 Oct., 2007, pp. 544-557.

[92] Schaefer, H. F. and Klemm, R. A., "Atomic Hyperfine Structure. II. First-Order Wave Functions for the Ground States of B,C,N,O, and F", *Physical Review*, Vol. 181, No. 1, doi:10.1103/PhysRev.181.137, May, 1969, pp. 137-143.

[93] Fuhr, J. R. and Wiese, W. L., "Tables of Atomic Transition Probabilities for Beryllium and Boron", *Journal of Physical and Chemical Reference Data*, Vol. 39, No. 1, doi:10.1063/1.3286088, 12 Mar., 2010, pp. 013101.

[94] Zhang, J., Bhattacharjee, S., Shutthanandan, V., and Ray, R. K., "Sputtering investigation of boron nitride with secondary ion and secondary neutral mass spectrometry", *Journal of Vacuum Science and Technology A*, Vol. 15, No. 2, doi:10.1116/1.580519, Mar., 1997, pp. 243-247.

[95] Yalin, A. P., et al., "High-Sensitivity Boron Nitride Sputter Erosion Measurements by Continuous-Wave Cavity Ring-Down Spectroscopy", *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2008-5091, Hartford, CT, 21-23 Jul., 2008.

[96] Berden, G., Peeters, R., and Meijer, G., "Cavity ring-down spectroscopy: Experimental schemes and applications", *International Reviews in Physical Chemistry*, Vol. 19, No. 4, 2000, pp. 565-607.

[97] Hargus, W. A., Jr. and Charles, C. S., "Near Plume Laser Induced Fluorescence Velocity Measurements of a 600 W Hall Thruster", *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, AIAA-2008-5004, Hartford, CT, 21-23 Jul., 2008.

[98] Hanson, K. M., *A Bayesian approach to nonlinear inversion: Abel inverison from x-ray data*, Transport Theory, Invariant Imbedding, and Integral Equations, Lecture Notes in Pure and Applied Mathematics, edited by P. Nelson, et al., Vol. 115, Marcel Dekker, New York, 1989, 363-368.

[99] Dasch, C. J., "One-dimensional tomography: a comparison of Abel, onion-peeling, and filtered backprojection methods", *Applied Optics*, Vol. 31, No. 1, doi:10.1364/AO.31.001146, 1992, pp. 1146.

[100] Saleh, B. E. A. and Teich, M. C., *Resonator Optics*, in *Fundamentals of Photonics*, Ch. 10, John Wiley & Sons, Hoboken, NJ, 2007.

[101] Kogelnik, H. and Li, T., "Laser Beams and Resonators", *Applied Optics*, Vol. 5, No. 10, doi:10.1364/AO.5.001550, Oct., 1966, pp. 1550-1567.

[102] Verdeyen, J. T., *Laser Electronics*, Ch. 2, Prentice-Hall, Englewood Cliffs, N.J., 1994.

[103] Morville, J., Romanini, D., Chenevier, M., and Kachanov, A., "Effects of laser phase noise on the injection of a high-finesse cavity", *Applied Optics*, Vol. 41, No. 3, Nov. 20, 2002, pp. 6980.

[104] Lehmann, K. K., Berden, G., and Engeln, R., *An Introduction to Cavity Ring-Down Spectroscopy*, in *Cavity Ring-Down Spectroscopy*, G. Berden and R. Engeln, Editors, Wiley, Chippenham, U.K., 2009.

[105] Kim, Y.-K. and Stone, P. M., "Ionization of boron, aluminum, gallium, and indium by electron impact", *Physical Review A*, Vol. 64, No. 5, doi:10.1103/PhysRevA.64.052707, 4 Oct., 2001, pp. 052707.

[106] Haas, J. M., "Low-perturbation Interrogation of the Internal and Near-field Plasma Structure of a Hall Thruster Using a High-Speed Probe Positioning System", Ph.D. Dissertation, Aerospace Engineering, University of Michigan, Ann Arbor, MI, 2001.