Copyright
by
Adam Alexander Libson
2012

The Dissertation Committee for Adam Alexander Libson
certifies that this is the approved version of the following dissertation:

General Methods of Controlling Atomic Motion: Experiments with Supersonic Beams as a Source of Cold Atoms

Committee:

Mark G. Raizen, Supervisor

Michael F. Becker

Herbert L. Berk

John W. Keto

Greg O. Sitz

General Methods of Controlling Atomic Motion: Experiments with Supersonic Beams as a Source of Cold Atoms

by

Adam Alexander Libson, B.A.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

December 2012

For Kirsten

Acknowledgments

First and foremost, I would like to thank my advisor, Mark Raizen. It has been a privilege to work with Mark for my graduate career, and I am continually amazed by his creative approach to physics research. I have also been thankful for the supportive environment he has fostered in the lab, which makes the daily frustrations of experimental work much easier to bare.

I would also like to thank my wife, Kirsten. Having a fellow physicist at home, and my partner at work, led to additional stress, but the space provided by our work on different experiments made all the difference. My ability to discuss physics "after hours" with her provided me with the solutions to many problems, and her emotional support in the lab got me through quite a few crises.

Much of my time in the lab was spent with my co-workers on the beam slowing experiment, which has truly been a team effort. Tom Mazur is a gifted experimentalist, and has been a fantastic person to work with. Very few people hold themselves to as high a standard as Tom does, and this certainly shows in his work and in his approach to physics. Isaac Chavez started working in the lab at the same time I did, and his mechanical aptitude and no fear approach helped us in the lab on many occasions. Travis Bannerman joined the experiment after working with single photon cooling and laser cooling in general on the rubidium experiment, and this experience was essential in setting up the hydrogen laser, since he was without doubt our laser guru.

I also had the pleasure to work with two fantastic post-docs, Ed Narevicius and Rob Clark. Ed was a fantastic person to work with, and his guidance certainly

provided me with plenty of insight. Ed, additionally, makes a mean grilled rack of lamb, and I greatly enjoyed spending time with him, his wife Julia, and their son, Yoav. Rob worked with me on the hydrogen experiment, and his laser experience laser helped us to get UV light. I also had many fun evenings at Clark Tavern.

I also wish to thank each of the Wurzberg students I worked with during my time here. Max Riedel and Christoph Schaeff worked with me on the rotor experiment. Max worked with me getting the rotor off the ground, and I would like to thank him for his contributions there. More important to me though is the friendship he provided during my first year. Whether playing badminton, listening to music, or just eating ice cream, Max was a great guy to hang out with. Christoph worked to get the LiF aspect of the rotor working, and his devotion to the experiment and to the Austrian body building method both impressed me. Finally, I must thank Christian Parthey. Christian worked with me on the magnetic slowing of neon, and was a fantastic person to work with in the lab. He is also a great friend and someone I spent a lot of my non-lab time with while he was here. I would also like to point out that his skill set is not limited to the lab as I've observed some impressive moves on the rock.

As I am leaving, the hydrogen experiment will be run by Maria Becker and Tom. Maria started at the beginning of this year and has quickly become an essential part of the team, as well as an arduino interlock expert. I have no doubts that the experiment is in good hands.

Besides the people I worked with directly, there are many other Raizen lab members who provided me with good times and good physics when I needed them. Uzi Even worked with us on the supersonic beam side and greatly shorted my learning curve regarding beam physics. Gabe was a great source of knowledge regarding single photon cooling and he also was the driving force behind the first and only Raizen lab Salt Lick Contest. David is a fantastic physicist, with a meticulous approach

to problems, and I frequently went to him if there was a concept that confused me. Hrishi, Tongcang, and Kevin were all wonderful people to work with. Simon and Akarsh are sure to do great things with their bead experiment, and Jamie and Roderigo are quickly getting the focusing experiment going. Jianyong is rapidly learning the lithium experiment, and he is sure to do great things with it in the future. Melissa was our lab's high energy person, and I had many wonderful discussions with her while sharing an office. I am sure she will find success in her new role leading the company that has grown out of her research here. I thank all of you for your help and friendship.

I have received considerable support from other members of the UT physics community as well, and I would particularly like to thank the machine and cryoshop here. It is thanks to the skills of the machinists that some of our more complicated designs came together and worked as intended, and the shop was certainly there for me to fix things that had gone wrong. I would particularly like to thank Allan for discussing my designs with me and going over the best way to approach a part. Also deserving of special mention are Lanny and Ed, who provided me with excellent support from the cryo and vacuum shop. Finally, I would like to thank Jack, who taught me everything I know about machining. When we needed a part RIGHT NOW, Jack was there to help me figure out how to make it. My heartfelt thanks to everyone on the third floor. I would also like to thank Elena, Mary Beth, and Olga on the 14th floor. Their tireless efforts make the lab run much smoother and facilitate much of what we do downstairs.

In addition to my fellow lab members, I would like to thank my parents for providing me with the support and love to get where I am now. Visits home frequently provided a much needed break, and life here was certainly less stressful knowing that they were there to support me. I would also like to thank my new family for all of

their support and for welcoming me into their lives.

Finally, I would like to acknowledge the University of Texas at Austin for their support in the form of the William S. Livingston Graduate Fellowship.

Adam Libson

General Methods of Controlling Atomic Motion: Experiments with Supersonic Beams as a Source of Cold Atoms

Adam Alexander Libson, Ph.D.
The University of Texas at Austin, 2012

Supervisor: Mark G. Raizen

This dissertation discusses several recently developed experimental techniques for controlling the motion of neutral atoms. While laser cooling and evaporative cooling have been extremely successful and have been in widespread use for many years, these techniques are only applicable to a few atomic species. Supersonic beams provide a general method of producing cold atoms in the co-moving frame, but their speeds are typically several hundreds of meters per second in the lab frame. Methods to slow and control atoms cooled by supersonic expansion are detailed.

A method for controlling the velocity of a cold beam of ground state helium using specular reflection from single crystal surfaces is demonstrated. The velocity of the beam is shown to be continuously tunable, and beam velocities as slow as 265 m/s are created from an initial beam speed of 511 m/s.

Magnetism is a nearly universal atomic phenomenon, making magnetic control of atomic motion a very general technique. Magnetic stopping of supersonic beams of metastable neon and molecular oxygen is demonstrated using a series of pulsed electromagnetic coils. Neon is slowed from 446 m/s to 56 m/s, and oxygen is slowed from 389 m/s to 83 m/s, removing over 95% of the kinetic energy. The experimental technique is described in detail, and the theory and principle are discussed. An

experiment for slowing and trapping of atomic hydrogen isotopes at around 100 mK using a room temperature apparatus is described. A method for further cooling of magnetically trapped hydrogen ensembles, single-photon cooling, is proposed.

Table of Contents

Acknowledgments v

Abstract ix

List of Tables xiv

List of Figures xv

Chapter 1. Introduction 1

1.1 Cooling an Effusive Beam . . . 2

1.1.1 Laser Cooling . . . 2

1.1.2 Buffer-Gas Cooling . . . 3

1.2 Methods of Controlling Supersonic Beams . . . 4

1.3 Dissertation Outline . . . 5

Chapter 2. An Overview of Supersonic Beams 7

2.1 Thermodynamics of Ideal Gases . . . 8

2.1.1 Steady State Gas Flow in 1D . . . 9

2.1.2 Mach Number and Supersonic Beams . . . 11

2.1.3 Supersonic Beam Velocities . . . 12

2.2 The Supersonic Beam as a Bright General Source of Cold Atoms and Molecules . . . 13

2.3 The Even-Lavie Nozzle for Generating Pulsed Beams . . . 16

Chapter 3. Slowing Supersonic Beams via Specular Reflection: The Atomic Paddle 20

3.1 Using Helium for Atom Optics Experiments . . . 22

3.2 Reflecting Atoms from Surfaces: Atomic Mirrors . . . 23

3.2.1 Crystal Choice: Si(111) . . . 25

3.2.2 Crystal Preparation and Wet Etching . . . 26

3.2.3 Crystal Characterization . . . 28

3.3 The Experimental Apparatus . . . 28

3.3.1 The Vacuum System . . . 30

3.3.2 The Rotor . . . 34

3.3.2.1 Initial Design . . . 34
3.3.2.2 Vibration and Balancing . . . 39
3.3.2.3 Improved Spindle Design . . . 42
3.3.3 Detection . . . 44
3.3.4 Control Systems . . . 44
3.4 Data and Results . . . 46
3.4.1 Static Reflection . . . 47
3.4.2 Beam Slowing . . . 48
3.4.3 Crystal Lifetime . . . 52
3.5 Potential Applications . . . 54

Chapter 4. The Atomic and Molecular Coilgun: Using Pulsed Magnetic Fields to Slow Supersonic Beams 56
4.1 Atoms in External Magnetic Fields . . . 57
4.1.1 The Zeeman Effect . . . 58
4.1.2 The Paschen-Back Effect . . . 61
4.1.3 Adiabatic Following and Spin Flips . . . 62
4.2 Magnetic Fields from Electromagnetic Coils . . . 63
4.2.1 Switching of Electromagnetic Coils . . . 65
4.3 Principle of Operation . . . 66
4.3.1 Phase Stability . . . 69
4.3.2 Electromagnetic Coils as a Waveguide . . . 72
4.4 Proof-of-Principle Experiment: Slowing Metastable Neon . . . 74
4.4.1 Beam Creation . . . 75
4.4.2 18 Stage Coil Design . . . 77
4.4.3 18 Stage Electronics and Coil Switching . . . 78
4.4.4 18 Stage Apparatus and Vacuum Chamber . . . 83
4.4.5 18 Stage Data and Results . . . 87
4.5 The Atomic Coilgun: Stopping Metastable Neon . . . 89
4.5.1 64 Stage Coil Design . . . 90
4.5.2 64 Stage Electronics and Coil Switching . . . 94
4.5.3 64 Stage Apparatus and Vacuum Chamber . . . 100
4.5.4 64 Stage Data and Results . . . 105
4.6 The Molecular Coilgun: Stopping Molecular Oxygen . . . 109
4.6.1 Oxygen's Magnetic Levels . . . 110
4.6.2 Oxygen Beam Creation and Detection . . . 111
4.6.3 Oxygen Data and Results . . . 113

Chapter 5. Towards Trapping and Cooling of Atomic Hydrogen Isotopes 116
5.1 Hydrogen Motivation and Structure 117
5.2 The Hydrogen Apparatus 119
5.2.1 Slowing Coils and Electronics 119
5.2.1.1 Slowing Coil Geometry 120
5.2.1.2 Slowing Coil Electronics and Switching 121
5.2.2 Trapping Coils and Electronics 125
5.2.2.1 Principle of Operation of the Trap 126
5.2.2.2 Trapping Coil Geometry and Fields 128
5.2.2.3 Trapping Electronics and Switching 130
5.2.3 Beam Creation and Detection 135
5.2.4 Vacuum Chambers 137
5.3 Simulations of Hydrogen Trapping 141
5.3.1 Slowing and Trapping 142
5.3.2 Ejection Simulations 145
5.4 Future Directions 147
5.4.1 Laser Detection and Spectroscopy of Hydrogen Isotopes 148
5.4.1.1 Two-Photon Excitation of Hydrogen 148
5.4.1.2 The Hydrogen Laser 152
5.4.1.3 Spectroscopic Goals 155
5.4.2 An Adiabatic Coilgun 156
5.4.3 Single-Photon Cooling of Hydrogen 157

Appendix 164

Appendix A. Reflection From a Spinning Rotor 165

Bibliography 168

Vita 189

List of Tables

4.1 Peak Magnetic Fields Measured for Different Lengths of TGG Crystal 100
4.2 Neon Slowing Data Summary 106
4.3 Oxygen Slowing Data Summary 114

List of Figures

2.1 Comparison of effusive and supersonic velocity distributions 14
2.2 Section view of the Even-Lavie supersonic nozzle 17

3.1 Calculated helium elastic scattering fraction from Si(111) 24
3.2 AFM image of a passivated Si(111) crystal 29
3.3 CAD image of nozzle chamber 30
3.4 CAD image of rotor chamber 31
3.5 CAD image of the detection chamber 32
3.6 CAD image of entire apparatus 33
3.7 CAD image of the rotor 36
3.8 Finite element calculation of the forces on a spinning rotor 36
3.9 CAD image of the rotor with spindle, feedthrough, and motor 37
3.10 CAD image of the large crystal holder 38
3.11 Measurements of rotor vibration at different rotational speeds 40
3.12 Finite element calculation of rotor vibrational modes 41
3.13 CAD image of the redesigned spindle 43
3.14 Rotor triggering and reflection positions 45
3.15 Time-of-flight profiles for direct and reflected beams 47
3.16 Time-of-flight profiles for different rotor speeds 49
3.17 Plot of slowed beam velocity vs. crystal velocity 51
3.18 Plot of reflected intensity vs time 53

4.1 Angular momentum coupling in low field 59
4.2 Example Zeeman shift for a $^3P_2$ state 60
4.3 Angular momentum coupling in high field 61
4.4 Magnetic Field from a Single Loop Coil 64
4.5 Principle of operation of the coilgun 68
4.6 Phase Stability From Coil Timing 71
4.7 Magnetic Field Magnitudes as a Function of Radial Position 73
4.8 Images of the Discharge Apparatus and Discharge 76
4.9 CAD Overview of Proof-of-Principle Experiment Coils 78
4.10 Finite Element Calculation of Magnetic Field in the 18 Stage Slowing Coils 79

4.11 Schematic of the 18 Stage Coil Driver Circuit 80
4.12 Current in 18 Stage Slowing Coil 82
4.13 Temporal Profile of the Magnetic Field in the 18 Stage Coil 83
4.14 CAD Image and Photo of 18 Stage Coilgun 84
4.15 CAD Image of the 18 Stage Coilgun in the Rotor Chamber 85
4.16 Schematic of the MCP Bias Circuit 86
4.17 Photos of the 18 Stage Coilgun and Chamber 86
4.18 Results of Phase Variation in 18 Stage Coilgun 87
4.19 Results of Current Variation in 18 Stage Coilgun 88
4.20 Exploded CAD View of 64 Stage Coilgun Coil 90
4.21 64 Stage Slower Coil Assembly Process 92
4.22 Finite Element Calculation of the 64 Stage Coilgun Magnetic Field . 93
4.23 Schematic of the 64 Stage Coil Driver Circuit 94
4.24 Photographs of the 64 Stage Coil Driver Board 95
4.25 Schematic of Faraday Rotation Setup 98
4.26 Temporal Profile of the Faraday Rotation in a 64 Stage Coilgun Coil 99
4.27 64 Stage Aluminum Coil Support Structure 101
4.28 Cut-Away CAD Image of the 64 Stage Coilgun and Vacuum Chamber 102
4.29 CAD Image of the 64 Stage Coilgun Vacuum System 103
4.30 Neon Slowing Data from 64 Stage Coilgun 105
4.31 Slowing Flux with Variable vs. Constant Phase 108
4.32 Schematic of LN2 Boil Off Valve Cooling System 112
4.33 Oxygen Slowing Data from 64 Stage Coilgun 114

5.1 Hyperfine Structure of Ground State Hydrogen 118
5.2 Numerical Calculation of the Magnetic Field in the Hydrogen Coilgun Coil 120
5.3 Photograph of the Hydrogen Coilgun Coil on the Chamber 121
5.4 Schematic of the Coilgun Coil Switching Circuit 122
5.5 Photograph of the Hydrogen Coilgun Coil Driver Circuit 123
5.6 Faraday Rotation Measurement of the Magnetic Field in the Hydrogen Coils 124
5.7 Switching Sequence of the Trapping Coils with Numerical Calculations of the Fields 127
5.8 Photograph of the Trapping Coils Wound around the Chamber 128
5.9 Measured On-Axis Magnetic Fields in the Trapping Coils 129
5.10 Schematic of the Trapping Coil Switching Circuit 131
5.11 Photograph of the Hydrogen Trap Coil Driving Circuit 132

5.12 Faraday Rotation Signal During the Switching of the Front Trapping Coil 134
5.13 Ratio of D to D_2 with and without the discharge 136
5.14 CAD image of the Hydrogen Valve Chamber and Differential Pumping Chamber 138
5.15 CAD image of the hydrogen coilgun, trap, and respective vacuum chambers 139
5.16 Photograph of the Assembled Slowing and Trapping Chambers with Coils 139
5.17 CAD Image of the Detection Chamber with a Quadrupole Mockup . 140
5.18 CAD Image of the Complete Hydrogen Apparatus 141
5.19 Simulated Time-of-Flight and Velocity Distribution Profile 143
5.20 Simulated Phase Space Distributions of Trapped Hydrogen 144
5.21 Simulated Time of Flight Profile for Atoms Ejected from the Trap into the Ardara Ionizer 146
5.22 Simplified Hydrogen Energy Level Diagram 149
5.23 Schematic of the Original Hydrogen Laser 153
5.24 Schematic of the Tellurium Spectroscopy Setup 154
5.25 Cooling and Compressing Phase Space with a One-Way Wall 158
5.26 Scheme for RF-Dressed Single-Photon Cooling of Hydrogen 162

A.1 Geometry Used for Rotor Calculations 166

Chapter 1

Introduction

General control of atomic motion is a long standing goal of atomic physics, which is motivated by a desire for increased precision in spectroscopy, tests of fundamental symmetries, studies of entanglement, many-body physics, and even solid state physics. The ultimate goal is the ability to cool and trap any atom or molecule to quantum degeneracy. Laser cooling, for example, has been extremely successful, but has been limited to a few atomic species. While degenerate gases of alkali atoms has provided a rich field of study, the lack of generality have limited studies of other atoms, which may be hiding interesting physics.

The starting point for many atomic physics experiments is an atomic beam. A simple effusive beam can be created from almost any species, simply by allowing gas to escape from an aperture into vacuum. This may be done with a gas reservoir, an oven, or any number of other configurations, making the effusive beam an extremely general technique. While they are simple to create, effusive beams do have several drawbacks. Most notably, the temperature of the atoms in an effusive beam is the same as the temperature in the gas reservoir. This disadvantage is particularly problematic when an experiment requires cold atoms for greater precision. Effusive beams also generally have somewhat low brightness, which can limit the number of atoms available for an experiment.

To address these problems, there are two approaches one can take. One solution is to start with a hot beam of atoms and to subsequently cool it, thereby

producing a cold sample. Alternatively, by changing the conditions under which the beam is created, a supersonic beam can be produced. Supersonic beams are significantly colder and more directional than effusive beams, and temperatures as cold as a few tens of millikelvin have been achieved in the co-moving frame. Supersonic beams are also very general sources of atoms, as most species can be entrained into the beam. Since the supersonic beam provides a significant degree of cooling, the development of methods to control the atoms in the beam is an exciting field which holds the promise of fulfilling the goal of general control of atomic motion.

1.1 Cooling an Effusive Beam

Starting with a hot beam and using interactions with external fields or collisions with other atoms is the typical method of producing cold samples of atoms. The methods described here are some of the techniques used for cooling and trapping the atomic beams created by effusive sources.

1.1.1 Laser Cooling

The use of laser induced transitions to cool atomic samples has developed into a mature and widely used technique [1]. This cooling method uses momentum transfer from photons to atoms to create a dissipative force that is cools the sample. While this method has found great success, and is used in many labs in the form of the Zeeman slower [2], and the Magneto-Optical Trap (MOT) [3], there are fundamental limits to its applicability. Because the momentum of a photon is typically less than that of an atom, many scattering events are required to cool each atom, usually on the order of a few thousand. Each scattering event consists of an atom absorbing a photon and transitioning from a lower state into an excited state, and then undergoing spontaneous decay back to the original lower state.

An atom in an excited state may, in most cases, decay to any one of a number of lower level states. For laser cooling to work properly, the atom must decay back to the original lower level state, where the laser wavelength is near resonant. If it does not, the atom can fall into a dark state, where the laser does not interact with the atom, and be lost. The branching ratio governs the rates at which atoms in a particular excited state will decay into the allowed lower level states. For most atoms, the internal structure is complex enough that the branching ratio has non-zero components in several lower level states that are not resonant with the cooling laser, and most of the atoms will be lost after several thousand scattering events. This limits the number of atoms which can be laser cooled. In practice, only the alkali atoms, and a few others are amenable to laser cooling due to this limitation. Another cooling method is required for those atoms that are not laser coolable due to a more complex internal structure, which has motivated the work found in this dissertation.

1.1.2 Buffer-Gas Cooling

Instead of using light to cool atoms, another approach is to use collisions with a cold background gas. This is known as Buffer Gas Cooling [4–7]. The entire experiment takes place inside a dilution refrigerator, with helium typically used as the buffer gas, to reach the cold temperatures desired. Using this method, a gas jet is allowed to expand into the cold vacuum chamber with the helium buffer gas. As the injected gas collides with the helium, it thermalizes with the background gas. This is typically done in the region of a magnetic trapping potential. When the temperature of the injected gas is below the trap depth, it will be trapped if it is in the correct state. The helium is then pumped away, leaving a cold sample in the trap.

While this is a general method, there are several drawbacks that have kept buffer gas cooling from seeing widespread use. Needing to have the experiment take

place inside a dilution refrigerator adds significant cost and complexity to the experiment. Optical access to the trapped sample is also limited because of the cryogenic environment. This lack of access limits the experiments which can be performed, as well as the ability to effectively observe the sample. Thus, a general method which is able to trap the sample in a room temperature chamber with good optical access would decrease complexity, while enhancing the experimental utility of the apparatus.

1.2 Methods of Controlling Supersonic Beams

Because supersonic beams make excellent sources of cold atoms, numerous groups have developed methods to control atoms and molecules cooled by supersonic beams. The challenge is that while the atoms or molecules have been cooled by the expansion, they are still moving fast in the lab frame. Even in this form, supersonic beams are a staple of physical chemistry [8–10]. For many experiments though, the high speed of the beams is a serious detriment, which has led to many efforts to control the motion of supersonic beams.

One approach to controlling the beam is to employ mechanical methods. The velocity of the supersonic flow is relative to the nozzle, not the vacuum chamber itself, suggesting that it should be possible to produce slower beams by moving the nozzle opposite to the expansion. This is the approach taken by Gupta and Hershbach [11], who mounted their nozzle at the end of a rapidly rotating arm. Further mechanical control methods have included focusing by reflecting the beams from bent crystals [12], and stopping via collisions between atoms in crossed beams [13].

Another approach is to use external fields to control the motion of atoms or molecules in the beam. Pulsed laser fields have been used to stop atoms in a beam [14]. Pulsed electric fields are used to slow and trap molecules with permanent electric dipole moments [15–17]. The pulsed electric field approach is also used to slow and

even trap Rydberg atoms [18–20].

Instead of using external optical or electric fields, external magnetic fields can be applied to the control of the atoms or molecules in the beam. Since magnetism is nearly universal for atoms in the ground state or a low-lying metastable state, this a very general approach to controlling atomic motion. Several groups have pursued this method, and this dissertation presents one approach to using magnetic fields to control atoms and molecules in supersonic beams. A similar approach has been pursued independently and in parallel by the group of F. Merkt [21–26].

1.3 Dissertation Outline

Chapter 2 introduces the physics of the supersonic beam. The thermodynamic properties of a one dimensional adiabatic expansion are derived, and the properties and advantages of the beam are discussed. The Even-Lavie nozzle, which is the device used to produce supersonic beams in these experiments, is introduced.

Chapter 3 of this dissertation discusses the Atomic Paddle experiment. In this experiment, a supersonic beam of helium is slowed by specular reflection from a moving crystal. While ground state helium is very difficult to control due to its low polarizability and lack of a magnetic moment, it is known to reflect well from atomically flat single crystal surfaces. By mounting an externally prepared crystal on the tip of a spinning rotor that is used to move the crystal along the beam propagation direction, the velocity of the reflected beam is reduced. Beam temperatures are shown to be conserved by this process. The experimental apparatus, vacuum chambers, and crystal preparation and characterization processes are described. The data presented demonstrates experimental control of the velocity of a helium beam.

Chapter 4 discusses the Atomic Coilgun experiment. Supersonic beams of

atoms or molecules are slowed using pulsed magnetic fields. Since most atoms and some molecules are magnetic dipoles in their ground state or an easily accessible metastable excited state, this method of producing cold slow samples is very general. Pulsed magnetic fields are created using high current electromagnetic coils, which the atoms or molecules travel through. The magnetic moment of the atom or molecule means that some are repelled by magnetic fields, leading to a loss of energy as the atom enters the coil and slowing them down. Switching off the coil before the atoms can exit means they do not regain this energy. Multiple coils in sequence are used to bring the beam to rest in the lab frame. The coilgun apparatus, including coil design and characterization, electronics, and beam preparation, is described in this chapter. The data presented shows control of the velocity of beams of metastable neon and molecular oxygen.

Chapter 5 discusses the ongoing project to trap atomic hydrogen. Hydrogen is particularly interesting since it is the simplest atom and has served as a Rosetta Stone for atomic physics. Since the magnetic moment to mass ratio of hydrogen is the highest among all atoms, it is ideally suited to the coilgun method. The hydrogen specific coilgun and trap are described and characterized. Simulations of the hydrogen trapping process are presented. Prospects for two-photon spectroscopy of trapped hydrogen are discussed, and the laser to excite the two-photon transition is described. Future improvements to the coilgun apparatus are presented. Finally, a proposed adaptation of single photon cooling to atomic hydrogen using RF-dressed states and the two-photon transition to further cool a magnetically trapped sample is discussed.

Chapter 2

An Overview of Supersonic Beams

The atomic or molecular beam is the basis of many experiments, and the starting point of all of the experiments described in this dissertation. As such, a basic description of the properties of ideal gas flows is provided. An atomic or molecular beam is generally created by allowing gas to flow through an aperture into another chamber held at lower pressure. The flow of gas generally falls into one of two regimes, creating either an effusive beam, or a supersonic beam.

A effusive beam can be created for almost any species, simply by allowing gas to escape from an aperture into vacuum. This may be done with a gas reservoir, an oven, or any number of other configurations, making the effusive beam an extremely general technique. While they are simple to create, effusive beams do have several drawbacks. Most notably, the temperature of the atoms in an effusive beam is the same as the temperature in the gas reservoir. This disadvantage is particularly problematic when an experiment requires cold atoms for greater precision. Effusive beams also generally have low beam brightnesses, which limits the number of atoms available for an experiment. The supersonic beam is somewhat more difficult to produce, but it has several advantages over the effusive beam. In particular, the supersonic beam is more directional, and is colder than an effusive beam, though it does have a higher mean velocity. The brightness of supersonic beams can also greatly exceed those produced by effusive sources.

2.1 Thermodynamics of Ideal Gases

For an effusive beam, the gas is in the free molecular flow regime, where the gas reservoir is kept at a relatively low pressure and the gas escapes through an aperture $d < \lambda$, where d is the diameter of the aperture, and λ is the mean free path in the gas. Starting from the ideal gas law,

$$PV = nk_BT, \tag{2.1}$$

where P is the pressure of the gas, V is the volume of the reservoir, n is the number of gas molecules, T is the temperature of the reservoir, and k_B is the Boltzmann constant, it is relatively simple to show that the mean free path is

$$\begin{aligned} \lambda &= \frac{1}{\sqrt{2}\pi d_{mol}^2 n_V} \\ &= \frac{k_BT}{\sqrt{2}\pi d_{mol}^2 P}, \end{aligned} \tag{2.2}$$

where d_{mol} is the diameter of a gas particle, and n_V is the number density of the gas. For an effusive beam, $d < \lambda$ so the last collision of an atom leaving the reservoir is far from the aperture, and the flow of gas leaving the reservoir does not have a significant effect on the beam. Thus, the expected velocity distribution is Maxwellian

$$p(v) = \frac{4}{\sqrt{\pi}} \left(\frac{m}{2k_BT} \right)^{\frac{3}{2}} v^2 e^{\frac{-mv^2}{2k_BT}}, \tag{2.3}$$

where m is the mass of a gas molecule, and v is its velocity. Since the flow does not have a significant effect on the atoms, the beam will be distributed over a large solid angle, typically a full 2π str. To give some idea of the scales involved, at 300 K and and a pressure of $1 * 10^{-3}$ atm, the mean free path of a helium atom is ≈ 1 mm. For a reservoir under these conditions, the beam produced will be effusive for an aperture smaller than 1 mm.

As the pressure in the gas reservoir is increased, the mean free path decreases accordingly. When the pressure is high enough, the mean free path will decrease to the point that $d > \lambda$. When this condition is met, the gas enters the continuum regime near the nozzle and can be treated as a compressible fluid. The principle advantages of this regime are the greatly decreased temperature of the atoms in the beam, and increased brightness due to higher reservoir pressures, larger nozzle diameter, and increased directionality of the beam.

2.1.1 Steady State Gas Flow in 1D

To examine what can be expected in the $d > \lambda$ regime, it is instructive to examine the flow of an ideal gas through passages with changing cross-sectional area (this is presented in detail in [8–10] and the discussion here follows the presentation in [8]). To do this effectively, some of the basic thermodynamic properties of ideal gases are needed. To start with, the ideal gas law (equation 2.1) can be expressed as

$$P = \frac{\rho k_B T}{m}, \tag{2.4}$$

where ρ is the density of the gas. Letting e_{int} be the internal energy of the gas per unit mass, the heat transfered to an unchanging quantity of an ideal gas, $\mathrm{d}q$, at constant pressure is

$$\mathrm{d}q = \mathrm{d}e_{int} + P\mathrm{d}\left(\frac{1}{\rho}\right), \tag{2.5}$$

and the enthalpy per unit mass is

$$h = e_{int} + \frac{P}{\rho} = e_{int} + \frac{k_B T}{m}. \tag{2.6}$$

Using equation 2.5 and equation 2.6, along with the fact that the internal energy of an ideal gas depends only on the temperature, produces

$$\mathrm{d}q = \mathrm{d}h - \frac{1}{\rho}\mathrm{d}P. \tag{2.7}$$

Since there will be no heat exchange between the gas and the environment during the expansion, equations 2.5 and 2.7 become

$$0 = \mathrm{d}e_{int} + P\mathrm{d}\left(\frac{1}{\rho}\right), \tag{2.8}$$

and

$$0 = \mathrm{d}h - \frac{1}{\rho}\mathrm{d}P. \tag{2.9}$$

The specific heats at constant pressure, c_P, and at constant volume, c_V, are given by $c_P = \left(\frac{\mathrm{d}h}{\mathrm{d}T}\right)_P$ and $c_V = \left(\frac{\mathrm{d}e_{int}}{\mathrm{d}T}\right)_V$, and their ratio, $\gamma = \frac{c_P}{c_V}$, is a constant for an ideal gas that depends only on the number of degrees of freedom of a gas molecule. Using these relations, equations 2.8 and 2.9, can be re-expressed and combined as [8]

$$\gamma\frac{\mathrm{d}\rho}{\rho} - \frac{\mathrm{d}P}{P} = 0. \tag{2.10}$$

This last relation is of particular use in examining the gas flow.

There are several conservation laws that can be used to simplify the problem of gas flow through a nozzle. Clearly, conservation of mass applies in the flow, gas particles are not being spontaneously created or annihilated. This is expressed as

$$\nabla\left(\rho\vec{w}\right) = 0 \tag{2.11}$$

or simplified for 1D flow as

$$\frac{\mathrm{d}w}{w} + \frac{\mathrm{d}\rho}{\rho} + \frac{\mathrm{d}A}{A} = 0 \tag{2.12}$$

where w is the velocity of the flow, ρ is the density of the gas, and A is the cross-sectional area of the flow. Momentum is also conserved in the flow, which is essentially an expression of Newton's second law in the gas. Assuming no external forces on the gas

$$\frac{\mathrm{D}\vec{w}}{\mathrm{D}t} = \frac{-\nabla P}{\rho} \tag{2.13}$$

which for steady state flow in 1D can be reduced to

$$w\mathrm{d}w = \frac{-\mathrm{d}P}{\rho}. \tag{2.14}$$

Lastly, energy is conserved in the flow. This means that

$$\frac{\mathrm{D}}{\mathrm{D}t}\left(h + \frac{w^2}{2}\right) = 0 \tag{2.15}$$

which can be simplified to

$$w\mathrm{d}w + \mathrm{d}h = 0. \tag{2.16}$$

2.1.2 Mach Number and Supersonic Beams

A final concept that must be introduced before addressing the flow problem is information's propagation in the gas. Information is transmitted in the flow as a pressure disturbance or pressure wave, and the speed of propagation is simply the speed of sound. In an ideal gas, the speed of sound is

$$v_s = \sqrt{\frac{\gamma P}{\rho}}. \tag{2.17}$$

This equation assumes that there is no collective motion in the gas. This can be addressed by simply changing coordinate systems into one that is moving along with the flow. Thus information will travel in the static lab frame at a velocity $v_s \pm w$, where the sign depends on whether the pressure wave is traveling with the flow or against the flow. If the flow is faster than the speed of sound, then information cannot travel backwards in the flow, and the flow is supersonic. The ratio of the flow velocity to the speed of sound is called the Mach number M,

$$M = \frac{w}{v_s}. \tag{2.18}$$

Since the speed of sound depends on the mean molecular speed, the Mach number can be thought of as a measure of the collective motion of the beam relative to the thermal motion of the individual atoms in the ensemble.

Using equation 2.17 to modify equation 2.10 gives $\frac{\mathrm{d}P}{\rho} = v_s^2 \frac{\mathrm{d}\rho}{\rho}$ which when substituted into the momentum conservation law stated in equation 2.14 gives

$$\frac{\mathrm{d}\rho}{\rho} = -\frac{w}{v_s^2}\mathrm{d}w. \tag{2.19}$$

This relationship can be substituted into the mass conservation law of equation 2.12 along with equation 2.17 giving

$$\frac{\mathrm{d}A}{A} + \frac{\mathrm{d}w}{w}\left(1 - M^2\right) = 0. \tag{2.20}$$

Much of the behavior of a 1D flow can be determined from this equation. A cold beam, or one with a high degree of directed motion as opposed to internal thermal motion, will have a high Mach number. Starting in the reservoir, the Mach number is clearly much less than one, and so $(1 - M^2) > 0$. This means that if the flow is subsonic, the flow velocity increases with decreasing cross-sectional area, and vice versa. Alternatively, if the flow is supersonic, then $(1 - M^2) < 0$ and the flow will increase in velocity with increasing area. Finally, if there is a sufficient pressure gradient to accelerate the flow to supersonic velocities, at the location where $M = 1$, $(1 - M^2) = 0$ and so $\frac{\mathrm{d}A}{A} = 0$, implying that the cross sectional area is at a minimum. This means that to create a supersonic beam, a nozzle must converge to a minimum before diverging again.

2.1.3 Supersonic Beam Velocities

Having determined the qualitative character of the supersonic flow, it is helpful to turn to the specifics of beam velocity and temperature. The energy conservation law, equation 2.16, can be integrated to find the flow velocity at various locations,

$$\frac{w_1^2 - w_0^2}{2} = h_1 - h_0 = \int_{T_1}^{T_0} c_P \mathrm{d}T = c_P T_0 \left(1 - \frac{T_1}{T_0}\right), \tag{2.21}$$

assuming that c_P is a constant. Selecting the initial location to be well inside the gas reservoir, where the flow velocity is negligible and the gas is at the reservoir

temperature and pressure, permits the velocity to be related to the temperature at any point along the flow. Using the relations $c_P = \frac{\gamma}{\gamma-1}\frac{k_B}{m}$, and $\frac{P_a}{P_b} = \left(\frac{T_a}{T_b}\right)^{\frac{\gamma}{\gamma-1}}$ gives a flow velocity of

$$w = \sqrt{2\frac{\gamma}{\gamma-1}\frac{k_B T_{res}}{m}\left(1-\frac{T}{T_{res}}\right)} = \sqrt{2\frac{\gamma}{\gamma-1}\frac{k_B T_{res}}{m}\left(1-\left(\frac{P}{P_{res}}\right)^{\frac{\gamma-1}{\gamma}}\right)}, \quad (2.22)$$

where T_{res} and P_{res} are the temperature and pressure in the reservoir respectively. Assuming that the gas is expanding into vacuum, where there is negligible pressure, the maximum flow velocity in a supersonic expansion is

$$w = \sqrt{2\frac{\gamma}{\gamma-1}\frac{k_B T_{res}}{m}}. \quad (2.23)$$

This maximal velocity is slightly greater than the average velocity of a gas molecule in the reservoir.

2.2 The Supersonic Beam as a Bright General Source of Cold Atoms and Molecules

While this discussion has illustrated many of the properties of adiabatic expansion in 1D, the reality is that physical flows are expanding in three dimensions. Much of the discussion above still holds, however the details of the expansions change. A detailed discussion of adiabatic expansion in 3D may be found in [8]. Since the situation in 3D is significantly more complicated, numerical simulations are often used to model the beam [27]. One change that must be addressed is that as the gas expands the regime returns to free molecular flow and collisions between gas molecules cease, to good approximation. A simple, yet useful, model divides the expansion region in two. Near the nozzle, the gas is assumed to be a continuous medium and still collisional, while outside of this region, collisions are assumed to no longer take place. This is known as the sudden freeze model, and the surface of this division is

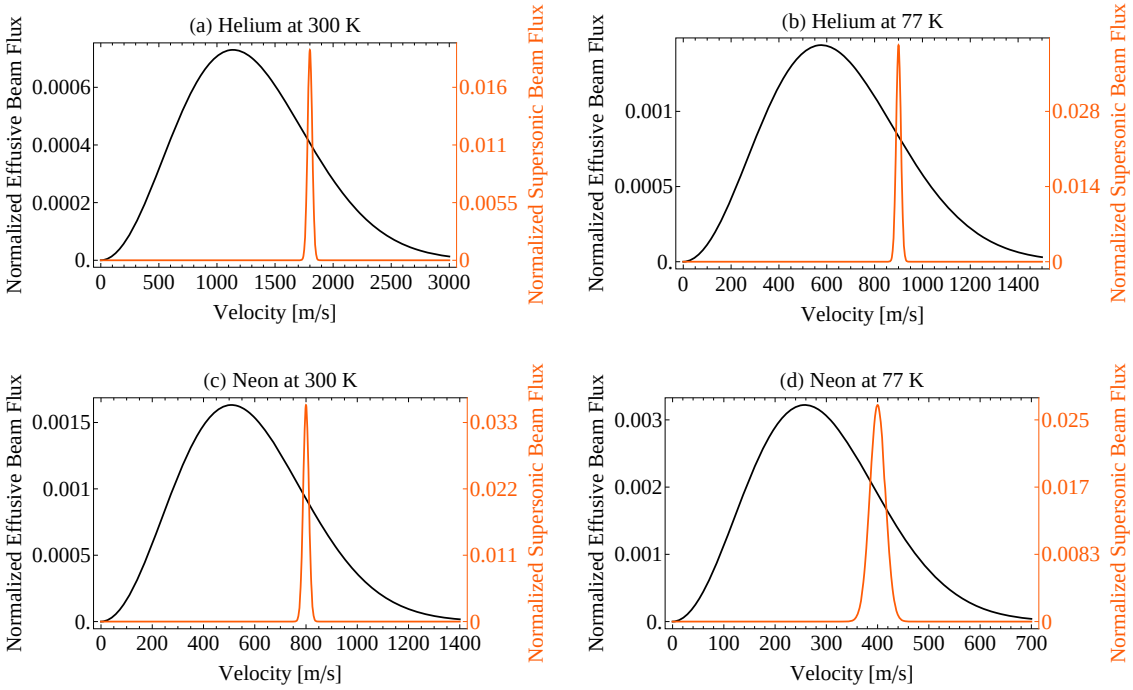

Figure 2.1: The velocity distribution for supersonic and effusive beams are compared for Helium and Neon beams at reservoir temperatures of 300 K and 77 K. The distributions are normalized to aid the comparison of flux at a particular velocity. The temperatures of the supersonic beams are taken from experimental data from the UT supersonic valve and experiments performed in Uzi Even's lab at Tel Aviv University [28].

called the quitting surface. Since the beam is no longer collisional downstream of the quitting surface, finding the velocity distribution at the quitting surface is sufficient to characterize the beam. The velocity distribution can be modeled as an anisotropic Maxwellian

$$p(v) = \sqrt{\frac{m}{2\pi k_B T_\parallel}} \frac{m}{2\pi k_B T_\perp} e^{-\frac{m v_\perp^2}{2k_B T_\perp} - \frac{m\left(v_\parallel - w\right)^2}{2k_B T_\parallel}}, \tag{2.24}$$

where $T_\parallel$ and $T_\perp$ are the temperatures along the axis of the beam and perpendicular to the beam respectively, with the same notation used for $v_\parallel$ and $v_\perp$.

In practice, the perpendicular temperature manifests as a loss of flux, since the beam generally goes through several apertures before being used in an experiment.

The apertures set the transverse temperature, and hotter atoms are removed from the beam by the apertures, leaving only the colder atoms in the beam. For this reason, it is appropriate to concentrate on the parallel component of the velocity, giving the Gaussian velocity profile

$$p(v_\parallel) = \sqrt{\frac{m}{2\pi k_B T_\parallel}} e^{-\frac{m\left(v_\parallel - w\right)^2}{2k_B T_\parallel}} = \frac{1}{\sigma_\parallel \sqrt{2\pi}} e^{-\frac{\left(v_\parallel - w\right)^2}{2\sigma_\parallel^2}}, \tag{2.25}$$

where $\sigma_\parallel$ is the standard deviation in the velocity parallel to the beam axis. This velocity distribution is compared to the Maxwellian velocity distribution for a few relevant gas species and reservoir temperatures in figure 2.1. As can be seen in the figure, the width of the velocity profile is significantly narrower for the case of a supersonic beam. Expressed in terms of temperature, beams of a few tens to a few hundred milliKelvin are common. This is significantly colder than the reservoir temperatures, and the supersonic expansion has effectively cooled the beam.

The cooling achieved by supersonic expansion is extremely general and can be used to cool any desired species. In the above discussion on isotropic expansion, it is assumed that the gases were truly ideal gases. Supersonic beams are generally created with a noble carrier gas that generally behaves as an ideal gas at higher pressures and lower temperatures than many other gas species. Other species of interest are mixed in the reservoir before the expansion, or added just outside the nozzle during the expansion. If the species of interest is a gas at the reservoir temperature, it is typically seeded into the carrier gas in the reservoir. A species that is not a gas at the reservoir temperature is typically entrained into the beam at the exit of the nozzle. Laser ablation can be used to produce an atomic sample of a solid at the exit of a nozzle, or alternatively, an oven can produce an atomic beam that is then entrained in the supersonic beam. In either case, the added species is cooled by thermalizing with the dominant carrier gas and is carried along with the supersonic beam, providing a

general method for producing cold atoms.

Different carrier gases have their own advantages and disadvantages that must be weighed when selecting a species to use in an experiment. The mass term in equation 2.23 means that a heavier carrier gas will produce a slower beam. This can be a large advantage, since supersonic beam velocities can range from several hundred to over a thousand meters per second. There are, however, some disadvantages to using heavier atoms as the carrier gas. Due to the dynamics of the collisions in the beam, light atoms being carried by a heavier carrier gas will tend to be pushed to the edges of the beam, and so the centerline flux of the species of interest will be reduced. The other possible disadvantage of a heavier carrier gas is the possibility of a hotter beam due to clustering. The binding energy of a cluster is released as heating in the beam itself. The cluster fraction can be estimated using the Hagena parameter [29, 30], and clustering becomes increasingly likely for the heavier noble gases.

2.3 The Even-Lavie Nozzle for Generating Pulsed Beams

The experiments described by this dissertation all use a pulsed valve for generating supersonic beams [31–33]. The valves used are made by Uzi Even and Nachum Lavie at Tel Aviv University. There are several advantages to using a pulsed valve. The methods used to control the beams are pulsed and having a pulsed beam matches well with the control mechanisms. Additionally, the vacuum pumping requirements are significantly reduced by using a pulsed beam. Instead of using large diffusion pumps to maintain vacuum in the valve chamber, the experiments are able to use turbomolecular pumps that cannot contaminate the chamber with pump oils the way a diffusion pump can.

The Even-Lavie uses a nozzle with a 200 μm diameter and a trumpet shaped expansion region. This shape helps to remove hot spots in the beam and reduces

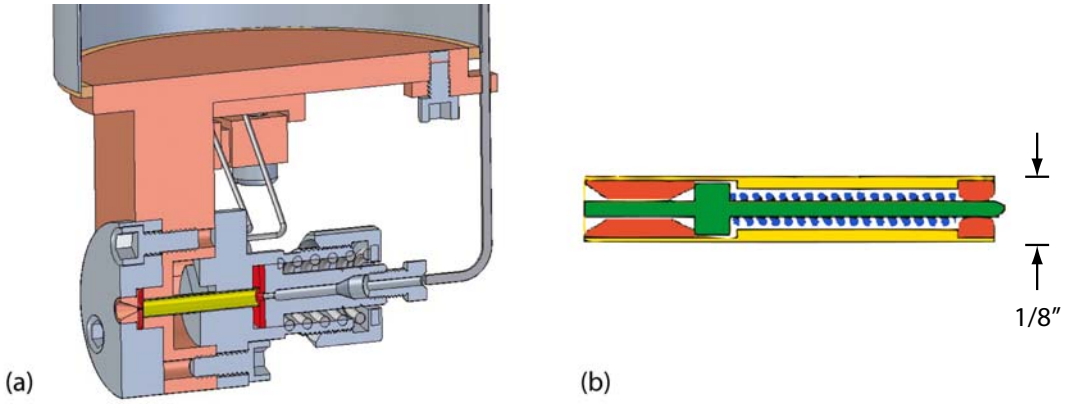

Figure 2.2: A CAD representation of the Even-Lavie supersonic nozzle. The complete nozzle apparatus is shown in (a) with a pool type cryostat which allows the nozzle to be cooled. Gas flows into the valve apparatus through the stainless steel tube on the right side of the image into a stainless steel pressure tube (yellow). The front and back of the tube are sealed using Dupont Kapton washers (red). The trumpet shaped nozzle where the gas expands and exits is on the left. An empty section can be seen surrounding the stainless steel pressure tube, and this is where the electromagnetic drive coil (not shown in this image) is located. The interior of the pressure tube is shown in (b). Gas enters the pressure tube (yellow) from the right and flows past the plunger (green), spring (blue), and guiding ceramic inserts (orange) to the nozzle exit (not depicted). The plunger forms a seal on the left side of the image with the leftmost Kapton washer from (a), and is held in place by the spring. Current in an electromagnetic coil produces a magnetic field that pulls the plunger to the right, allowing gas to expand adiabatically into the vacuum until the magnetic field is turned off and the plunger is pushed back into place by the spring. The motion of the plunger is guided by the two alumina pieces. Figure Courtesy Max Riedel.

the overall temperature, as well as producing a very directional beam with a FWHM opening angle of 16° [33]. The nozzle is unique in that it can produce pulses which are only 10 μs FWHM in length, and the valve can operate at repetition rates over 40 Hz. The valve can operate with backing pressures of 100 atm and at temperatures as low as 20 K. The short opening time is essential for these reservoir conditions, as the vacuum system would be overwhelmed by the resulting gas load otherwise.

The gas throughput of the nozzle can be estimated by

$$\Phi = \frac{n_0 w}{4} \tag{2.26}$$

where Φ is the throughput in molecules/m^2/s, n_0 is the number density of the molecules in the reservoir, and w is the speed of the flow. Consider neon at 77 K with a reservoir pressure of 5 atm as an example. At this pressure and temperature, the ideal gas law (equation 2.1) gives a number density of $5 \cdot 10^{26}$ molecules/m^3. The velocity of such a flow is approximately 400 m/s giving an estimated throughput of $5 \cdot 10^{28}$ molecules/m^2/s. Using a nozzle opening time of 10 μs and accounting for the 200 μm diameter of the valve gives a flux of $1.5 \cdot 10^{16}$ molecules/shot. Operating the valve under optimal conditions results in a peak brightness of $4 \cdot 10^{23}$ molecules/sr/s [33], which is greater than other sources [34–36] by at least an order of magnitude.

A CAD representation of the valve is shown in figure 2.2. The nozzle hangs from a pool type cryostat that allows the nozzle to be cooled. The 200 μm diameter trumpet shaped nozzle is depicted on the left side of figure 2.2 (a), with the gas entering from the stainless steel tube on the right. The gas flows through stainless steel pressure tube that is sealed at both ends by DuPont Kapton washers. The actual valve mechanism is depicted in figure 2.2 (b). The plunger forms a seal with the Kapton washer on the valve exit side of the high pressure tube and is held in place by a spring which sits inside the pressure tube. The plunger is moved by $\approx$150 μm via an electromagnetic coil that produces fields of around 2.5 T, allowing gas to adiabatically expand into the vacuum. The plunger is guided by two alumina pieces that keep it centered in the pressure tube. The plunger is returned to the sealing position by the spring when the coil is turned off. The entire plunger cycle takes 15 μs [37].

The Even-Lavie valve is a good match for the experiments described in the remainder of this dissertation. The methods of slowing and control described here

are pulsed methods, and cannot operate on a continuous basis. As such, the pulsed source provided by the Even-Lavie valve limits the amount of wasted gas released into the vacuum chamber. The uniquely short pulse lengths that the valve is capable of producing also provide a precise start time for the slowing experiments, permitting a simplification in the control electronics. Finally, since the overall goal of these experiments is to produce cold atoms that can be used for precision measurements, the high beam brightness produced by the Even-Lavie valve is an important asset, as this maximizes the number of atoms available.

Chapter 3

Slowing Supersonic Beams via Specular Reflection: The Atomic Paddle

As experimental control of atomic motion has improved, the wave nature of atomic beams has become an important factor in many experiments. Quantum mechanics states that like light, matter is both a particle and a wave. This has given birth to a new class of experiments which exploit the wave nature of atoms and manipulate beams of atoms in a particle analog to optical manipulation of beams of light. This growing field is known as atom optics, and much research has recently gone into developing a toolbox for manipulating atomic beams.

Most of the work in atom optics has used laser cooled atoms [38, 39], which provide several advantages. Since the atoms are laser cooled they are quite cold, and beams can be produced at low velocities. This accentuates the wave nature of the particles, as

$$\lambda = \frac{2\pi\hbar}{p} \tag{3.1}$$

where λ is the de Broglie wavelength of a particle, and p is the particle's momentum. Furthermore, laser cooled atoms have an accessible transition an experimentalist can use to both manipulate and detect the atoms in the beam. There are some drawbacks to using laser cooled atoms however, most notably their sensitivity to stray fields.

It is rather difficult to effectively control the atomic motion of ground state helium, due to the lack of accessible transitions with current laser technology. However, helium is known to reflect well from single crystal surfaces [40], which suggests

a method by which beams of helium may be controlled. In the experiment described here, specular reflection from receding single crystal surfaces is used to control the velocity of a supersonic beam of helium atoms. The crystal is mounted on the tip of a spinning rotor that provides the necessary crystal velocity to effectively slow the beam. The original concept for this experiment was first proposed by Doak [41], but the design studied had an expected slow beam flux that is significantly lower than in the experiment described here. The work presented in this chapter is also described in [33, 42, 43].

The slower itself is similar to the beam paddle that was developed for neutrons [44]. For helium atoms reflecting from a linearly moving single crystal surface, the velocity of the atoms after hitting the mirror will be

$$v_f = -v_i + 2v_m, \tag{3.2}$$

where v_f is the final velocity of the atoms, v_i is the initial velocity of the incoming beam, and v_m is the velocity of the mirror. The mirror is assumed to be perpendicular to the incoming beam. Since the typical velocity of a supersonic beam is 500 m/s, reducing the velocity of the beam by a factor of two requires a crystal velocity of 125 m/s. While it is possible to create a device to repeatably move a mirror linearly at these speeds (perhaps by mounting the mirror on a projectile and launching and catching it using electromagnetic coils), it is simpler to let rotary motion approximate linear motion and mount the crystal on the tip of a spinning rotor arm. This facilitates the high mirror velocities required to significantly slow the beam. The introduction of rotary motion complicates equation 3.2, and the full description of the velocity change due to an atom reflecting from the rotor is shown in appendix A.

3.1 Using Helium for Atom Optics Experiments

Most atom optics experiments have been done with laser cooled atoms. The advantage of this is a cold sample which can easily be manipulated. However, because of the large polarizability of most laser coolable atoms, they are sensitive to stray electric fields. Similarly, many laser coolable atoms are quite sensitive to stray magnetic fields. While their sensitivity is an advantage when controlling the atoms, stray fields will add systematic effects limiting the sensitivity of many measurements.

Though ground state noble gases are difficult to manipulate, doing so is not impossible and there are several advantages to using them for experiments in atom optics. One such advantage is that they are known to reflect well from single crystal surfaces, and even if they do not reflect from a surface, they are unlikely to stick to it. Because of this, an aperture will not clog when subjected to a beam of noble gases, even at high intensity. The other principle advantage is their low polarizabilty. Ground state noble gases are relatively insensitive to stray electric fields, and they are insensitive to magnetic fields to first order. This makes ground state noble gases, and ground state helium in particular, very attractive for certain precision measurements and applications such as atom interferometry.

A final advantage which is available when using ground state noble gases for atom optics experiments is the large fluxes available. Supersonic beams typically use noble gases as carriers, and a significant fraction of the flux of nearly any supersonic beam will consist of a ground state noble gas. Fluxes of up to $4 \cdot 10^{23}$ molecules/sr/s have been measured for pulsed beams as described in chapter 2, and for continuous beams, fluxes of $1.1 \cdot 10^{20}$ molecules/sr/s have been observed [45]. These brightnesses provide a large advantage and can reduce the uncertainty of many precision measurements, particularly for atom interferometry where the shot noise contributes significantly to the uncertainly. Helium is the atom most resistant to clustering, re-

sulting in the brightest supersonic beams and providing further motivation to find methods of controlling ground state helium beams.

In large part because of the advantages described above, there is a long history of using helium beams and helium scattering as an experimental probe. Stern and Estermann used helium diffraction from lithium fluoride crystals to demonstrate the wave nature of atoms and to experimentally confirm equation 3.1 [46]. Detailed examination of the diffraction of helium from surfaces has provided many insights into the details of surface structure and the atom-surface potential [47]. More recent work has used transmission diffraction gratings to study helium clustering in supersonic beams [35].

While ground state noble gases are more difficult to control than atoms which are laser coolable, several techniques have been developed which can be applied to the control of ground state helium. A standard optical toolbox includes mirrors, beamsplitters, and lenses. To effectively use ground state helium for atom optics, analogs of these tools are required. One such approach to creating these tools has been to use species independent methods such as transmission diffraction gratings [48–50] and transmission Fresnel lenses [51, 52]. Another approach is to make use of the high elastic scattering probability of helium from crystal surfaces. Flat surfaces create mirrors and the equivalent of a lens can be formed by bending the crystal, creating a curved mirror [12]. Finally, the periodic structure of a crystal surface can be used as a diffraction grating [53], creating a beamsplitter for the atoms. In this manner, the entire optical toolbox can be created using crystal surfaces.

3.2 Reflecting Atoms from Surfaces: Atomic Mirrors

Having a large fraction of the helium beam be elastically scattered from the crystal surface is important when manipulating a supersonic beam using reflection

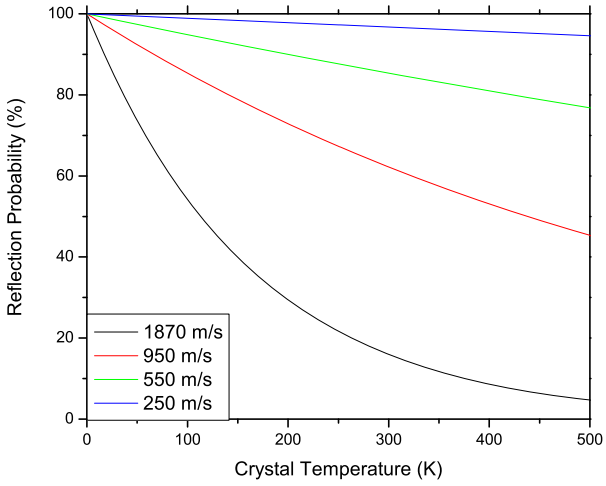

Figure 3.1: Calculated elastic scattering probability for helium incident on Si(111) for different beam velocities as a function of crystal temperature. These are calculated using the Debye-Waller factor described in equations 3.3 and 3.4. While the elastic fraction is well approximated by the Debye-Waller factor, atoms may still be diffracted and these numbers do not represent the fraction of specularly reflected atoms. Figure Courtesy Max Riedel.

from crystals. Assuming the atom surface potential has zero well depth (which means that there is no Beeby correction), the elastic scattering fraction from single crystalline surfaces is governed by the Debye-Waller factor [54],

$$\frac{I}{I_0} = e^{-2W} \tag{3.3}$$

where I is the elastically reflected intensity, I_0 is the incident intensity, and

$$W = \frac{12mE_{\perp}T_s}{Mk_B\theta_s^2}. \tag{3.4}$$

Here m is the mass of the incident atom, $E_{\perp}$ is the incident energy, T_s is the temperature of the crystal, M is the mass of the atom in the crystal, and θ_s is the Debye temperature of the crystal.

To put equations 3.3, and 3.4 into context, consider the following example. The experiment uses Si(111) as the reflecting surface, for reasons described in section 3.2.1, and the following calculation is based on this surface. The average mass of a single silicon atom in the crystal is $4.7 \cdot 10^{-26}$ kg and the crystal Debye temperature is 690K (at room temperature). For a supersonic beam of helium from a 77K nozzle, the beam velocity is ≈ 900 m/s and the mass of a helium atom is $6.7 \cdot 10^{-27}$ kg. From this, the Debye-Waller factor indicates that 65% of the beam will be elastically scattered. While this is a good approximation of the elastic scattering amplitude, it does not indicate the fraction of the beam which will be specularly reflected, as the beam may be distributed into many diffraction channels. Figure 3.1 shows what fraction of a helium beam will be elastically scattered as a function of crystal temperature, for several incident beam velocities.

3.2.1 Crystal Choice: Si(111)

The choice of crystal substrate is important when creating a good atomic mirror. Ideally, an atomic mirror will be atomically flat, since diffuse scattering will result from surfaces that are rough on a length scale above the wavelength of the incident beam. Also, since a mirror requires high specular reflection, the crystal substrate should have a high Debye temperature, as well as low surface corrugation to minimize diffracted flux. Additionally, the surface must be stable against reconstruction and corrosion, and will ideally be passive, reducing the accumulation of dirt and adsorbates on the surface that lower the specularly reflected intensity. Finally, the surface needs to be simple and inexpensive to prepare.

Crystal surfaces for helium diffraction studies are typically prepared by cleaving or annealing *in-situ*. Because the crystal must be mounted to the tip of a rotor, these methods are not ideal for this experiment. However, some surfaces are so inert

that they can be prepared *ex-situ* and will remain clean enough to still reflect helium well. The crystal chosen for this experiment, is hydrogen passivated Si(111). Lithium fluoride was also examined as a possible crystal substrate, and the details of this examination may be found in [55].

To use Si(111) as an atomic mirror which is prepared *ex-situ*, the crystal must be passivated. A freshly cleaved Si(111) surface will have dangling bonds and the surface will not be passive enough to keep contaminants from sticking to the surface. By wet etching the surface, the dangling bonds are terminated with hydrogen, which passivates the surface, resulting in Si(111)-(1×1)H [56, 57]. The mirrors are created from phosphor doped 200 μm thick wafers with a pre-existing oxide layer and a miscut angle of $\pm 0.1^\circ$ acquired from Virginia Semiconductor. The 100 mm diameter waferes were pre-cut into smaller circles [58], but the cutting process scratched the crystal surfaces and rendered the area within 2 mm of a cut unusable. The wafers are cut to their final size for etching and insertion into the vacuum chamber by a dicing saw, removing the damaged regions. The wet etching process used to create this surface is now described.

3.2.2 Crystal Preparation and Wet Etching

To produce an atomically flat passivated Si surface, the crystal must first be cleaned as much as possible to remove contaminants, especially organic residues. The procedure described below is based on suggestions from the group of Dr. Ken Shih at The University of Texas at Austin, and the group of Dr. Bill Allison at the Cavendish Lab in Cambridge [57]. All chemicals used are LP grade or better, and the ammonium fluoride is MB grade. Additionally, any water used in the procedure is ultra pure and has a resistivity of greater than 17 MΩ cm. Finally, all beakers, bottles, containers, and tools used to handle the wafers or in contact with the chemicals are made of

PTFE or PFA. The procedure takes place under the fume hood, and the working surface is covered by a layer of clean room wipes.

The initial cleaning is done by sonicating the crystal in isopropanol for 10 minutes, followed by another 10 minutes in acetone. The crystals are then rinsed with water to remove these organic solvents before continuing the cleaning process. The next step is to place the wafers in an RCA I solution which removes organic residues and grows a uniform oxide layer. This solution consists of 1 part 30% ammonium hydroxide, 1 part 30 % hydrogen peroxide, and 5 parts water. The water is heated in a water bath on a hot plate to 80°C and the ammonium hydroxide and hydrogen peroxide are added just before the wafers. The solution is stirred gently with a magnetic stirrer on the hot plate. The crystal is left in the RCA I solution for 5 minutes, at which point it is removed from the solution and rinsed with water. Each rinsing step described here involves a significant quantity of water, typically at least half a liter is poured over each crystal during the rinsing phases.

The next step is the wet etching, which is done in a 40 % solution of ammonium fluoride. Before the crystals are etched, the disolved oxygen in the ammonium fluoride must first be removed by bubbling ultra pure argon through the solution for at least 40 minutes prior to inserting the crystals. For the crystals used here, etch times between 10 and 15 minutes are optimal. Shorter times do not completely remove the surface oxide layer, while longer times result in deep etch pits in the surface. The crystals used in the experiments described below are etched for 12 minutes. Unlike in the previous steps, the ammonium fluoride solution is not stirred or disturbed in any way while etching the crystals. Even continuing to gently bubble argon in the solution negatively affects the resulting surface.

After the etching is complete, the crystal is rinsed a final time. The passivated surface is hydrophobic, and the water does not wet the surface if the etching is

successful. Any water that remains on the crystal can be removed by blowing it off with argon. Storing the crystal in an argon environment until it can be inserted into vacuum preserves the pristine state of the surface.

3.2.3 Crystal Characterization

To characterize the passivated surfaces, they are examined using an atomic force microscope (AFM). The AFM is operated in tapping mode, and scanning 2μm by 2μm regions of the wafer was found to provide a good compromise between examining a large area and scan speed. Figure 3.2 shows AFM images of a wafer that was wet etched. Image (a) shows a clean surface, with large atomically flat regions. The steps seen in the image are bilayer steps from one crystal plane to another, and are 3.8 Å high. The image is of the wafer 15 minutes after removing it from the etching solution. Image (b) shows the same wafer after waiting an additional 28 minutes. The apparent increase in surface roughness is likely due to the oxidation of the crystal. The height profiles shown in images (c) and (d) are taken along the white lines shown in (a) and (b) respectively. The surface roughness shown in (c) is less than an Å in magnitude, and is probably noise in the AFM rather than real artifacts on the crystal surface. However, the increased surface roughness shown in (d) shows grains approximately 2Å in height, and is likely showing real contamination of the surface. This indicates the importance of minimizing the time the crystals spend in atmosphere before inserting them into the vacuum chamber.

3.3 The Experimental Apparatus

This section describes the apparatus constructed to create, slow, and detect a beam of supersonic helium. The description of the apparatus follows the path of the beam, starting with the vacuum system, followed by the slowing apparatus, and

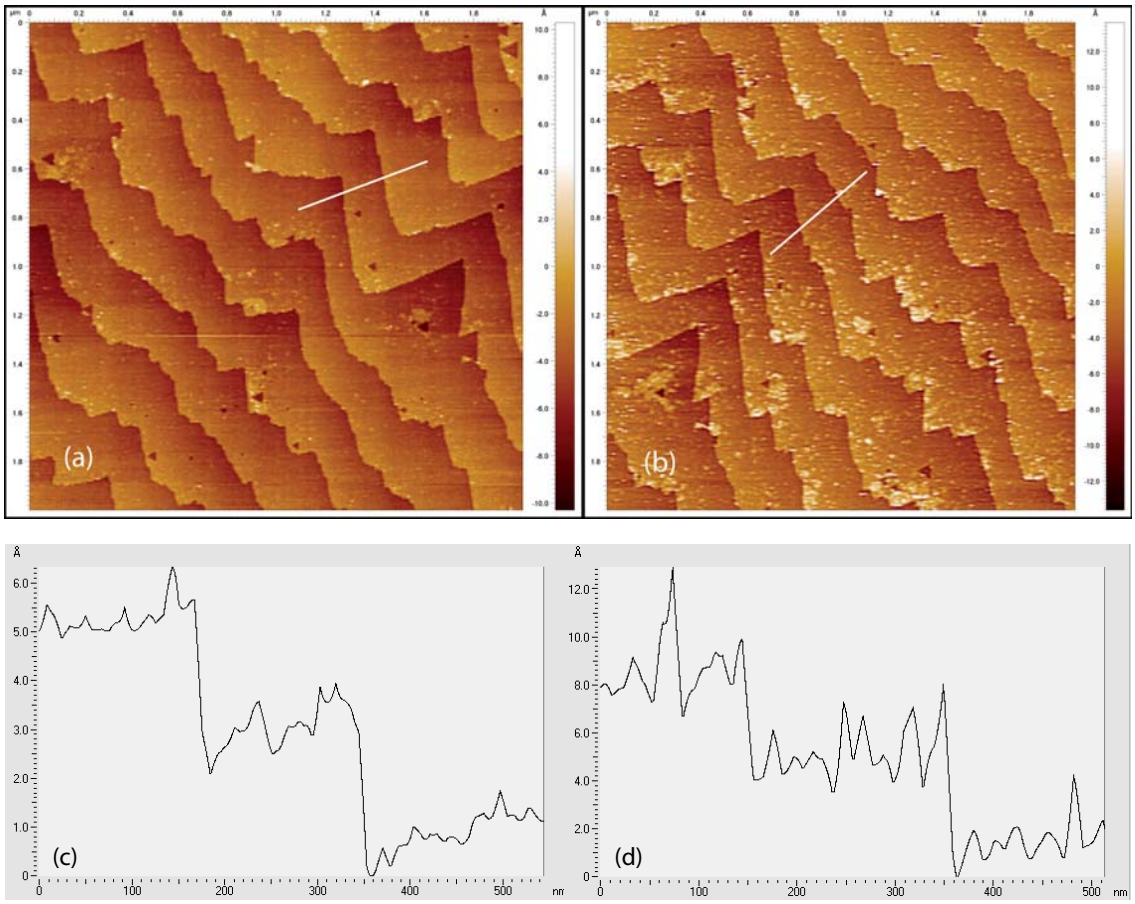

Figure 3.2: An AFM image of a passivated Si(111) crystal. Image (a) and (b) are 2 μm by 2 μm regions of the same wafer, 15 minutes and 43 minutes after etching respectively. Image (a) shows atomically flat regions separated by bilayer steps between crystal planes, each of which are 3.8 Å high. The increased noise seen in image (b) shows the surface contamination (most likely oxidation) from the additional 28 minutes the wafer spent in atmosphere. The height profiles shown in (c) and (d) are taken along the white lines in (a) and (b) respectively. The roughness of the crystal terraces in (a) and (c) are likely noise in the AFM, but the increased roughness seen in (b) and (d), showing grains approximately 2Å high, show the increased contamination of the surface.

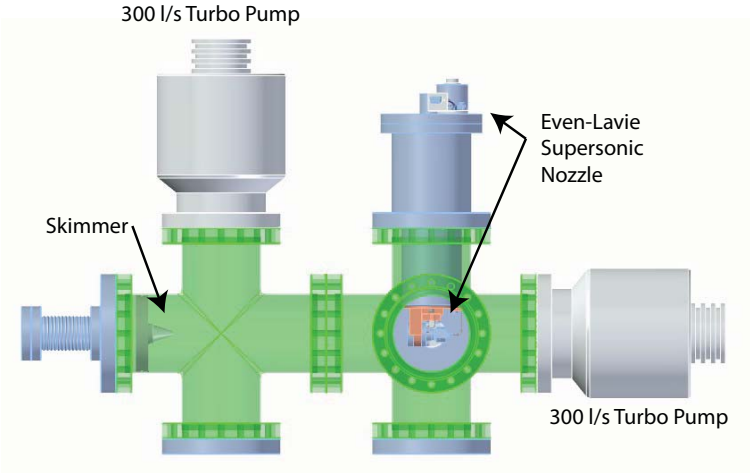

Figure 3.3: A CAD image of the nozzle chamber. The chamber is pumped by two 300 l/s turbo pumps (indicated). The Even-Lavie nozzle hangs from a pool type cryostat into the center of the right 6 inch cross (the beam propagates left in this image). The skimmer has a 5 mm aperture and is 5 cm long and sits $\approx$ 30 cm from the nozzle.

finally a description of the detection system and experimental control system.

3.3.1 The Vacuum System

The vacuum system for this experiment is divided into three distinct sections, the beam creation chamber, the rotor chamber, and the detection chamber. The beam creation chamber consists of one 4-way 6 inch cross, and one 6-way 6 inch cross. These were purchased from HPS [59] and are custom made to have a tolerance of .5° in parallelism and perpendicularity of the flanges, and are constructed to .02 inch tolerances. These tolerances ensures that the Even-Lavie nozzle sits in the center of the chamber and that the beam propagates along the center-line of the chambers.

An image of the beam creation chamber with the Even-Lavie valve is shown in figure 3.3. The chamber is pumped by two 300 l/s Leybold turbo pumps, which maintain a pressure of 10^{-6} torr when the nozzle is operating at a repetition rate

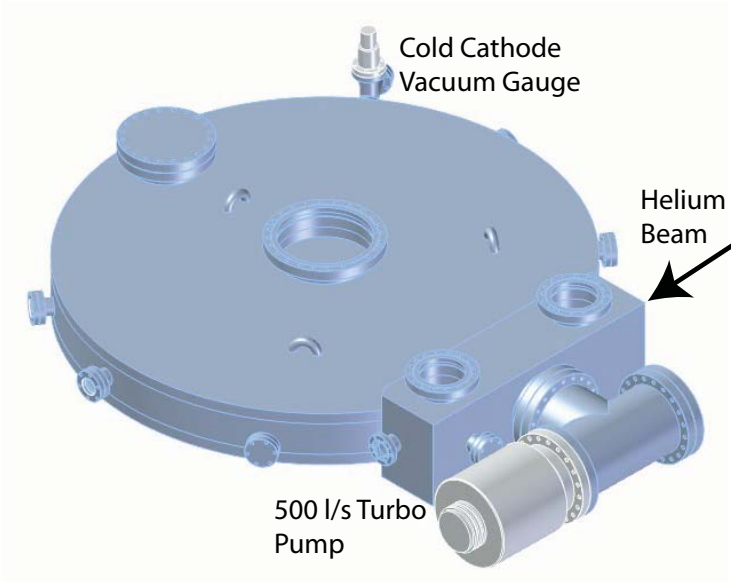

Figure 3.4: A CAD image of the rotor chamber. The chamber is pumped by a 500 l/s turbo pump (indicated). The chamber walls are 1 inch thick 304 stainless steel, and the chamber disc is 112cm in diameter and 12.7cm high.

of 5 Hz (10^{-8} torr when the nozzle is not operating). This is monitored by a cold cathode vacuum gauge (not shown). The nozzle hangs from a pool type cryostat into the center of the 6-way cross. The beam propagates 30 cm to the 5 cm skimmer, which has a 5 mm aperture. From the skimmer, the beam passes into the rotor chamber, which is connected to the nozzle chamber by a flexible bellows, allowing the helium beam to be aligned to the rest of the experiment.

The rotor chamber has to accommodate the rotor, which is the key element in the design of the chamber as a whole. As described in section 3.3.2, the rotor is 100.8 cm in diameter, and so the chamber needs to be large enough to allow it to spin freely. In addition, the beam needs to enter the interaction region and be detected. This led to a custom chamber design, as seen in figure 3.4. The chamber is pumped by a 500 l/s Varian turbo pump that keeps the chamber at a pressure of $2 \cdot 10^{-8}$ torr while the nozzle is firing (10^{-9} torr when the nozzle is not operating). The chamber

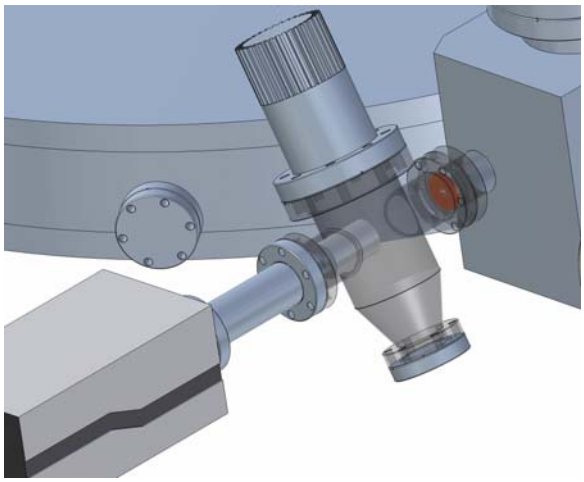

Figure 3.5: A CAD image of the detection chamber. The chamber is separated from the rotor chamber by a 5 mm aperture. This provides differential pumping, keeping the detection volume from experiencing too high a gas load from the direct beam. The chamber is pumped by a 70 l/s turbo pump. Figure Courtesy Max Riedel.

is constructed from 1 inch thick stainless steel, and can be seen in figure 3.4. The chamber disc is 112 cm in diameter and 12.7 cm high. Due to the large surface area of the top and bottom of this disc, atmospheric pressure moves the top and bottom of the chamber together by 4 mm when the chamber is under vacuum.

The final chamber section is the detection chamber, illustrated in figure 3.5. The detection chamber is separated from the rotor chamber by a 5 mm aperture constructed from a copper disc (used instead of an annular gasket where the rotor chamber connects to the detection chamber). The detection chamber is pumped by a 70 l/s Varian turbo pump, and houses the detector used to observe the beam when it does not interact with the rotor. This detector is located 1.42 m from the nozzle. The direct beam detector was originally directly connected to the rotor chamber, but this negatively impacted detector performance, probably due to the significant gas load entering the detection region without adequate pumping.

The entire apparatus is shown in figure 3.6. The support structures are made

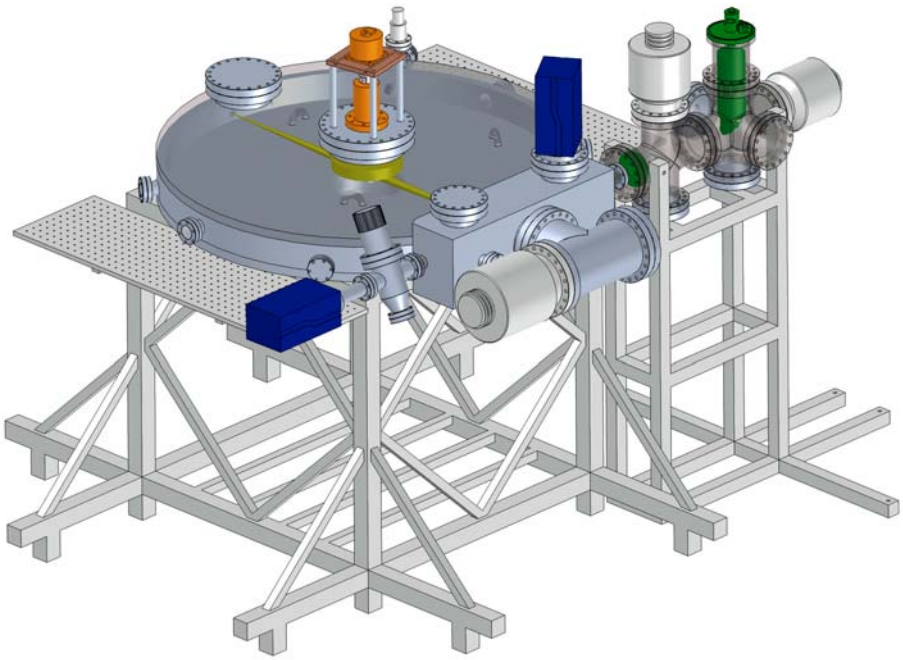

Figure 3.6: A CAD image of the entire apparatus, including the chamber support structure and optical breadboards. The Even-Lavie nozzle and skimmer are shown in green, the SRS residual gas analyzers used to detect the beam are shown in blue, the rotor and spindle are shown in yellow, and the motor and rotary feedthrough which drive the rotor are shown in orange. Figure Courtesy Max Riedel.

from tubular aluminum welded together to support the 600 kg chamber mass. The Even-Lavie nozzle is shown in green, as is the skimmer and skimmer mount. The SRS residual gas analyzers used for beam detection are shown in blue. The rotor and spindle are depicted in yellow, while the rotary feedthrough and motor which drive the rotor are shown in orange.

3.3.2 The Rotor

There are several design requirements for the rotor itself, which informed its final construction. As can be seen from equation 3.2, to significantly slow a supersonic beam, the crystal should be capable of reaching velocities up to half of the initial beam speed. A supersonic beam of pure helium from a nozzle cooled with liquid nitrogen to 77 K will have a velocity of 900 m/s and so the rotor was originally designed to reach tip velocities of up to 450m/s. Another design constraint comes from the use of rotary motion to approximate linear motion of the crystal. Clearly, over small distances, a larger rotor will more closely approximate linear motion. This is important because the helium beam has a finite spatial and temporal extent when it arrives at the crystal ($\approx 100\,\mu$s). The angle of reflection changes during the arrival of the beam, limiting the fraction of the reflected beam which travels in the desired direction. This fanning effect leads to a loss of flux, and should ideally be minimized by making the rotor as large as practical. Finally, an effective and simple method of mounting the crystals to the tip of the rotor, which does not overly compromise the rotor's strength, needs to be found.

3.3.2.1 Initial Design

The centripetal force needed to hold a mass in circular motion, $F = \frac{mv^2}{r} = m\omega^2 r$, where F is the force, m is the mass of the body, r is the radius of the orbit, v

is the velocity, and ω is the angular velocity, dictates the design of the rotor, which must be strong enough at every point to support the force exerted by the section of the rotor at larger radius. The force at each position can be approximated (assuming no transverse forces in the rotor) as

$$F = \int_{r_c}^{r_{max}} \rho A(r) \omega^2 r \mathrm{d}r, \tag{3.5}$$

where r_{max} is the maximum radius of the rotor, r_c is position the force is being measured, ρ is the density of the rotor material, and $A(r)$ is the cross sectional area of the rotor. The force at each position must be compared to the yield tensile strength of the rotor material times the cross sectional area. Equation 3.5 shows that a light and strong material is ideal for the rotor. While composites would be ideal, their lack of vacuum compatibility poses a serious problem. The chosen rotor material is grade 5 titanium, a high strength alloy which is also vacuum compatible. Another constraint on the rotor design is the vacuum itself, since the rotor must fit into a practical vacuum chamber. This led to a compromise design with a rotor radius of 50.4 cm at the center of the crystal.

A CAD image of the rotor can be seen in figure 3.7. The overall length of the rotor is 101.5 cm. Since the inside of the rotor needs to be stronger than the tips, as seen in equation 3.5, the rotor is tapered linearly from a minimum thickness of 0.3175 cm where the rotor joins the mirror holder, to a maximum thickness of 2.54 cm, with the rotor a constant thickness over its center 5.08 cm. This taper reduces the mass at large radius, reducing the force exerted, while placing material where it is needed near the center of the rotor. The crystal mirrors, which fit directly into the rotor are 5 mm in diameter and are held in place by an aluminum set screw, allowing for an exposed crystal diameter of 4 mm.

While an analytical calculation of the forces placed on a the rotor is prohibitively complex, finite element analysis permits a numerical calculation of the

Figure 3.7: A CAD image of the rotor. The rotor has 50.4 cm radius at the center of the crystal holder, and a maximum thickness of 2.54 cm in the center. The diameter of the rotor itself is linearly tapered from 0.3175 cm where the rotor joins the crystal holder, to the maximal thickness 2.54 cm from the middle of the rotor. The small hole seen in the middle of the rotor allows precise centering of the rotor via a pin in the spindle. A 5 mm diameter crystal (4 mm exposed) is held in place by a set screw. Figure Courtesy Max Riedel.

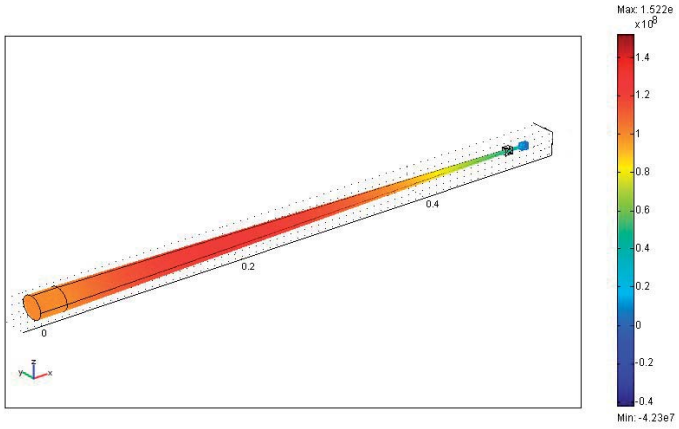

Figure 3.8: A finite element calculation of the stress exerted on the rotor by rotation at 140 Hz. This is expressed in MPa, since it is the force per unit area of material, rather than the absolute force, which is relevant. For reference, the peak stress place on the rotor is only a factor of 5 smaller than the yield tensile strength of the titanium alloy used.

Figure 3.9: A CAD image of the rotor as held by the initial spindle. The spindle is rigidly coupled to the shaft of the ferrofluidic feedthrough. The feedthough attaches to a $10 - 4.5$ inch zero length reducer, which supports the spinning parts, as well as the motor, which is mounted above the feedthrough via a flexible coupling. Figure Courtesy Max Riedel.

forces. Figure 3.8 shows a finite element calculation of the stresses placed on the rotor by rotation at 140 Hz, measured in MPa. This is expressed as a pressure because the strength of the material is proportional to the cross sectional area, and thus the relevant number is not the absolute force placed on the rotor, but rather the force per unit area. The maximum stress placed on the rotor by this speed of rotation is 152 MPa, which is well below the yield tensile strength of the titanium allow used (880 MPa).

The rotor is held in vacuum by a stainless steel spindle which is rigidly coupled to the shaft of a ferrofluidic rotary feedthrough [60]. Since an object will rotate stably around the axis of lowest or highest inertia, one goal in the construction of the spindle was to ensure that the rotor would be spinning about the axis with the highest moment of inertia (the lowest moment being rotation about the centerline of the rotor itself). To do this, a disc was added to the spindle, with most of the mass on the outside

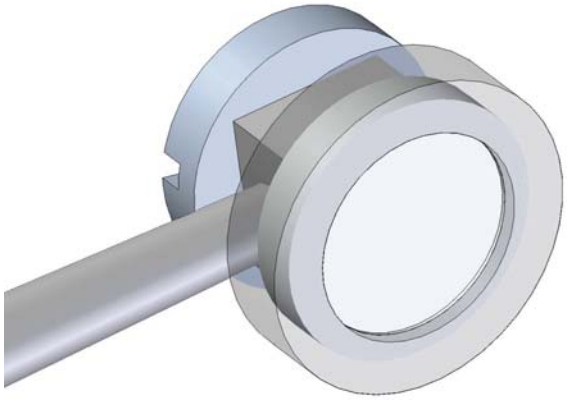

Figure 3.10: A CAD image of large crystal holder attached to the rotor. This enables 12.7 mm diameter crystals to be mounted to the rotor tips. The crystals are held between a threaded aluminum circular plate and a titanium locking ring, exposing a 9 mm diameter region of the crystal. The circular plate connects via a threaded hole to an adapter, which fits in the rotor where the smaller crystals would sit. Figure Courtesy Max Riedel.

of the disc. The spindle also has screw holes drilled around the outside of the disc, which were added to allow for the balancing of the rotor and spindle should this be necessary. The rotating parts of the experiment, along with the $4-10$ inch zero length conflat reducer they are suspended from, are shown in figure 3.9.

As is explained in section 3.3.2.2, the maximum speed of the rotor was initially limited to 42 Hz due to the vibrational stability of the spindle. Since the stress placed on the rotor scales quadratically with the rotation rate, being limited to 42 Hz reduced the expected stresses by an order of magnitude. This allowed the size of the crystal to be increased, becauase the rotor is strong enough support the increased mass of a larger crystal and crystal holder at the tip for these velocities. The larger crystal increases the area over which the beam can reflect from the rotor tip, increasing the reflected flux. An image of the larger crystal holder can be seen in figure 3.10. The larger crystal holder increases the diameter of the crystals from 5 mm to 12.7 mm. These crystals are clamped between a threaded aluminum circular plate,

and a titanium locking ring, with an exposed crystal diameter of 9 mm. The back of the threaded plate has a tapped hole, which screws onto a third piece which fits into the rotor via the smaller crystal mounting threads. This system has the additional advantage of simplifying the mounting and removing of crystals from the rotor, reducing the time the crystals and the vacuum chamber are exposed to air during a crystal exchange.

3.3.2.2 Vibration and Balancing

The spinning parts of the experiment are all constructed with dimensional tolerances of .001 inches, which ensures that the rotor apparatus should be nearly balanced to begin with. To verify the initial balancing, the rotor was mounted with the motor axis horizontal, instead of vertical as it is in the chamber, and statically balanced. This provides the rough balancing needed to initially spin the rotor, since even a small static imbalance would lead to potentially large vibrations when spinning.

To balance the spinning system, the vibration of the rotor must be measured. This is accomplished in two ways. The first method uses a laser beam to measure the amplitude of the pendulum modes of the shaft. The laser is directed though the chamber, with half of the beam blocked by the lower part of spindle. As the shaft vibrates, the spindle moves in the laser beam, blocking more or less of the beam, which is detected by monitoring the intensity of the light leaving the chamber via a photo diode. Vibrations are detected to an accuracy of $5\,\mu$m using this setup. A knife edge mounted to a micrometer is used to calibrate the readings. Having assembled the system with the greatest accuracy possible, using feeler gauges to ensure that parts are properly aligned, a runout of $70\,\mu$m is observed (over the course of one slow rotation the spindle moves over a range of $70\,\mu$m).

The second method of monitoring the vibrations does not directly measure

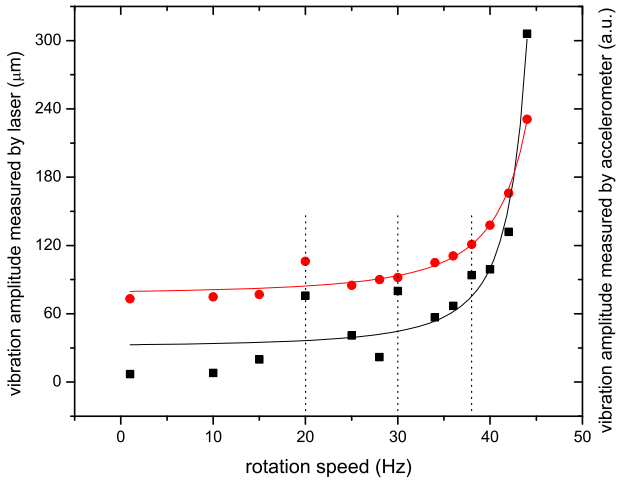

Figure 3.11: This plot shows the amplitude of the rotor vibration at several rotational speeds. The vibrational amplitude measured by the laser is shown in red, while the amplitude measured by the accelerometer outside the chamber is shown in black. The dashed lines indicate observed resonances as observed by either the laser or the accelerometer. The solid line is a Lorentzian fit for the resonance which ultimately limits the rotor velocity.

the motion of the rotor, but instead monitors the vibrations which are transmitted to the chamber. This is done by rigidly mounting an accelerometer chip to the top of the vacuum chamber. While an externally mounted accelerometer cannot measure the amplitude of vibrations on the rotor, or determine the mode of vibration, it does have the advantage of detecting vibrational modes the laser is insensitive to. The experiment is set up to monitor the level of vibration measured by the laser and the accelerometer, and automatically spins down the rotor if the vibrations grow too large.

The vibrations observed using these detection systems are significantly larger than anticipated when the rotor was designed. A plot of vibrational amplitude for different rotation speeds is shown in figure 3.11. The plot shows the vibrational amplitude measured by the laser in red, and the amplitude measured by the accelerometer

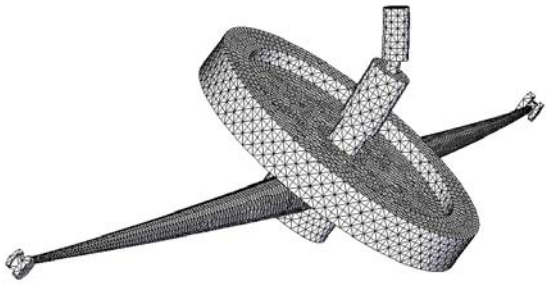

Figure 3.12: Finite element calculation of rotor vibrational modes. This was performed using the built in finite element package in Solid Edge. The vibrational mode shown has a resonance frequency of 47 Hz. The calculation does not take into account the fact that the rotor and spindle are made from different materials, as only one material can be simulated. As such the entire structure is simulated as stainless steel, since most of the weight is in the stainless steel spindle. Additionally, the shaft of the spindle, which is found to be region that flexes most, is stainless steel. Reducing the weight of the rotor (it is titanium in the experiment) would likely result in a slightly higher resonant frequency in the real system.

in black. The dashed lines indicate resonances by at least one of these methods. Additionally, these resonances can be heard as increased noise from the chamber, and the chamber vibrations are easily felt by placing a hand on the chamber. The solid lines in figure 3.11 are Lorentzian fits of the limiting resonance, which is predicted be lie around $48 - 50$ Hz. From these fits, the peak vibrational amplitude can be estimated to be $1000 - 2000\,\mu$m which is too large to be comfortably passed through, and thus limits the rotor velocity to 42 Hz.

A finite element calculation of the vibrational modes of the rotor and spindle was performed to investigate the source of this unexpected resonance using the built in finite element package of CAD program Solid Edge. This analysis shows the nature of the resonance which limits the rotor velocity, and the results of this calculation are shown in figure 3.12. The finite element calculation predicts two pendulum modes, one at 34 Hz and one at 47 Hz, in addition to a twisting mode at 35 Hz. Several

approximations are made in this calculation. The first approximation is that the coupling between the spindle and feedthrough is truly rigid. The second approximation comes from the fact that this finite element package cannot simulate two materials simultaneously. The entire rotor structure is simulated as stainless steel, since the stainless steel spindle comprises most of the mass, and because the point which the calculations indicate is flexing the most is the stainless steel shaft of the spindle. In the experimental setup, the rotor is made from titanium and weighs about 700 g less than the stainless steel rotor of the finite element calculation. This decrease in mass should result in increased resonant frequencies in the actual rotor. A final uncertainty lies in the fact that this calculation does not include the shaft of the feedthrough, as the rigidity of the bearings is unknown. However, given that the finite element calculation predicts resonant frequencies in the same region where they are observed leads to a belief that the pendulum modes of the spindle shaft are the cause of the resonance which limits the rotor velocity.

3.3.2.3 Improved Spindle Design

Led by the finite element calculation described in section 3.3.2.2, the spindle was redesigned with the goal of increasing the resonant frequencies of the vibrational modes. To do this, the spindle material is changed to titanium instead of stainless steel, and redesigned to have minimal mass, which reduced the weight of the spindle from 4.07 kg to .88 kg. The reduction in spindle mass means that the rotor does not spin around the axis with the greatest moment of inertia, but this is viewed as an acceptable trade off. Additionally, the method of clamping the spindle to the shaft of the feedthrough is completely redesigned. Rather than having a thin shaft extend from the top of the spindle, which is connected to the shaft of the feedthrough by a clamp, the spindle is designed to be a clamp itself. This design is similar to the

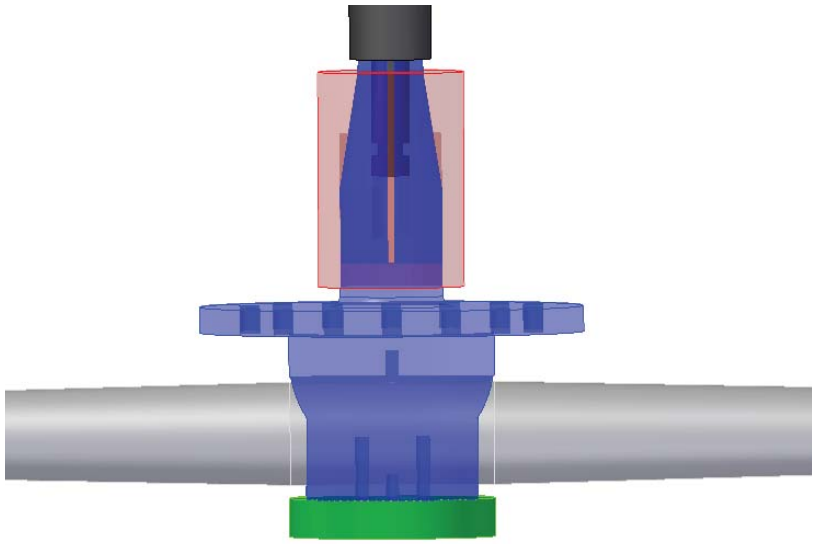

Figure 3.13: A CAD image of the redesigned spindle. The shaft of the feedthrough is shown in black. The spindle is blue, with the clamp that tightens the collet of the spindle onto the feedthrough shaft shown in red. The rotor is silver, and the clamp that holds the rotor in the V-groove of the spindle is green.

collet used in machine tool holders, and is be much more rigid, as well as being self centering. The method of clamping the rotor is also changed to a V-groove, instead of a cylindrical clamp, which is self centering and easier to machine with high accuracy. A CAD image of the new spindle can be seen in figure 3.13.

Finite element calculations performed on the new spindle predict resonant frequencies well above 100 Hz, which is a significant improvement on the old spindle. With the new spindle in place, the rotor has successfully been spun up to 60 Hz without any serious resonances appearing. The large spindle holders are still in place, and the increased mass at the rotor tip means that higher speeds are deemed to be too large a risk. With the vibrational stability of the rotor improved by the new spindle, the rotor can spin at the speeds required for potential future experiments.

3.3.3 Detection

An SRS [61] residual gas analyzer (RGA) is used to detect the beams. The experiment uses two RGAs, one to detect the slowed helium beam (RGA 200), and one which detects the beam when the rotor is out of the way (RGA 100). Both RGAs work by ionizing the helium via electron bombardment in an ionization region. The ions are then focused into a quadrupole mass spectrometer which provides mass filtering, and detected by a continuous dynode electron multiplier. The output of the electron multiplier is fed directly into a current pre-amplifier, and from there the signal is read by the data acquisition card in the computer.

Because the beams are detected by ionizing the helium, the detection efficiency for helium is on the order of 10^{-5} [62]. The ionization efficiency should scale with the time the atoms spend in the ionization volume, a slow beam should have a greater detection efficiency than a fast one.

3.3.4 Control Systems

The experiment is controlled by a data acquisition card made by National Instruments (NI PCI-6120), run by a custom control program written in LabVIEW. The motor driving the rotor is a 380 watt Moog motor fitted with an 8000 line encoder, and is servo controlled by a Galil controller. This permits the rotational frequency of the rotor to be controlled to a few thousandths of a Hz. The servo controller is also capable of controlling the static position of the rotor to within 0.25°, and communicates with the computer control system through a serial connection.

The data acquisition card also controls the timing of the experiment. The nozzle must pulse at the correct phase of the rotor's rotation for the beam to reflect at the correct position. Figure 3.14 illustrates the method used to achieve the required phase locking between the rotor and the nozzle. The motor controller delivers a

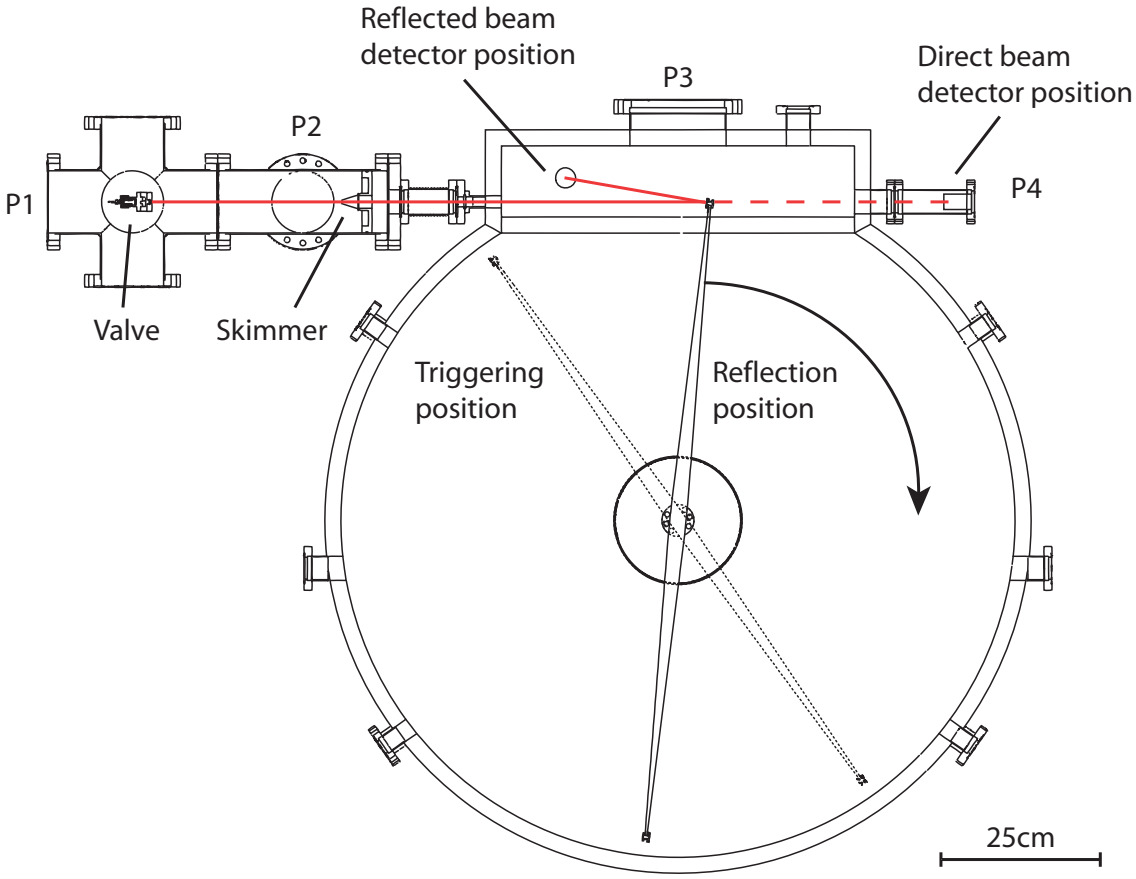

Figure 3.14: Illustration of rotor triggering and beam reflection positions. The motor servo controller delivers a TTL pulse to the data acquisition card when the rotor passes through the triggering position. The card then waits an empirically determined time before pulsing the Even-Lavie nozzle. The beam is illustrated in red and the vacuum pumps and detectors are indicated in the schematic.

TTL pulse to the data acquisition card when the rotor passes though the triggering position, which is 30° before the static reflection position. The card then waits an empirically determined time before pulsing the Even-Lavie nozzle to create the beam. This ensures that the beam will arrive at the nozzle at the correct time for the beam to be reflected to the slowed beam detector. The control program is capable of running the experiment at one shot per rotor rotation, but the vacuum pressure in the nozzle chamber rises too much when this is done at high rotor velocity. As such, the control system can be told to skip rotations, limiting the nozzle to a certain maximum repetition rate.

3.4 Data and Results

Because the maximal rotation rate of the rotor is limited by vibrations to speeds well below the initial design parameters, as explained in section 3.3.2.2, it is necessary to reduce the speed of the helium beam below that obtained by cooling the nozzle to 77 K. At this temperature, a pure helium beam has a velocity of around 900 m/s. Even at the maximum rotation rate achieved of 60 Hz the rotor tip velocity is 190 m/s, which is not enough to significantly slow the beam. However, by mixing the helium with neon in a 50:50 mixture, the beam velocity is reduced to around $500 - 550$ m/s, depending on the backing pressure of the nozzle. Lowering the velocity of the incoming helium beam also has the advantage of improving the Debye-Waller factor, which should lead to a larger fraction of the helium being specularly reflected. The data described here is taken with the old spindle, limiting the rotor to 42 Hz and a tip velocity of 133 m/s.

For the experiments and data described below, the specific gas mixture used is a 48:52 mixture of helium and neon, with a nozzle reservoir pressure of 42 psi. This mixture produces a helium beam with a velocity of 511 ± 9 m/s.

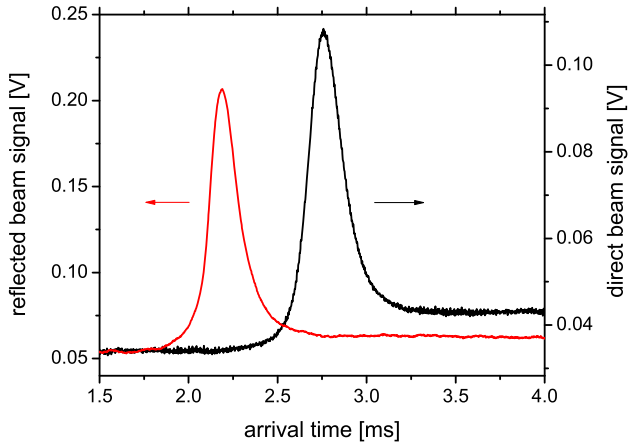

Figure 3.15: Time-of-flight profiles of the direct and reflected beams. The direct beam is shown in black, while the reflected beam is shown in red. Both beams have the same velocity and temperature, indicating good specular reflection of the beam from the crystal. The reflected beam arrives earlier in time because the distance from the nozzle to the rotor and on to the reflected beam detector is shorter than the distance between the nozzle and the direct beam detector. This also explains the slightly narrower time-of-flight profile of the reflected beam. Since the beams are detected by different detectors, setup in different configurations, the amplitudes are not comparable. Each curve is an average of 50 beam pulses.

3.4.1 Static Reflection

The first test of the rotor is to verify that the passivated Si crystals act as good atomic mirrors. Time-of-flight profiles of the beam are examined for both the direct and reflected beams, without spinning the rotor. Since the rotor is static, the velocity profile of the beam should not change for atoms specularly reflected from the crystal. To find the rotor position where the beam is reflected to the detector, the motor controlling the rotor is turned off and the rotor is adjusted by hand to locate the reflected signal on the detector. At this point, the servo motor is turned on and the angle of the rotor is adjusted in .225° steps to empirically optimize the reflected signal.

A comparison of the time-of-flight profiles for a direct and statically reflected beam is shown in figure 3.15. The direct beam is shown in black and the reflected beam is shown in red. Both curves are averages of 50 beam pulses. The temperatures of these beams are calculated from Gaussian fits to be 249 ± 36 mK for the direct beam and 254 ± 41 mK for the reflected beam. These temperatures are effectively the same, indicating that the beam is specularly reflecting from the crystal and is not being heated by inelastic processes. The reflected beam arrives earlier than the direct beam in the because the distance traveled by the reflected beam is shorter than that traveled by the direct beam. This also explains the slightly narrower time-of-flight profile of the reflected beam. Because different detectors, set up in different configurations, are used to detect the direct and reflected beams, the amplitude of the signals is not comparable.

3.4.2 Beam Slowing

To spin up the rotor and reflect the beam from the receding crystal, the beam must arrive at the correct time so that the rotor is at the correct position and angle to reflected the beam to the detector. A pulse is sent from the motor controller when the rotor reaches a position 30° from the optimized static reflection position. After a delay, the nozzle is pulsed, allowing the beam and rotor to interact at the optimal position. To find the correct delay for a particular rotor speed, the delay times which cause the rotor to just miss the beam on either side are found, and the optimal delay is taken to be the average of these delays. At the highest rotor speeds, these delays are optimized to the microsecond.

Time-of-flight profiles of the reflected beam for different crystal velocities are shown in figure 3.16. The arrival time of the reflected beam is delayed with increasing mirror velocity, which shows that the beam is being slowed by reflection from the

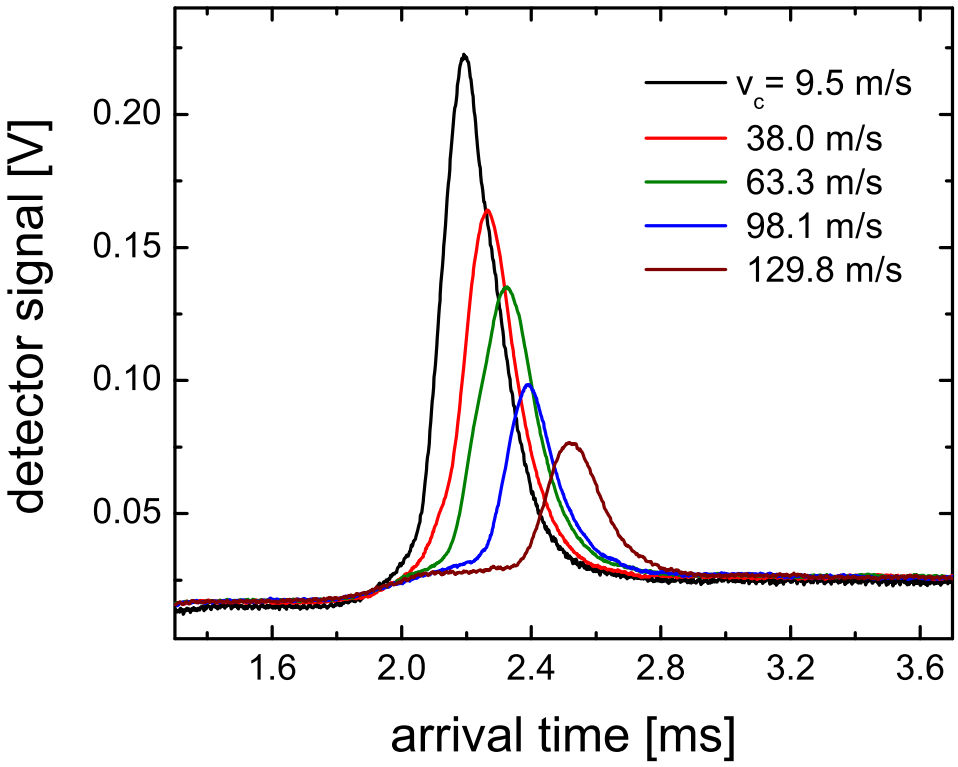

Figure 3.16: Time-of-flight profiles of reflected beams for different crystal velocities (v_c). The later arrival times of the beam for greater crystal velocity show that the beam is being slowed by reflection from a receding crystal. The decrease in amplitude of the beams is the result of the rotor changing in angle during the time the beam arrives, creating a fanning effect. Each curve is an average of 50 beam pulses.

receding crystal. Each curve is the average of 50 signals, which artificially broadens the peak. The time-of-flight profiles of the individual shots are slightly narrower, but there is some slight variation in their arrival time due to slight differences in the rotor velocity and a small jitter in the valve triggering signal sent from the encoder.

At the higher crystal velocities, a shoulder appears in the signal before the arrival of the slow beams in figure 3.16, and in all cases the detected signal is higher after the beam passes through the detector than before. This is due to a rise in the background helium in the chamber, which results from the atoms in the beam that do not hit the rotor and are inelastically scattered from the chamber walls, or from atoms that inelastically scatter from the rotor itself. Since the chamber and rotor are at room temperature, these atoms are faster than the atoms in the beam, especially when the beam has been slowed by reflecting from a fast receding crystal, which is why the rise in background helium pressure in the chamber can be observed even before the slowed beam arrives for the fastest crystal velocities.

Another element of time-of-flight profiles shown in figure 3.16 is the decreasing detected amplitude of the slowed beam for increasing crystal velocity. One cause of this is the temporal extent of the incoming beam, which arrives at the position of the rotor over a period of around 115 μs. During this time, the angle of the rotor changes in proportion to the angular velocity, and the reflected beam is spread in plane of rotation. This fanning effect is mitigated by a large rotor, which allows for lower angular velocities for the same corresponding crystal velocity, but at the highest crystal velocities, the approximation of linear motion by a rotating object can still become important. The fanning effect was quantified in the simulations performed in [33]. The other cause for the decrease in the signal height is the decay of the crystals over time in the vacuum chamber. This is discussed in more detail in section 3.4.3.

A summary of the slowing results is shown in figure 3.17. The slowed beam

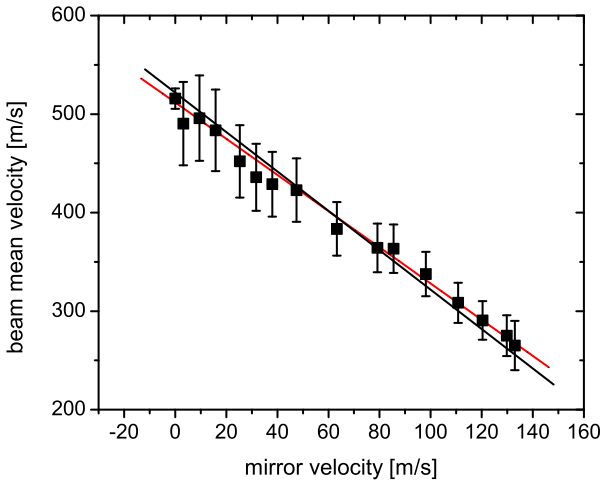

Figure 3.17: The calculated velocity of the reflected beam plotted against the velocity of the crystal. The red line is the least square fit of the data, while the black line shows the expected velocity dependence from simulation.

velocity is calculated from the arrival time of the beam and from the beam path length calculated from the nozzle to the rotor and from the rotor to the detector. This path length depends on the phase of the rotor, and is calculated from the nozzle delay time used to optimize the reflected beam signal. The beam velocities are calculated by performing a Gaussian fit of the data for each individual nozzle pulse. The mean of the center of the Gaussian fits is taken as the arrival time, while the standard deviation of the centers is used to calculate the error.

The red line shown in figure 3.17 is the least square fit of the data, while the black line is the velocity dependence predicted by simulation. While the dependence expected from simulation lies within the error bars for the measurement, there seems to be a real difference in the slope of the two lines. This is most likely a result of the manner in which the optimal nozzle delay is found. Since the distance between the reflection point and the detector is short compared with the total travel distance, the

calculated slow beam velocity is quite sensitive to the rotor position calculated from the optimal nozzle delay. Thus, a problem with the method of choosing the optimal delay could result in the difference in slope observed.

The maximum velocity removed by the rotor is 246 m/s, resulting in a 265 m/s helium beam. This is not a fundamental limit, but is the result of the 42 Hz rotation limit imposed by vibrations.

3.4.3 Crystal Lifetime

The reflected intensity of the helium beam is observed to decay over the course of the experimental runs. The crystal itself appears to be the source of this decay, since attempts to continue to use a crystal after it had been in the vacuum chamber overnight resulted in significantly reduced reflected intensity. The crystal surface should be passive, but it is possible for hydrocarbons to stick to the surface, reducing the intensity of the reflected beam. A cold trap was inserted into the vacuum chamber to try to improve the lifetime of the crystal. Another possibility is that something in the vacuum system is damaging the passivation. Any stray electrons or metastable atoms would deposit enough energy to remove the hydrogen terminating the silicon bonds, reducing the passivation and increasing the probability of background gas sticking to the surface.

A possible source of electrons and metastable atoms in the chamber is the filament of the RGA used to ionize the atoms for detection. Some of the electrons from the filament could escape the ionization region and hit the crystal surface. Additionally, atoms passing through the detection region could be excited by the electrons to a metastable state, creating atoms which will remove the passivating hydrogen if they strike the surface of the crystal. To investigate whether the RGA filament is the cause of the crystal decay, a data set is taken with the RGA kept off, and only turned

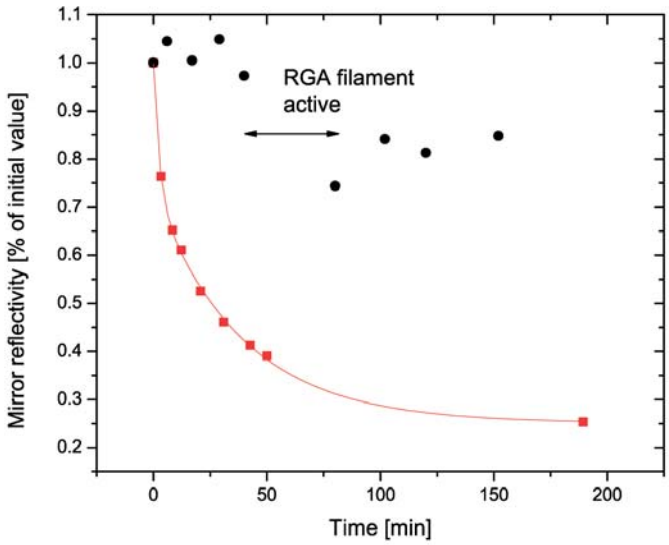

Figure 3.18: This shows the decay in the reflected intensity from a passivated silicon crystal over time. The red data comes from a normal data run with the rotor spinning. The rotor is periodically stopped to measure the reflected intensity. The red curve is a least square fit of this decay as a double exponential. The data shown in black is the reflected intensity from a dedicated intensity run without spinning the rotor and where the RGA is only on to measure the reflected intensity, except for the period noted in the figure where the RGA is left on deliberately to observe its effect.

on to measure the reflected intensity at specific intervals. The RGA is also left on for a specific interval to observe its effect directly.

This data set, along with the measurement of the decay from a normal data run is shown in figure 3.18. The red data comes from a normal data run with the rotor spinning. The rotor is stopped at certain times to measure the static reflected intensity, and these measurements are shown in red. The red curve is a least square fit of the decay as a double exponential decay with time constants of 2.5 and 38.5 minutes. The data shown in black is the reflected intensity where the RGA is only on measure the reflected intensity (about 1 minute) for each measurement, except for the period noted in the figure where the RGA is left on deliberately to observe its effect. The rotor is not spun during the taking of this data set. As seen in the figure,

the RGA does have an effect on the reflected intensity of the beam, and electrons or metastable atoms from the RGA may be de-passivating the surface of the crystal.

Since it appears that de-passivation of the silicon surface is at least partially responsible for the decay in the reflected intensity, other crystal options were considered. Lithium fluoride is also a very passive crystal and can be cleaved to create clean smooth surfaces. Lithium fluoride was studied as a possible replacement for silicon as a crystal surface and this work is described in [55].

3.5 Potential Applications

The ability to reduce the velocity of a supersonic beam of helium without heating is an enabling technology and this section discusses several potential applications of this technology. One possible use of the slowed beam is to study atom-surface interactions. The short nozzle pulse compared to the slow beam propagation time leads to an expected energy resolution of $2\,\mu$eV [33]. This energy resolution, combined with ability to quickly and precisely tune the velocity of the beam, makes it ideal for studies of the atom-surface potential. The slowing would also permit this potential to be mapped in a new energy range. There is no reason that this technique would not work with Helium 3 as well, permitting even greater energy resolution using the spin echo technique [47], along with precise control over the velocity and an extended energy range of interrogation.

Another potential application is in the field of atom interferometry [63–65]. Ground state noble gases are insensitive to magnetic fields to first order, and helium's low polarizability makes it comparatively insensitive to electric fields, reducing the potential for errors due to stray fields. These traits make it an ideal candidate for use in interferometry. For a white light geometry Mach-Zender interferometer of length L rotating at a rate Ω, the phase shift ϕ imparted on the beam is $\phi = \frac{2\pi\Omega L^2}{vd}$

where v is the velocity of the beam and d is the grating spacing [66]. Similarly, for an acceleration a perpendicular to the direction of atomic motion the phase shift is $\phi = \frac{2\pi a L^2}{v^2 d}$ [66]. Reducing the velocity clearly leads to a larger phase shift, and greater resolution. The rotor is an excellent candidate for providing a reduced velocity helium beam for atom interfometry. In interfometry, the flux is also quite important as shot noise may otherwise limit sensitivity. Methods for increasing the flux downstream of the rotor are discussed in [33]. Finally, resolution can also be increased by reducing the grating spacing. In fact, the periodic structure of a crystal lattice can provide an extremely short grating which is well matched to the de Broglie wavelength of the incoming helium beam. It may be possible to construct the atomic analog of the perfect single crystal neutron interferometer [44, 66], where instead of using the crystal as a transmission grating, it would instead act as a reflection grating.

Chapter 4

The Atomic and Molecular Coilgun: Using Pulsed Magnetic Fields to Slow Supersonic Beams

As discussed in chapter 2, supersonic beams are remarkable for their general applicability. Almost any atomic species, and many molecules as well, can be seeded into a supersonic beam and cooled by the expansion. Because they serve as a nearly universal source of cold atoms, they are the ideal starting point for a general method of controlling atomic motion.

Since supersonic expansion is a very general method for producing cold atoms and molecules, a general method of slowing and control provides a nearly universal tool for producing cold samples for further study. Magnetic fields provide this ability. On the atomic level nearly all atomic species are magnetic either in their ground state or in an easily accessible metastable state. The physical mechanism through which magnetic fields allow for control of atoms is the Zeeman effect.

Using a series of pulsed electromagnetic coils, and inspired by the macroscopic coilgun [67], it is possible to produce the large fields needed to slow and stop supersonic beams in a table-top scale experiment. The principle of operation of the atomic coilgun is described in detail. Experiments using two generations of coilgun are presented. The first coilgun has 18 pulsed coil stages, and slowed metastable neon from 461 m/s to 403 m/s. The apparatus, including the coils and electronics used to drive them, is described. The second generation coilgun, which slowed and stopped metastable neon and molecular oxygen is also presented in this chapter. The exper-

iments are described in detail, as are the methods used to produce and characterize the magnetic fields produced in the coils. The concepts and experiments presented in this chapter are also described in [68–73].

4.1 Atoms in External Magnetic Fields

In the presence of an external magnetic field, the internal energy levels of a particle can undergo a shift, which is referred to as the Zeeman effect in low field, and the Paschen-Back effect in high field. The description of these effects presented here is limited to the fine structure level, though the mechanics described can easily be extended to the hyperfine structure.

At the fine structure level, each energy level is described by the total angular momentum $J = L + S$ where L is the orbital angular momentum and S is the spin angular momentum of the state. The perturbing Hamiltonian which leads to the fine structure splitting couples L and S, making J a good quantum number. Each state with total angular momentum J has $2J + 1$ sublevels, labeled by m_J, giving states $|J, m_J\rangle$ which are eigenvectors of the operators J^2 and J_z (assuming z is the quantization axis) with eigenvalues

$$J^2 |J, m_J\rangle = J(J+1) |J, m_J\rangle , \tag{4.1}$$

and

$$J_z |J, m_J\rangle = m_J |J, m_J\rangle . \tag{4.2}$$

The sublevels of the state J are degenerate, but an external magnetic field breaks this degeneracy due to the additional perturbing Hamiltonian

$$H_B = -\vec{\mu}_{\text{atom}} \cdot \vec{B}, \tag{4.3}$$

which shifts the sublevels. The magnetic moment of the atom, $\vec{\mu}_{\text{atom}}$, must be determined in order to calculate the shifts in the energy levels caused by the magnetic

field. In the $L-S$ coupling regime (as opposed to jj coupling which is more common for heavier atoms) the magnetic moment of an atom is

$$\vec{\mu}_{\text{atom}} = -\mu_B \sum_i \left\langle g_l \vec{l_i} + g_s \vec{s_i} \right\rangle = -\mu_B \left\langle g_L \vec{L} + g_S \vec{S} \right\rangle, \tag{4.4}$$

where μ_B is the Bohr magneton, and g_L and g_S are the "g-factors" for the orbital and spin angular momenta respectively. Thus the Hamiltonian becomes

$$H_B = \mu_B \left\langle g_L \vec{L} + g_S \vec{S} \right\rangle \cdot \vec{B}. \tag{4.5}$$

The manner in which the magnetic moment is calculated depends on the ratio of this Hamiltonian to the Hamiltonian which produces the fine structure splitting.

4.1.1 The Zeeman Effect

If the interaction due to the Hamiltonian in equation 4.5 is much smaller than the fine structure splitting, then J will remain a good quantum number. In this regime, called the Zeeman regime, the individual orbital angular momentum vector $\vec{L}$ and the spin angular momentum vector $\vec{S}$ precess around the sum of these vectors, the total angular momentum vector $\vec{J}$. This in turn precesses around the external magnetic field, where the projection of $\vec{J}$ in the direction of the magnetic field (quantization axis) determines the magnetic moment of the atom. In this picture, as $\vec{L}$ and $\vec{S}$ precess around $\vec{J}$, their projections along the direction of the magnetic field are rapidly changing, meaning that these projections (L_z and S_z) are not good quantum numbers. However, since the projection of $\vec{J}$ along the quantization axis remains constant, the eigenvalues of the operator J_z will remain constant, making m_J a good quantum number. This situation is illustrated in figure 4.1.

Since m_J is a good quantum number, but not m_L and m_S, the equation for the Hamiltonian (equation 4.5) must be expressed in terms of $\vec{J}$. This is done using

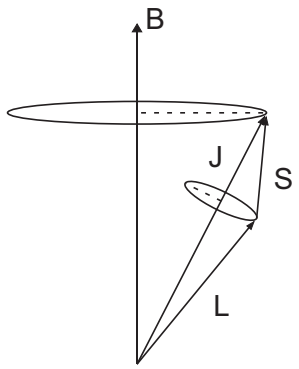

Figure 4.1: A pictorial illustration of the angular momentum coupling of $\vec{L}$ and $\vec{S}$ which takes place in the low field limit. The spin and orbital angular momenta are coupled by the spin-orbit interaction, and as such $\vec{L}$ and $\vec{S}$ precess around the total angular momentum $\vec{J}$. This in turn precesses around the magnetic field B. As $\vec{L}$ and $\vec{S}$ precess, their projection along B changes, meaning that m_L and m_S are not good quantum numbers, but the projection of $\vec{J}$ along B remains constant and so m_J is a good quantum number.

the relations

$$\vec{S}_{ave} = \frac{\vec{S}\cdot\vec{J}}{\vec{J}^2}\vec{J} \tag{4.6}$$

and

$$\vec{L}_{ave} = \frac{\vec{L}\cdot\vec{J}}{\vec{J}^2}\vec{J}. \tag{4.7}$$

These provide the needed projections onto $\vec{J}$, and can be calculated using

$$\vec{S}\cdot\vec{J} = \frac{\left(\vec{J}^2 + \vec{S}^2 - \vec{L}^2\right)}{2} \tag{4.8}$$

and

$$\vec{L}\cdot\vec{J} = \frac{\left(\vec{J}^2 + \vec{L}^2 - \vec{S}^2\right)}{2}. \tag{4.9}$$

Using the above equations gives the Hamiltonian

$$\begin{aligned} H_B &= \mu_B g_L \vec{J}\cdot\vec{B}\left(\tfrac{J(J+1)+L(L+1)-S(S+1)}{2J(J+1)}\right) \\ &+ \mu_B g_S \vec{J}\cdot\vec{B}\left(\tfrac{J(J+1)+S(S+1)-L(L+1)}{2J(J+1)}\right). \end{aligned} \tag{4.10}$$

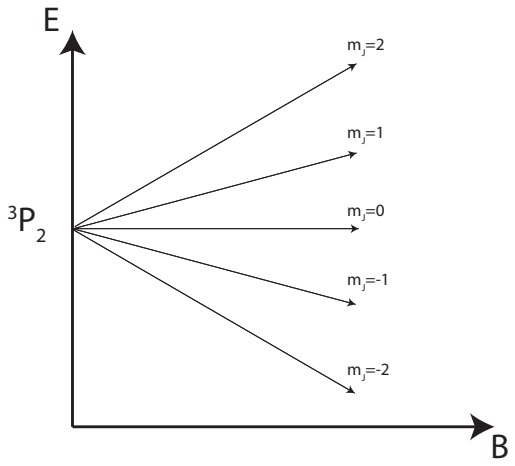

Figure 4.2: This graph gives an example of the Zeeman shift of the m_J sub-levels of a $^3P_2$ state. Since $J = 2$ there are 5 sublevels ($m_J = -2, -1, 0, 1, 2$). The g-factor for this state is $g_J = \frac{3}{2}$ which means that the negative m_J levels lose energy in the field and are high field seekers, while the positive m_J levels gain energy and are low field seekers.

Using the g-factors $g_L = 1$ and $g_S \approx 2$, and the relationship $\vec{J} \cdot \vec{B} = BJ_z$ allows the Hamiltonian to be simplified to

$$H_B = \mu_B B \left(1 + \frac{J(J+1) - L(L+1) + S(S+1)}{2J(J+1)}\right) J_z = \mu_B g_J B J_z, \tag{4.11}$$

where g_J is the Landé g-factor. This Hamiltonian shifts the energy of the m_J sub-levels by

$$\Delta E = \mu_B m_J g_J B. \tag{4.12}$$

This shift is illustrated in figure 4.2. Particles which gain potential energy in the magnetic field are called low-field-seeking, since magnetic field gradients will exert a force pushing them away from regions of high field. Particles that lose potential energy in the field will feel the opposite force and are referred to as high-field-seeking.

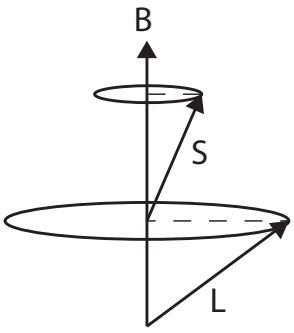

Figure 4.3: A pictorial illustration of the angular momentum coupling of $\vec{L}$ and $\vec{S}$ which takes place in the high field limit. The coupling of the spin and orbital angular momenta from the spin-orbit interaction is broken by the high magnetic field, and as such $\vec{L}$ and $\vec{S}$ precess around the magnetic field independently. As such, the projections L_z and S_Z along B are constant, meaning that m_L and m_S are good quantum numbers.

4.1.2 The Paschen-Back Effect

In high field, the perturbing Hamiltonian produces a shift which is larger than the fine structure splitting. In this limit the interaction is known as the Paschen-Back effect. Because the magnetic interaction is stronger than the spin-orbit coupling, $\vec{L}$ and $\vec{S}$ precess independently around the magnetic field. To correctly calculate the energy levels as a function of magnetic field, the fine structure Hamiltonian must be considered a perturbing interaction on top of the shifts induced by the magnetic field.

In this regime, the coupling between $\vec{L}$ and $\vec{S}$ is broken, and both precess independently, making the projections L_z and S_Z along $\vec{B}$ constant. This means that m_L and m_S are good quantum numbers. The angular momentum coupling of L and S in the Paschen-Bach regime is illustrated in figure 4.3. Starting from equation 4.5 the Hamiltonian is now

$$H_B = \mu_B B \left(g_L L_z + g_S S_z\right). \tag{4.13}$$

This results in an energy shift of

$$\Delta E = \mu_B B \left(m_L + 2m_S\right). \tag{4.14}$$

For intermediate fields, the full perturbing Hamiltonians from the magnetic field and the spin-orbit coupling must be accounted for by diagonalizing the sum of the individual Hamiltonians. However, for the special case where the state of interest has maximal J, m_J resulting from the maximal possible projections of $\vec{S}$ and $\vec{L}$, the state transitions smoothly between low and high field with the same slope of the shift as a function of B. This is convenient since it is desirable to target such a state for slowing by the coilgun as this state will have the greatest possible energy shift in the field. Thus, for sub-levels targeted by the coilgun, it is frequently unnecessary to diagonalize a Hamiltonian to calculate the energy shifts.

4.1.3 Adiabatic Following and Spin Flips

The Hamiltonian described in the previous sections, depends on the projection of the spin onto the axis of the magnetic field. An important consideration is what will happen when the apparent direction of the magnetic field changes. There are two possible outcomes of a change in the direction of the field: the projection of the spin may adiabatically follow the field for slow enough changes in the field, or the original quantization axis may be lost causing expectation value of the spin will change. This last event is referred to as a spin flip, as a particle can change from being low-field-seeking to high-field-seeking, or vice versa.

It is important to determine the limits in the rate of change for which the projection of the spin will adiabatically follow the field. As noted above, the angular momentum vector $\vec{J}$ in the Zeeman regime, or $\vec{S}$ and $\vec{L}$, in the Paschen-Back regime precess around the magnetic field. This is known as Larmor precession, and the rate of this precession is

$$\omega_L = \frac{\mu_{\text{atom}}|\vec{B}|}{\hbar}, \tag{4.15}$$

which is referred to as the Larmor frequency. The projection of the spin onto the

magnetic field will adiabatically follow the field so long as the Larmor frequency is much larger than the rate of change of the direction of the magnetic field. If this condition, expressed as

$$\omega_L \gg \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{\vec{B}}{|B|}\right), \tag{4.16}$$

is met, then the equations for the Hamiltonian and for the energy shift do not depend on $\vec{B}$, but instead will depend on $|\vec{B}|$, which is hereafter simply referred to as B. This condition is easily met for most magnetic fields described in this dissertaion, and special note is made of situation where it may not be met.

4.2 Magnetic Fields from Electromagnetic Coils

Pulsed electromagnetic coils are used to produce the high magnetic fields necessary to effectively control the velocity of a supersonic beam. From the differential form of Ampère's law

$$\nabla \times \vec{B} = \mu_0 \vec{J}, \tag{4.17}$$

where $\vec{J}$ is current density and μ_0 is the permeability of free space, it is possible to derive the Biot-Savart law

$$\vec{B} = \frac{\mu_0}{4\pi}\int_C \frac{I\left(\mathrm{d}\vec{\ell}\times\vec{x}\right)}{|\vec{x}|^3}, \tag{4.18}$$

where I is the current and $\mathrm{d}\vec{\ell}$ is the infinitesimal path of the current [74] (this also assumes a wire with zero cross-sectional area). From this, the magnetic field from a particular current path can be found. Of particular interest is the magnetic field from a circular loop of wire of radius r. The cylindrically symmetric field generated by this loop can be found exactly as [75]

$$B_z = \frac{\mu_0 I}{2\pi}\frac{1}{\sqrt{(r+\rho)^2+(z)^2}}\left[K(k^2)+\frac{r^2-\rho^2-z^2}{(r-\rho)^2+z^2}E(k^2)\right] \tag{4.19}$$

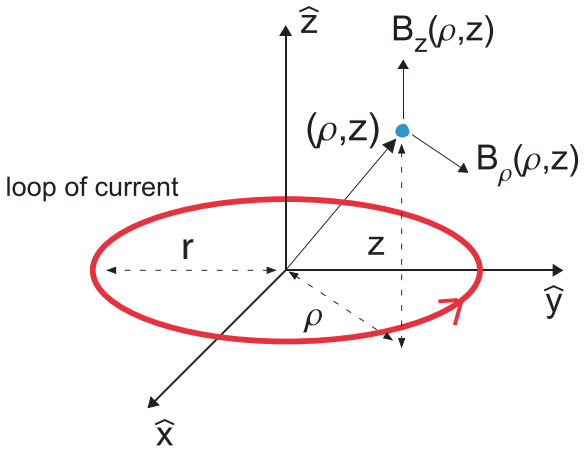

Figure 4.4: This figure illustrates the geometry of the coil which produces the magnetic field calculated in equations 4.19 and 4.20. The $\vec{B}$ field is expressed in its components along the ρ and z directions (there is no component in the ϕ direction).

$$B_\rho = \frac{\mu_0 I}{2\pi\rho}\frac{z}{\sqrt{(r+\rho)^2+(z)^2}}\left[-K(k^2)+\frac{r^2+\rho^2-z^2}{(r-\rho)^2+z^2}E(k^2)\right] \tag{4.20}$$

where

$$k^2 = \frac{4r\rho}{(R+\rho)^2+z^2} \tag{4.21}$$

is the argument of the complete elliptic integrals K and E of the first and second kind respectively. The situation described by the equations above is illustrated in figure 4.4. Finding exact analytical solutions to these equations is often difficult, or even impossible depending on the argument of the elliptic integrals, and solutions are typically found numerically.

Another reason the magnetic fields of the coils used in the experiments described here are typically found numerically is that the above equations are for coils and fields in vacuum. While the particles in the beam are traveling through vacuum and are of a low enough density not to appreciably change the character of the field, the materials near and around the coils will have a significant effect on the fields produced. To account for this, Maxwell's equations can be modified, resulting in the

macroscopic Maxwell's equations [74]. Doing this, a new field $\vec{H}$ is defined as

$$\vec{H} = \frac{1}{\mu_0}\vec{B} - \vec{M} \tag{4.22}$$

where $\vec{M}$ is the magnetization of the material, allowing Ampère's law to be rewritten as

$$\nabla \times \vec{H} = \vec{J}. \tag{4.23}$$

For many materials, $\vec{B}$ and $\vec{H}$ are linearly related via the equation $\vec{B} = \mu\vec{H}$, where μ is the magnetic permeability of the material. However, for ferromagnetic materials, $\vec{B}$ is a non-linear multivalued (due to hysteresis) function of $\vec{H}$, making the effects of the material on the field much more difficult to calculate. Calculating the fields produced by the coils numerically allows the full non-linear behavior of the real materials near the coils to be taken into account.

4.2.1 Switching of Electromagnetic Coils

The final aspect of producing magnetic fields via pulsed coils that needs to be addressed is the switching of the coils. Maxwell's equations give the relationship

$$\nabla \times \vec{E} + \frac{\partial \vec{B}}{\partial t} = 0, \tag{4.24}$$

which says that a changing magnetic field will produce an electric field. When this is applied to a coil, the result is Lenz's law, which states that

$$\mathcal{E} = -N\frac{\mathrm{d}\Phi}{\mathrm{d}t}, \tag{4.25}$$

where $\mathcal{E}$ is the electromotive force, N is number of turns in the coil, and Φ is the magnetic flux through the coil. Thus, switching a coil will produce a potential that induces a current in the direction that opposes changes in the field. This makes it difficult to quickly switch the current in a coil, and doing so creates large voltage spikes.

The same effect is also responsible for the creation of eddy currents in materials near a coil. These currents must be accounted for when computing the time dependence of the magnetic fields.

The switching of a coil can also be examined by looking at the voltage and current in the coil as a function of time. The problem can be simplified by approximating the coil as an ideal resistor in series with an ideal inductor. The voltage across this circuit is

$$V = IR + L\frac{\mathrm{d}I}{\mathrm{d}t}, \tag{4.26}$$

where L is the inductance of the coil and R is the resistance of the coil. For instantaneously switched drive voltages, the current when switching the coil on is

$$I(t) = \frac{V}{R}\left(1 - e^{\frac{-tR}{L}}\right), \tag{4.27}$$

and similarly the current when switching the coil off is

$$I(t) = \frac{V}{R}e^{\frac{-tR}{L}}. \tag{4.28}$$

To achieve fast switching times, it is desirable to have either high resistance, or low inductance. Because high currents are also needed to produce large magnetic fields, high resistance is unfavorable, and thus low inductance is an important aspect of the coil design. Similarly, greater applied voltages can produce larger $\frac{\mathrm{d}I}{\mathrm{d}t}$, both for turn on and for turn off, but this requires the electronics driving the coil to be capable of withstanding these large voltages.

4.3 Principle of Operation

As described in section 4.1, the energy levels of atoms or molecules with a non-zero magnetic moment undergo a shift in a magnetic field. This energy level shift results in a change in the potential energy of the atom or molecule. As an

atom or molecule in a low-field-seeking state moves from a region of low field to a region of high field, its potential energy increases. Since force $F = -\nabla U$ where U is the potential energy, a spatially varying magnetic field produces a force, and a low-field-seeking atom or molecule is repelled from regions of high field. Conservation of energy dictates that a particle loses kinetic energy equal to its gain in potential energy. Thus, as a low-field-seeking atom enters a region of high field it will lose a kinetic energy equivalent to the Zeeman shift of the internal states. For a particle moving in an electromagnetic coil, the loss of kinetic energy T in the Zeeman regime (assuming $L - S$ coupling) is

$$\Delta T = -\mu_B g_J m_J \Delta B, \tag{4.29}$$

and similarly in the Paschen-Back regime the change in kinetic energy is

$$\Delta T = -\mu_B \left(2m_S + m_L\right) \Delta B. \tag{4.30}$$

While a low-field-seeking particle loses kinetic energy moving from a region of low field to a region of high field, the reverse is also true. Hence, a low field seeking particle moving in an electromagnetic coil loses energy as it enters the coil, and if the field is unchanged it regains the kinetic energy as it exits the coil. However, if the coil is switched off quickly while the particle is in a region of high field, the particle loses kinetic energy. This process can be repeated through multiple coils, each removing kinetic energy from the beam, until the desired velocity is attained. The slowing sequence for a single coil is illustrated in figure 4.5.

While equations 4.29 and 4.30 are exact for an ideal coil that switches instantaneously, a real coil is more complicated. In a real coil, the magnetic field turns off over some period of time, during which the particle will continue to move in the coil. For this reason, it is important to optimize the coil and switching electronics

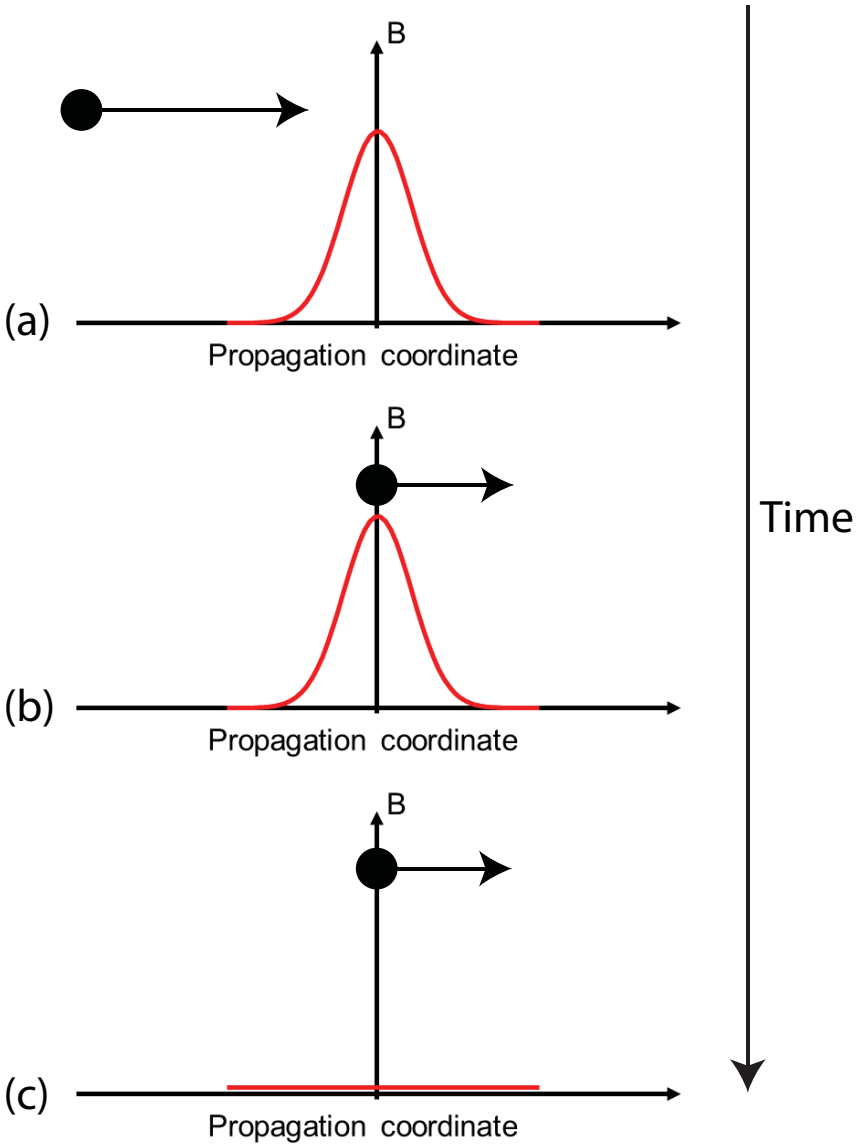

Figure 4.5: A pictorial illustration of the principal of operation of the coilgun. In (a) a low field seeking atom approaches an energized coil with the coil's field illustrated in red. The atom enters the magnetic field region and is slowed down as illustrated in (b). The field is then turned off, allowing the atom to exit the coil without regaining the lost energy, as shown in (c).

to reduce the switching time as much as possible. The time varying potential adds significant complication, and numerical simulation is used to determine the energy lost by a hypothetical particle.

Since the amount of kinetic energy lost by a particle depends on its position in the coil when the coil is switched off, the timing of the coil switching is critically important. Varying the switching time of a coil changes the velocity and thus the arrival time of the particles at the next coil. The switching times of the coils are calculated numerically by simulating the trajectory of a single particle traveling through the experiment. This particle is henceforth referred to as the synchronous atom or molecule, and can generally be thought of as being in the phase space center of the slowed bunch of atoms. The average initial velocity of the pulsed beam can be found before the start of the experiment by observing the beam with no magnetic fields present and simply calculating the velocity from the time of flight to the detector. The synchronous particle is then assigned this velocity as its initial velocity. The short opening time of the Even-Lavie supersonic nozzle means that the atoms leave the nozzle at a precise time, and so the time-of-flight measurement of the initial velocity of the beam is very precise. From this, the arrival time of the beam at the entrance of the coilgun can be computed easily and precisely, and the entire coil timing sequence can be referenced to the valve pulse.

4.3.1 Phase Stability

For an instantaneously switched coil, maximal slowing is achieved by switching the coil when the synchronous particle is at the peak field location. However, this is not the ideal position to switch a real coil for several reasons. The first reason is that a real coil will require some time to turn off, and thus the switching position that imparts maximal slowing is moved towards the front of the coil (the direction from

which the particles enter). The other reason to switch the coil when the synchronous particle is in front of the coil is to maximize the volume in phase space of the slowed bunch of particles.

The volume of phase space which the coilgun is able to slow is an important consideration, since the number of slowed particles is a key factor in many future uses the coilgun. It is desirable to maximize the phase space acceptance of the coilgun, where the ideal acceptance maximizes its overlap with the phase space occupied by the incoming supersonic beam. While it may not be possible to match the phase space acceptance of the coilgun to the beam, any increase in the accepted phase space volume is beneficial. For the purposes of the following discussion, it is assumed that the coils are switched instantaneously. While this is not a valid assumption in general, the qualitative behavior of the slowed bunch in this scenario is easier to ascertain, and the same behavior is exhibited, though with greater complexity, in the case of real coils with finite switching times.

As stated above, maximal slowing is achieved, by timing the coils to turn off when the synchronous particle is at the peak magnetic field. This timing scenario is illustrated in figure 4.6(a). In this scenario, the particles which are behind and ahead of the center of the slowed bunch are not slowed as much as the synchronous particle, since they do not experience as great a field when the coil is switched. Because particles which are ahead of the center of bunch are not slowed as much, they get out of phase with the coil pulse sequence and are eventually lost from the slowed bunch. Particles that are behind the center of the bunch are not slowed as much, so they catch up with the center of the bunch, but any particle that gets ahead of the synchronous atom is eventually lost, leading to a vanishingly small region of phase space which stays in phase with the full coil sequence.

Instead of switching when the synchronous atom is at the peak of the field,

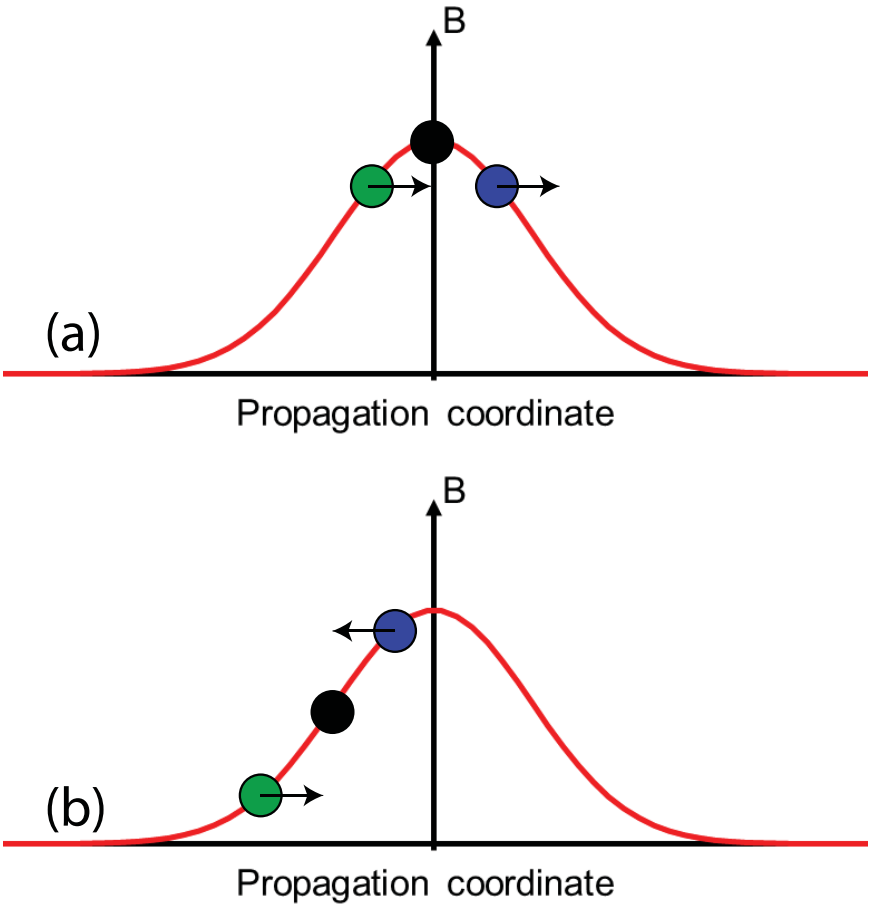

Figure 4.6: An illustration of how the phase space acceptance of the coilgun, and phase stability of the slowed bunch is determined by coil timing. In (a), the coil is switched when the synchronous particle (black) is at the peak of the magnetic field. Particles which are behind (green) or ahead (blue) of the synchronous particle are not slowed as much, since they do not see as large field when it is switched. The arrows indicate the relative shift in velocity of the particles compared to the synchronous particle. The phase space acceptance window is small because any particle which is ahead of the synchronous particle is not slowed as much, and is lost from the slowed bunch. In (b), the coil is switched before the synchronous particle arrives at the field peak. Here particles which are behind(ahead of) the center of the bunch are slowed less(more), leading to a region of phase stability in the slowed bunch.

the coil can instead be switched before it arrives at this position. This is illustrated in figure 4.6(b). In this case, the particles that are ahead of the synchronous atom in the bunch are slowed more, and thus will be pushed back towards the center of the bunch. Similarly, particles which are behind the center of the bunch are slowed less than the synchronous particle and catch up. This creates a region of phase stability, where the bunch remains in phase with the coil pulses for the entire coil sequence. This also reduces the necessary accuracy of the timing sequence for the coil pulses, as slight deviations in coil timing are sel-correcting. A slight error in coil timing may mean that the bunch is slowed slightly more(less) than predicted by a particular coil, but because the bunch was slowed more(less), the next coil in the coilgun slows it less(more) because the bunch is farther(closer) from(to) the peak field when the coil is switched. This makes the coilgun quite robust against small timing errors.

A brief comment on notation is useful at this point. When producing the timing sequence, the numerical simulation of the synchronous atom switches each coil at a pre-defined position in that coil. This position can also be referred to as a phase angle using the following conversion. The midpoint between two coils is considered to be a phase angle of 0° and switching when the synchronous atom or molecule is in the center of a coil is labeled 90°, with intermediate positions linearly converted to an angle. This convention comes from previous work with particle accelerators, and the concept of phase stability is similar to that found in colliders [76, 77], and the pulsed electric field decelerator [15, 78]. The concept of phase stability in the coilgun is discussed in [21, 25, 69–71].

4.3.2 Electromagnetic Coils as a Waveguide

While the longitudinal stability of the slowed bunch of particles is addressed in the previous section, the transverse behavior of particles in the coilgun also has a

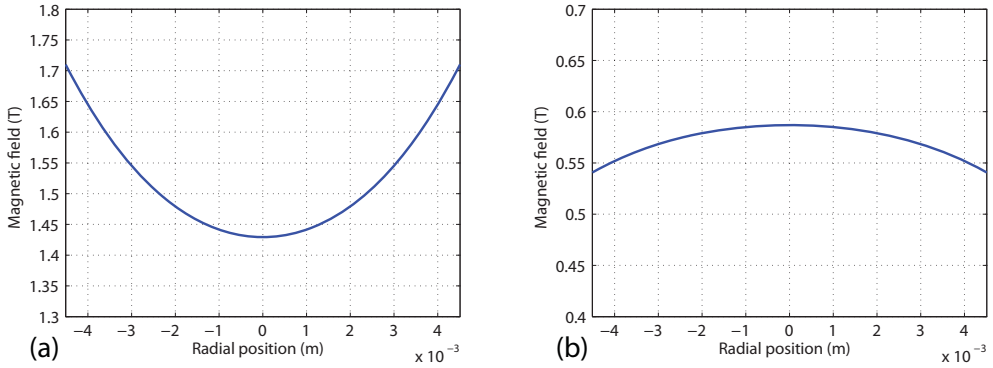

Figure 4.7: This figure illustrates the manner in which the magnitude of the magnetic field varies with radial position in a 1.05 cm diameter coil. The field at the center of the coil ($z = 0$) is plotted in (a). The magnitude of the field is at a radial minimum on the axis of the coil at this axial position, which provides a focusing force to low-field-seeking particles. The field, as a function of radial position well outside of the coil is plotted in (b). Here, the radial dependence of the field magnitude is at a maximum on the axis of the coil, providing a de-focusing effect for low-field-seekers. The interplay between the focusing effects inside the coil and the de-focusing outside the coil is complex, and depends on the velocity of the particles passing though the coil, as well as the switch-off timing.

large effect on the number of slowed atoms or molecules. Maxwell's equations dictate that there can be no free space maximum in magnetic field. Though the magnitude of the magnetic field on the axis of a coil (z-axis) is at a maximum at center of the coil, this cannot be a true local maximum. Instead, this point at the center of a coil is a saddle point, and is at a minimum of the field radially. The magnitude of the field of a coil as a function of radial position at $z = 0$ is plotted in figure 4.7(a). The fact that the field is at a minimum on axis here means that low-field-seeking atoms feel a force pushing them towards the axis of the coil, producing a focusing effect on the beam of particles. This allows the coilgun to act as a waveguide, and can provide a degree of transverse stability to the slowed bunch of atoms in the coilgun.

Though the magnitude of the field is radially minimal on axis at the center of the coil, this is not true at all locations. Examining the radial dependence of

the magnitude of the field some distance along the axis from the center, the field is eventually a maximum along the axis, instead of a minimum. This aspect of the field is illustrated in figure 4.7(b). This means that while low-field-seekers are focused at the center of the coil, outside the coil the magnetic field de-focuses low-field-seekers. The transverse stability of the particles in the slowed bunch depends on the interplay between these two effects.

For typical coils, the focusing effect in the center of the coil is much stronger than the de-focusing effect outside the coil. However, from the discussion in section 4.3.1, there is a benefit to turning the coils off before particles reach the center of the coil. The degree to which the atoms are focused instead of de-focused depends on their location in the coil when the field is switched off. While the longitudinal stability of the slowed packet is greatest when the coils are switched when particles are far from the center of the coil, this provides the least transverse stability, and may even lead to a de-focusing effect on particles being slowed. However, at the cost of additional laboratory space, it may be beneficial not to switch off some coils, enabling the particles to feel the full focusing effects of these coils and providing a counter to the de-focssing effects of switching far from the coil. This concept of transverse stability also shows up in the pulsed electric field decelerator [79, 80], and is further discussed for the coilgun in [25, 81]. The concept of using strong magnetic fields to focus atoms without slowing them is discussed in [82].

4.4 Proof-of-Principle Experiment: Slowing Metastable Neon

The goal of the initial experiment was to show that the concept presented above works, and that measurable slowing of a beam can be achieved over a short distance. To accomplish this, an apparatus consisting of 20 coils was produced, and this apparatus was inserted into the same vacuum chamber that houses the rotor,

described in section 3.3.1. While two of the coils failed in the course of the experiment, the remaining 18 coils are able to slow a supersonic beam of metastable neon from 461 m/s to 403 m/s. Importantly, the slowed beam is separated in time from the initial distribution, allowing it to be easily and clearly observed.

4.4.1 Beam Creation

As noted above, the species slowed in this experiment is metastable neon. In its ground state electronic configuration, $1s^2 2s^2 2p^6$, neon has a closed shell structure and $J = 0$, making it insensitive to magnetic fields (except on the order of the nuclear magnetic moment which is ≈ 2000 times smaller than electronic magnetic moments). However, neon can be excited to the $1s^2 2s^2 2p^5 3s^1$ configuration. The $^3P_2$ metastable state has a lifetime of 14.7 s [83], which is much longer than the few millisecond duration of these experiments.

In the $^3P_2$ state, both $m_J = 1, 2$ are low-field-seeking states, and since the Zeeman shift is directly proportional to m_J, it is the $m_J = 2$ sublevel which is targeted and slowed. The $^3P_2$ state has 5 magnetic sublevels, and atoms in the supersonic beam should be evenly distributed among these sublevels, at most 20% of metastable atoms in the beam are in the targeted state. There is also a $^3P_0$ metastable state in neon, which probably also makes up some percentage of the initial metastable beam. In the $^3P_2$ state, the Landé g-factor is 1.50 [84], which means that the magnetic moment of the atom is $3\mu_B$. The fine structure splitting of the states is around .05 eV and the magnetic field required to produce a shift of the same energy is around 300 T. Hence, the fields produced in these experiments are far lower than this, the Zeeman shift is the appropriate regime for any calculations. To give a sense of scale, the Zeeman shift at 1 T is .174 meV, and the kinetic energy of a metastable neon atom at 500 m/s is around 26 meV. Thus, fields of order several Tesla per stage are needed

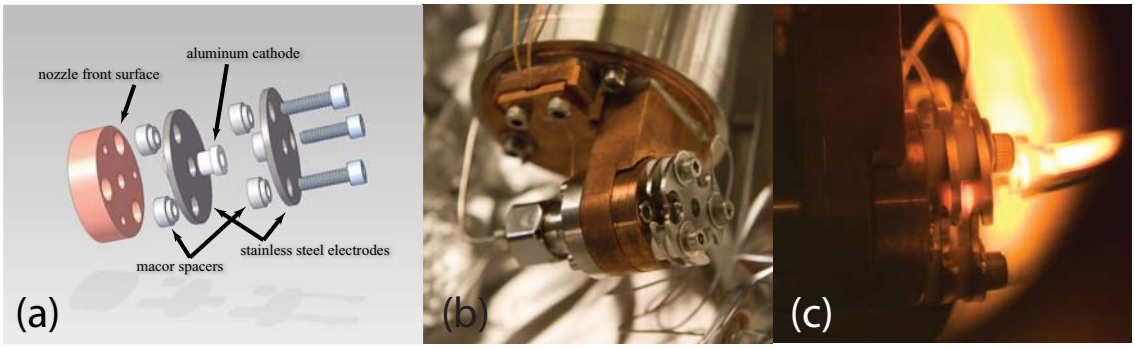

Figure 4.8: The discharge and discharge apparatus mounted to the valve are shown in these images. A CAD image of the discharge apparatus is shown in (a), and a photo of the assembled discharge apparatus mounted to the valve is shown in (b). A photo of the discharge is shown in (c). The bright source on the right side of this image is the filament, while the discharge is the red/purple glow between the discharge plates.

to significantly change the energy of the beam.

A pulsed DC discharge between a stainless steel plate and a hollow aluminum electrode is used to excite the neon atoms in the supersonic beam to the metastable state. The discharge apparatus is mounted to the front of the supersonic nozzle as seen in figure 4.8. The distance between the exit of the nozzle and the discharge region is 5 mm, which means that the beam is still under-expanded and quite dense, which improves the discharge efficiency. The discharge voltage is typically set at 1 kV and is pulsed for 3 μs. The pulse is produced using a Behlke HTS 31-03-GSM high voltage push-pull MOSFET switch, which gates a .112 μF high voltage capacitor bank. The current from the capacitor bank to the discharge is limited to 1 mA by a 1 $k\Omega$ resistor in series with the discharge. A filament biased at -100 V is mounted 2 cm from the discharge and 1 cm off the beam axis, which serves as an electron source that helps initiate and stabilize the discharge.

The supersonic beam of neon in which the discharge is produced is created by the Even-Lavie supersonic nozzle described in chapter 2. Since neon is a noble gas,

it does not need to be seeded into another carrier gas; a pure neon beam is used. To reduce the initial velocity, the nozzle is cooled to 77 K. A 50 psi backing pressure is used in these experiments.

4.4.2 18 Stage Coil Design

As noted above, to significantly slow a beam of metastable neon from $\approx$ 500m/s fields of several Tesla per coil are desired. To produce these fields, a pulsed electromagnetic coil is constructed which is modeled on the coil that drives the plunger in the Even-Lavie nozzle. The coil consists of 30 windings (6 layers of 5 turns) of .5 mm diameter Kapton insulated copper magnet wire. The wire is wound around a hollow Vespel cylinder with 50 μm wall thickness and 3 mm outer diameter. The coil is located between two Permendur discs (Permendur 49, ASTM-A-801 Type 1, from Carpenter Technology), one 3 mm thick, and the other 2.5 mm thick, both 10 mm in diameter, and is surrounded by a 3 mm thick soft magnetic steel shell. Permendur is a high magnetic saturation material with a saturation value of 2.3 T, which confines the field lines and increases the peak field inside the coil. The reason that one disc is made thinner than the other is that the bottom of the soft steel shell is .5 mm thick, and so there is 3mm of magnetic steel on each side of the coil. Additionally, a Kapton disc is inserted in between the coil and each Permendur disc, which helps to protect the insulation of the wire. The entire coil is held together using Epo-Tek H77 epoxy, which is a UHV compatible thermally conductive epoxy. A CAD image of the coil is shown in figure 4.9.

The peak current in the coil is measured to be 400 A. Because the coil is surrounded by non-linear ferromagnetic materials, the field profile of the coil is very difficult to calculate analytically. As such, the field is instead calculated numerically using finite element analysis using the program COMSOL Multiphysics. This allows

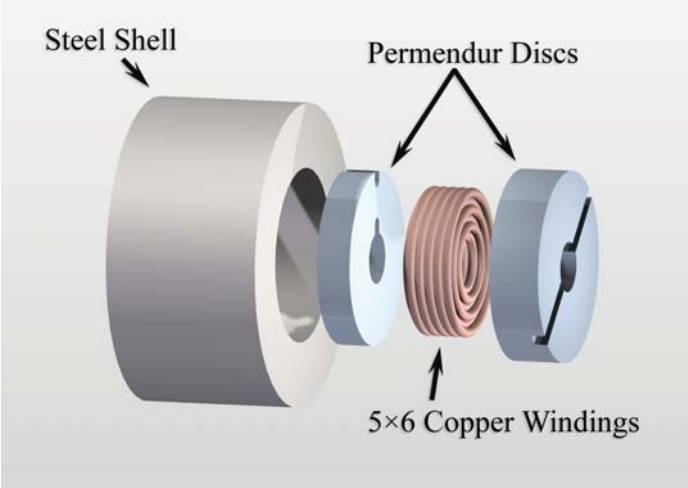

Figure 4.9: A CAD overview of the slowing coils used in the proof-of-principle experiment. The coils are 3 mm bore coils of 5×6 windings, surrounded by Permendur discs and enclosed in a magnetic steel shell.

the full non-linearity of the magnetic materials to be accounted for. Finite element analysis calculations of the field along the axis of the coil, and the radial field in the center of the coil are shown in figure 4.10. The peak field in the center of the coil is 3.6 T.

While great care was taken when winding the coils, during testing in vacuum two of the coils failed when the coil wire came into electrical contact with the permendur shell. These coils were numbers 17 and 19, and they were not used for the slowing. The Kapton insulation of the magnet wire was probably scratched during the winding process, and a test procedure was developed to prevent future coils from exhibiting this same behavior as described in section 4.5.1.

4.4.3 18 Stage Electronics and Coil Switching

The coil wire is only .5 mm in diameter and carries 400 A, so the current must be pulsed or the wires would overheat and melt. A pulse length of 80 μs is long enough

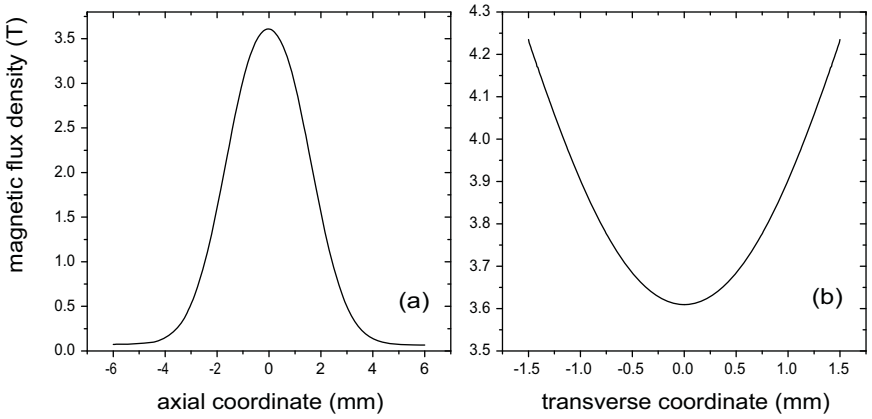

Figure 4.10: Finite element calculations of the magnitude of the field in the proof-of-principle experiment coils are shown. The field calculated along the axis of the coil is shown in (a), while the transverse field magnitude at the center of the coil is shown in (b).

that the atoms are still $3-4\,\mathrm{cm}$ from the coil when the current is turned on, and they are still well outside of the extent of the field when the current reaches its maximum. Thus the atoms will see a steady state field when they enter the coil, and the time variation in the field experienced by the atoms is only due to their motion and the deliberate switch-off of the field.

The current is produced by allowing a capacitor to discharge across the coil, with two solid state switches controlling the current. The switches used are an integrated gate bipolar transistor (IGBT), and a thyristor. The IGBT (Powerex CM400DY-12NF) is a fast switch, with well-defined switching characteristics, while the thyristor (Littelfuse S6040R) is slow, with relatively poorly defined switching profiles. Using both switches allows the fast switching but much more expensive IGBT to be used to switch multiple coil channels, while the thyristor is used to isolate the channels from each other.

A schematic of the switching circuit is shown in figure 4.11. The parts within the dashed box exist for each coil channel, while those outside the box are shared

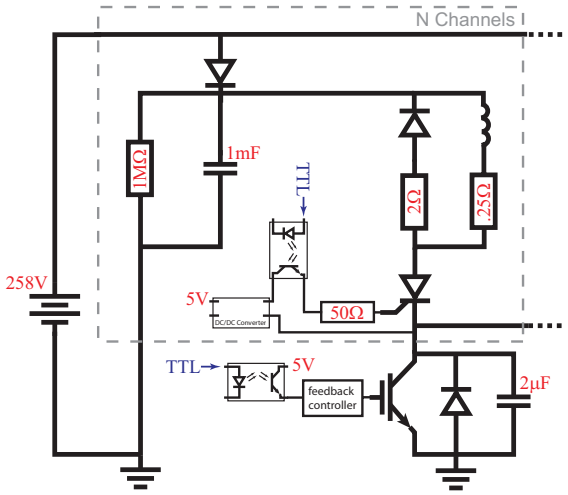

Figure 4.11: A schematic of the coil driver circuit used to drive the 18 stage coilgun coils. The coils are switched by two solid state switches, a thyristor and an IGBT. The elements inside the dashed box are unique to each coil, while those outside the box are shared by multiple coil channels.

among multiple channels. There are 8 total IGBTs, so each IGBT is responsible for switching 2-3 coils. Each coil has a 1 mF capacitor, charged to 258 V by a set of five 50 V switching power supplies connected in series, which is capable of 2.1 A steady state. A 1 MΩ resistor is placed across the capacitor to allow it to discharge even if power is disconnected. This allows the circuits to be safely examined after some time, even if there is a fault that does not allow the capacitors to be discharged across the coils.

The thyristor is controlled by the current flowing from the gate to the cathode, and requires 20 mA to close. Producing this current requires a positive voltage with respect to the cathode, and a DC/DC converter is used to create a 5 V bias between the gate and cathode. The bias is switched by an opto-coupler that protects the digital side of the circuit by physically disconnecting it from the high voltage side of the circuit. The thyristor requires $\approx 40\ \mu$s to effectively close after the current

between the gate and cathode, and the current from the anode to the cathode have both stopped. An additional 20 μs is added to this time to ensure that the thyristor is truly closed. For 80 μs coil pulses, an IGBT can be fired once every 140 μs. With 8 IGBTs, and a coil spacing of 1.41 cm, as described in section 4.4.4, this permits beams of up to 800 m/s to be slowed. Faster beams would require a larger number of IGBTs. The IGBT itself is controlled by a gate driver and DC/DC converter package (Powerex VLA-500), and mounted on a dedicated board (Powerex BG2A). When the IGBT is closed and current is allowed to flow, the gate driver lifts the voltage of the gate to 15 V compared to the emitter of the IGBT.

The timing of the coil pulses is a critical part of the magnetic slowing experiment. Each thyristor requires a 5 V signal to open, as does each IGBT. Thus the timing of the pulses going to the 18 thyristors, 8 IGBTs, the Even-Lavie nozzle, the discharge driver, and the data acquisition system must all be carefully controlled. This is done using a field programmable gate array (FPGA). The FPGA board (Xilinx DS-BD-3SxLC-PQ208 REV 2) has 140 digital outputs, and an external clock at 10 MHz providing 100 ns timing resolution. The outputs of the FPGA are buffered and distributed using a custom board, and all outputs from this board are passed through opto-couplers to isolate the digital electronics from the coil drive circuitry.

When the two TTL pulses arrive at the IGBT and the thyristor, each open and allow the capacitor to discharge. The measured peak current across the .25 Ω resistor in series with the coil is 400 A and it reaches this value after 35 μs. The current remains constant for the remaining 45 μs of the pulse. An oscilloscope trace of the voltage across the .25 Ω resistor is shown in figure 4.12. When the IGBT opens the circuit, the current in the IGBT stops very quickly (a few hundred nanoseconds), however the current in the coil does not shut off immediately, but instead drops linearly to 0 A over 7 μs. This continued current produces a large voltage spike on the collector of

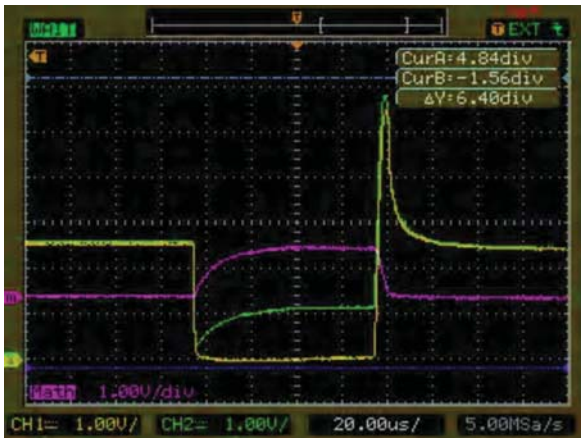

Figure 4.12: An oscilloscope trace measuring the voltage across the .25 Ω resistor in series with the slowing coil. The probes are 1:100, since the voltages would be too high for the oscilloscope to measure with 1:1 probes. The green trace is the voltage on the capacitor side of the resistor, while the yellow trace is the IGBT side. The purple trace shows the voltage drop across the resistor. Figure Courtesy Christian Parthey.

the IGBT of 550 V, due to the back EMF in the coil trying to maintain the current in the coil even though the connection to ground no longer exists. This voltage spike would be even higher, but a 2 μF snubber capacitor across the IGBT absorbs much of the current, as well as limiting voltage and current oscillations. Additionally, the voltage on the collector is reduced by a freewheel diode and 2 Ω resistor in parallel with the coil, permitting the current to briefly circulate while being damped by the resistor. These measures are necessary because the blocking voltage of the IGBT and thyristors is limited to 600 V.

While the measured current drops quickly, Lenz's law indicates that eddy currents may be formed in the permendur surrounding the coil, which could affect the switching of the magnetic field. For this reason, the field is measured directly using a pickup coil. A single turn coil is inserted into the bore of the slowing coil, and the current in the coil is monitored. Integrating this current gives a temporal profile of the field produced by the coil, as shown in figure 4.13. As seen here, though the

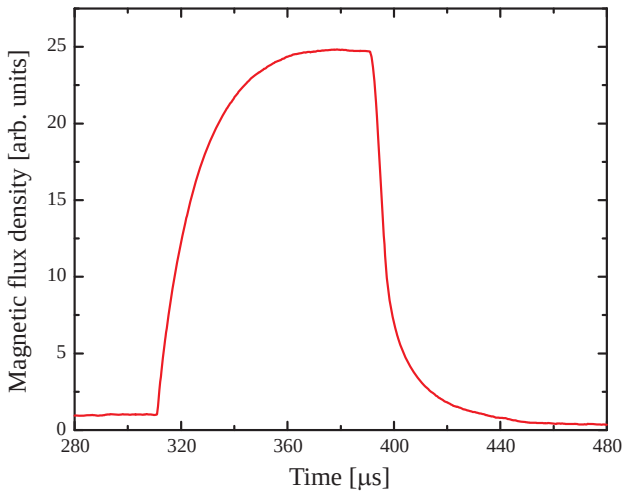

Figure 4.13: A temporal profile of the magnetic field in the center of the slowing coil. The trace is a numerical integration of the current induced in a pickup coil placed in the center of the slowing coil.

current in the coil reached its peak value after 35 μs, the field does not reach its peak for another $5 - 10$ μs. Also, while the current shuts off within 7 μs, the field drops linearly over this time to 35% of its peak but decays exponentially from there with a time constant of 11 μs.

4.4.4 18 Stage Apparatus and Vacuum Chamber

To make the coilgun work correctly, the coils must be held in alignment in vacuum. The coils are held in copper discs by vented screws. The copper discs have three holes drilled into them, and this is how the entire assembly is held together and aligned. The discs are slid onto three stainless steel rods, with stainless steel spacers between each copper disc. This produces a center to center spacing between each coil of 1.41 cm. Instead of spacers, a stainless steel post is used on two of the rods between coils 7 and 8, and between coils 13 and 14. These posts are mounted in a stainless steel block that is secured by set screws in the vacuum chamber. The

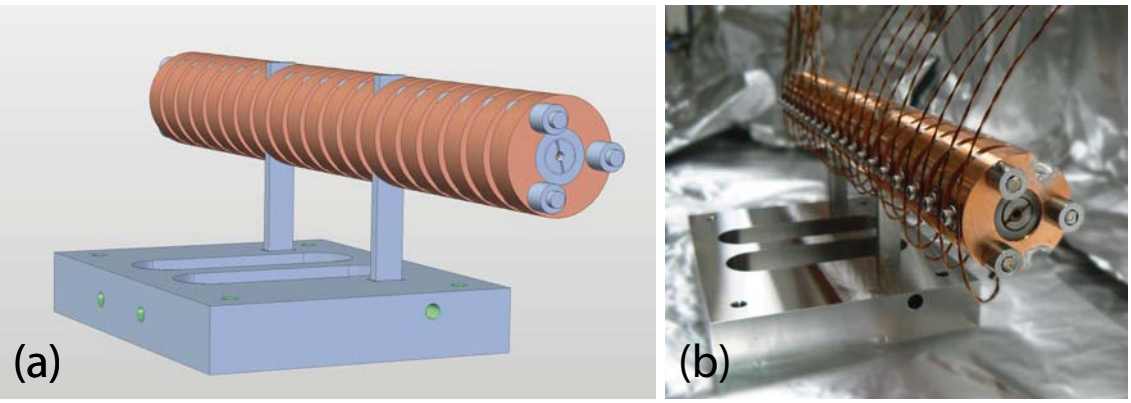

Figure 4.14: A CAD image of the 18 stage coilgun (a) and a picture of the apparatus (b). There are 20 stages pictured here, but two of the coils became inoperable after the apparatus was placed in vacuum. Figure Courtesy Christian Parthey.

copper discs are also water cooled by a copper water line that is indium soldered to the discs, allowing some of the heat to be removed from the coils. A CAD image and photo of this setup is shown in figure 4.14.

Since this coilgun was intended simply as a proof-of-principle experiment, a separate vacuum chamber was not built to house the coils. They were instead placed into the existing rotor chamber. The supersonic beam chamber is un-modified from the chamber described in section 3.3.1 (except that the discharge plates are mounted to the nozzle), and the rotor chamber is essentially unchanged. The 5 mm skimmer is mounted 30 cm from the nozzle, as it was in the paddle experiment. The center of the first coil is located in the rotor chamber 59 cm from the nozzle. The coilgun mounted inside the rotor chamber can be seen in figure 4.15.

The one significant change from the rotor vacuum chamber is in the detection system. The metastable atoms are detected using a micro-channel plate (MCP) instead of an RGA. The MCP (El Mul Technologies) detects the metastable neon atoms because the ionization energy of the coating on the MCP is lower than the excitation energy of the neon causing a free electron to be created when a metastable

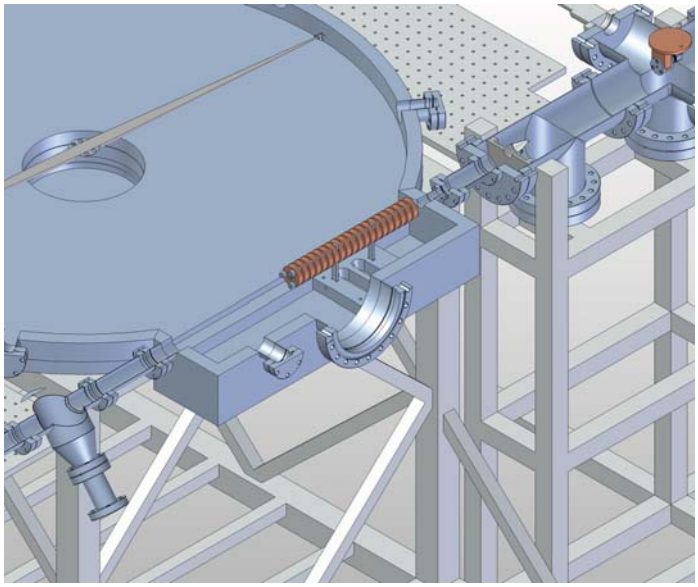

Figure 4.15: A CAD image of the 18 stage coilgun placed in the rotor chamber. The center of the first coil is 59 cm from the nozzle, and the coilgun is 26.8 cm from the center of the first coil to the center of the last coil. Figure Courtesy Christian Parthey.

neon atom hits the surface (known as Penning ionization). The MCP is an array of electron multiplier tubes that amplify the electron current produced by the impact of the metastables. A bias across the MCP sets the amplification gain: typical values are 1.6 kV for the bias resulting in a gain of 10^6. The circuit which biases the MCP is shown in figure 4.16, and the MCP used has an 18 mm diameter active area.

The position of the MCP is different from the position of the direct beam RGA from the rotor experiment. While the RGA was 1.42 m from the nozzle, it is desirable to have the MCP be even farther away. The coilgun only slows a small fraction of the beam, and slows that fraction by only a few tens of meters per second. Having a long distance between the exit of the coilgun and the MCP increases the separation between the slowed and unslowed parts of the beam. Additionally, the MCP is mounted on a 5.08 cm translation stage that allows it to be moved closer to or farther from the coilgun. By observing the difference in time-of-flight of the

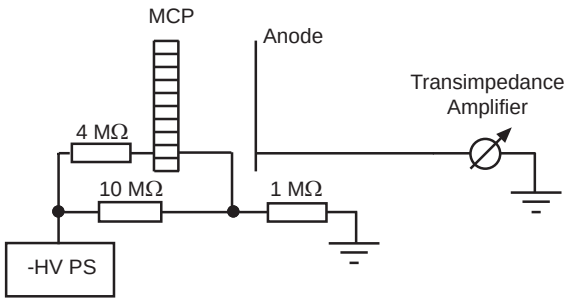

Figure 4.16: A Schematic of the MCP bias circuit. The voltage across the MCP sets the gain, typically set at around 1.6 kV. Figure Courtesy Christian Parthey.

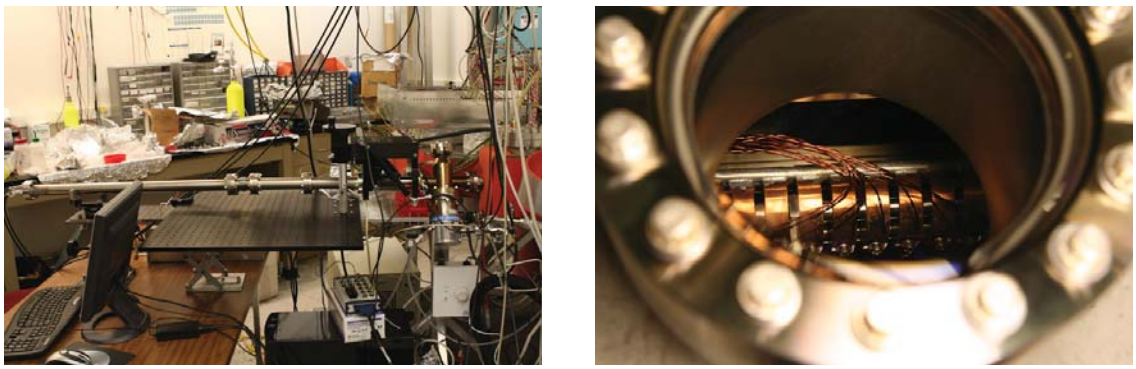

Figure 4.17: At left is a photo of the MCP extension tube, mounted on a translation stage. The increased distance between the coilgun and the MCP is needed to resolve the slowed beam from the main peak. At right is a photo of the coilgun in the vacuum chamber, seen through a viewport.

different features of the signal, the velocity of the atoms that produced that feature can be measured. The MCP is located 2.54 m (2.59 m) from the nozzle in the retracted (extended) position and 1.68 m (1.73 m) from the end of the coilgun. A photo of the MCP mounting extension tube and translation stage, along with a photo of the coilgun in the vacuum chamber seen through a viewport are shown in figure 4.17.

Becuase the coils are in vacuum, the water cooling and electrical connections to the coils are accomplished using feedthroughs. For the water lines, 1/4 inch Swage is used to make the in-vacuum connections. The electrical feedthroughs are 50 pin

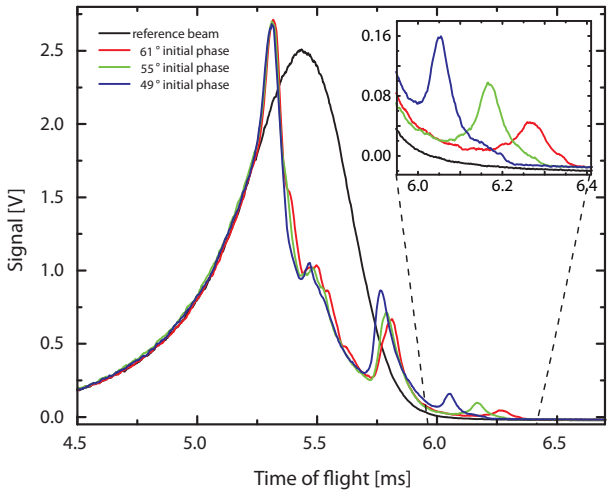

Figure 4.18: A plot of the time of flight results of the metastable neon beam, varying the phase in the 18 stage coilgun. Each curve is an average of 10 shots with a current in the coils of 400 A, and the reference beam is the signal measured without pulsing the coils. A greater phase angle leads to more slowing, but a smaller region of phase stability, leading to fewer slowed atoms.

D-Sub connector feedthroughs from Accu-Glass.

4.4.5 18 Stage Data and Results

As noted above, during testing of the coils in vacuum, two of the coils (17 and 19) failed as the coil wire became shorted to the Permendur shell. Though the coilgun was designed to use 20 coils, in the end only 18 are used for the slowing results presented here. The data show the slowing of the beam from an initial velocity of 461 m/s. This is the target velocity chosen from a beam with a center velocity of 470 ± 1.8 m/s and a FWHM velocity of 43 m/s. From this target velocity, the path of a synchronous atom is numerically integrated for a given phase, and the start time of the pulse sequence is empirically scanned to optimize the slowed signal by matching the target velocity simulated with the correct velocity group of the incoming beam.

The output velocity of the coilgun can be controlled by varying the phase used

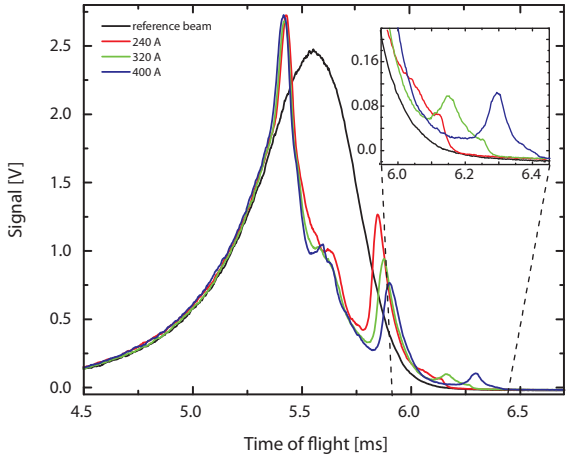

Figure 4.19: A plot of the time of flight results of the metastable neon beam, varying the current in the coils of the 18 stage coilgun. Each curve is an average of 10 shots and each timing file generated had an initial phase of 55°. The larger currents lead to greater magnetic fields, and thus more slowing.

in the coils, and the results of this investigation are presented in figure 4.18. Rather than set a specific constant phase angle for the entire slower however, the numerical simulation which creates the timing file is set to switch the coil off at a certain constant time before the atom would arrive at the center of the coil. Since the atoms are being slightly slowed, this is equivalent an almost linear variation of the phase angle for each coil. The phase angles used are $49° - 51°$, $55° - 58°$, and $61° - 64°$, which yielded final speeds of $431.2 \pm 6.0\,\mathrm{m/s}$, $409.3 \pm 9.1\,\mathrm{m/s}$, and $403 \pm 16\,\mathrm{m/s}$ respectively. As can be seen in the figure, greater phase angles lead to more slowing, but also to a smaller phase stable region, reducing the number of slowed atoms.

Varying the current in the coils allows the effects of magnetic field strength to be examined. The currents used are 400, 320, and 240 A, and the results can be seen in figure 4.19. The current was not reduced below 240 A because at lower currents the magnetic fields were no longer strong enough to separate the slowed peak from the main beam. These currents correspond to peak magnetic fields in the coils of 3.6,

3.0, and 2.4 A respectively. The initial phase used in generating the timing files for these results is 55°. The final velocities of the slowed beams for these magnetic fields are 409.3 ± 9.1 m/s for 400 A and 416 ± 22 m/s for 320 A. The speed of the beam created by operating the slower at 240 A is not determined as the slowed peak could not be sufficiently resolved.

4.5 The Atomic Coilgun: Stopping Metastable Neon

Having successfully slowed a beam of metastable neon using the proof-of-principle apparatus, the next step was to show the scalability of the coilgun method. Additionally, a goal of the next apparatus was to show that the coilgun could slow the beam enough to enable trapping of the slow packet. While this experiment did not attempt to trap the slowed beam, it was important to show that the coilgun method could produce beams slow and cold enough that trapping the beam at the exit of the coilgun is a realistic application of the coilgun.

To this end, a 64 stage coilgun was constructed using the lessons learned from the proof-of-principle experiment, and incorporating improvements in both the coil design and the electronics. A dedicated vacuum system was built to house the 64 stage coilgun, and the signal processing was optimized. Using the 64 stage coilgun with metastable neon allowed full control of the velocity of the slowed beam. Neon is slowed from an initial velocity of 446.5 ± 2.5 m/s to velocities as low as 55.8 ± 4.7 m/s, with this lower limit set by the transverse expansion of the beam from the end of the slower to the detector. The coilgun effectively stops the beam as more than 99% of the initial kinetic energy is removed and the remaining velocity is only kept to allow the atoms to propagate to the detector.

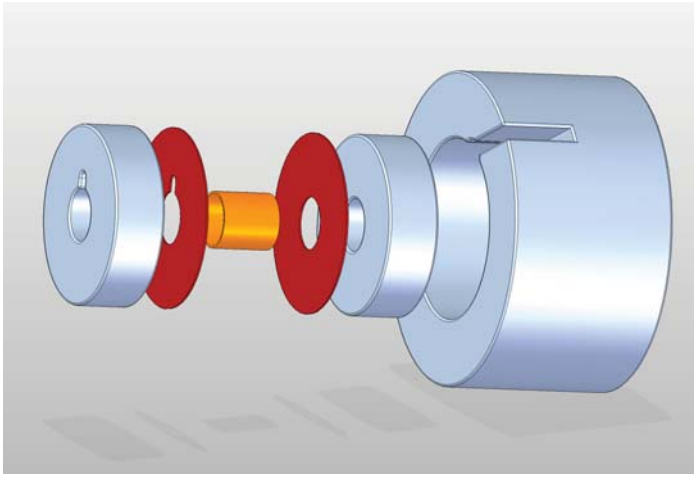

Figure 4.20: An exploded view of the 64 stage coilgun coil. Each Permendur disc is 3 mm thick, and 10.2 mm in diameter. The wire is wound around the 3 mm O.D. center Vespel cylinder. Figure Courtesy Christian Parthey.

4.5.1 64 Stage Coil Design

The design of the coils used in the 64 stage coilgun is slightly modified from those used in the proof-of-principle experiment and described in section 4.4.2. The main change is that rather than using a magnetic steel cup, with one Permendur end cap slightly thinner than the other, the end caps are made with the same thickness, and the magnetic steel shell is now a cylinder.

The full specifications for the new coil are as follows. The coil consists of 5×6 windings of 500 μm diameter kapton insulated copper wire wound around a 3 mm bore. The coil is wound around a Vespel cylinder, which has a wall thickness of 50 μm and an outer diameter of 3 mm. Each solenoid is sandwiched between two Permendur discs with 10.2 mm diameter and 3 mm thickness. 150 μm thick Kapton washers are placed between the wire and the permendur discs to protect the coils and to keep the insulation from scratching during the winding process. The coils are surrounded by a magnetic steel cylinder which is 9.6 mm long, with a 10.2 mm inner

diameter and 17.8 mm outer diameter. The coil is held together with Epo-Tek H77 UHV compatible, thermally conductive, and electrically insulating epoxy. A blow up of the components which surround the coil is presented in figure 4.20.

To wind the coils, the Permendur discs, Kapton washers, and Vespel cylinder are placed on a winding tool. The coil is then wound in the gap between the Kapton washers, around the Vespel cylinder. Shrink wrap Teflon is placed around the wire where it enters and exits the coil to provide further protect at these locations, which were found to be particularly prone to scratching. Once the coil is wound, it is removed from the winding tool, and is placed onto a 3 mm diameter Teflon rod. The rod is anchored in a Teflon mold which includes a groove that allows the wires to exit the coil. The top Permendur disc and Kapton washer are removed, and the magnetic steel cylinder shell is placed over the coil. At this point the epoxy is applied to the coil windings using a small Teflon spatula. The epoxy is quite viscous at room temperature, but heating the coil to 50° C with a heat gun reduces the viscosity enough that the epoxy flows around the coil windings. The 3 mm diameter Teflon rod the coil sits on keeps the epoxy from flowing into the coil bore. The top Kapton washer and Permendur disc are then put in place and the entire assembly is clamped in place with another Teflon mold held in position with a C-clamp. The coil is then baked for 3 hours at 150° C to cure the epoxy. Once the coils have cooled, the clamps are removed and the excess epoxy is removed using pliers. Lastly, any remaining epoxy is removed by sanding the outside of the coil by hand. The coil assembly sequence is shown in figure 4.21.

After the coils are assembled, they are tested for short circuits and for scratched insulation on the wires. This is done by placing the assembled coil in a beaker of methanol. A multimeter is used to measure the resistance between the coil wire and the methanol. Since the methanol will flow around the coil, and around the steel and

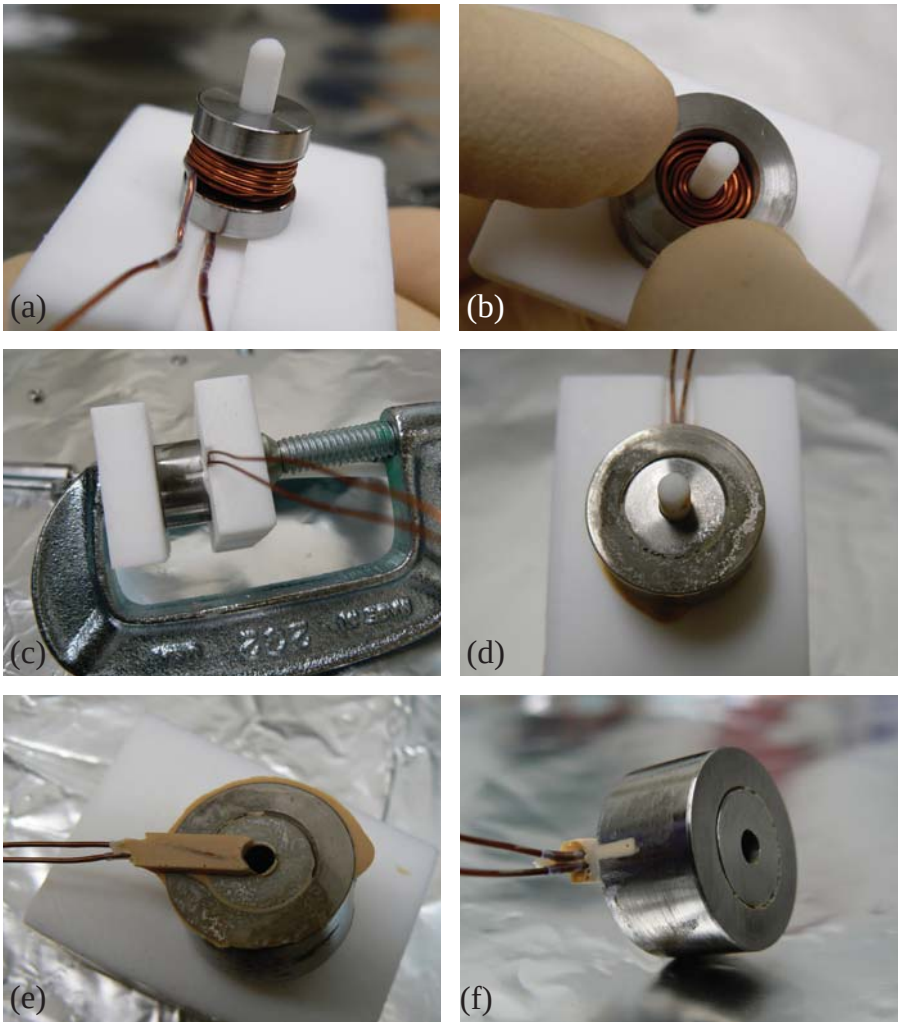

Figure 4.21: A pictorial overview of the coil assembly procedure for the 64 stage coilgun. The coil wound around the vespel cylinder between the kapton washers and permendur discs is shown in (a). The top kapton washer and disc are then removed and the magnetic steel shell is placed around the coils, seen in (b), so that the coil can be epoxied. The top kapton and permendur are replaced, and the entire assembly is clamped in place for baking as shown in (c). The coil after the epoxy has been cured is shown in (d) and (e), and an image of the coil after all excess epoxy has been removed is shown in (f). Figure Courtesy Christian Parthey.

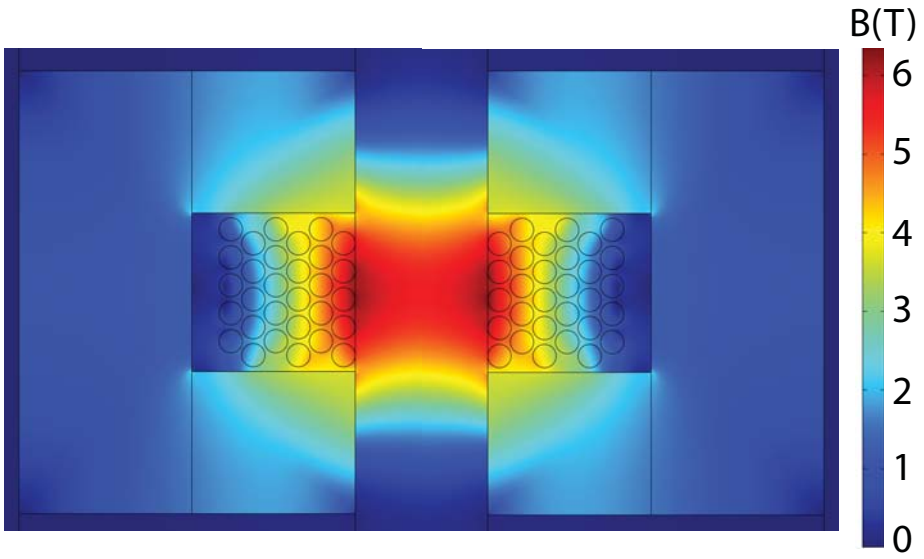

Figure 4.22: A finite element calculation of the field in one of the 64 stage coilgun coils at 750 A. The peak field on axis is 5.56 T.

Permendur enclosure of the coil, any breaks in the Kapton insulation of the wires which is not covered by epoxy will be detected by this method. Any coil with a measurable resistance is assumed to be shorted and is rejected. While two coils failed in the course of the proof-of-principle experiment, no coils that passed this test have failed to this point.

The spatial profile of the fields produced by the coil would be extremely difficult to determine analytically, and so they are instead computed numerically using finite element analysis. From the coil characterization discussed in section 4.5.2 the current in the coil is estimated to be 750 A. At this current, the peak magnetic field on the axis of the coil is calculated to be 5.56 T. A finite element analysis calculation of the field in the coil is shown in figure 4.22.

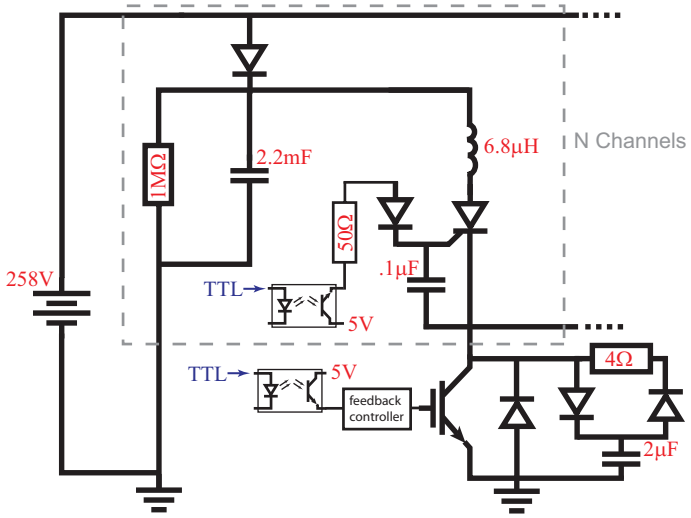

Figure 4.23: A schematic of the coil driver circuit used to drive the 64 stage coilgun coils. Similar to the driver used in the proof-of-principle experiment, the coils are switched by two solid-state switches, a thyristor and an IGBT. The elements inside the dashed box are unique to each coil, while those outside the box are shared by multiple coil channels.

4.5.2 64 Stage Electronics and Coil Switching

In addition to a new coil design, the electronics used to switch the coils are improved from the proof-of-principle experiment. Many aspects of the original circuit remain in place, and some elements are completely unchanged from the circuit used to drive the coils of the 18 stage coilgun. Those elements which are not changed are simply noted as such, and the information on these components can be found in section 4.4.3.

A schematic of the coil driving circuit is shown in figure 4.23 and a photography of the circuit is shown in figure 4.24. The capacitor that discharges across the coils is upgraded to a 2.2 mF capacitor, however the power supply and capacitor discharge resistor remain unchanged. The concept of using thyristors as slow switches to isolate individual coil channels from each other, and to use fast IGBTs to switch the current

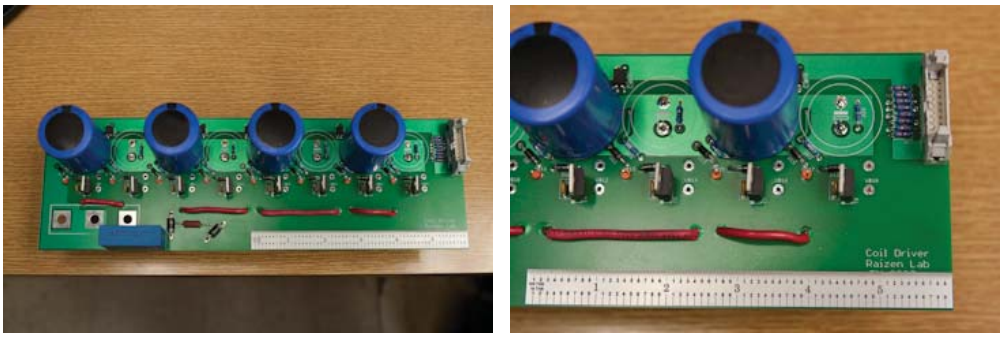

Figure 4.24: Photographs of the 64 stage coilgun coil driver board. At left, the entire board is shown (except for the IGBT which mounts to the three square pads). At right, a zoom-in view provides more detail.

for multiple channels is also unchanged, though many of the individual components of the system are upgraded. The IGBTs are upgraded to a 1.2kV blocking voltage model (Powerex CM200DY-24A) and the thyristor is likewise upgraded to a 1 kV voltage model (Littlefuse SK055R). These changes allow the freewheel diode and resistor to be taken out of the circuit, which decreases the switching time of the coil.

While the IGBT gate drive circuitry is unchanged, the method used to drive the thyristor gate is also modified. The DC/DC converter used to drive the gate was found to fail when subjected to the voltage spikes produced by the faster coil switching at higher current. Rather than replace the DC/DC converter with a different model, the solution implemented is simply to close the IGBT 20 μs before switching the thyristor. Closing the IGBT grounds the cathode of the thyristor, allowing a 5 V power supply, switched by an opto-coupler, to provide the necessary current to the gate.

The change in the thyristor drive circuitry necessitates a change in the snubber capacitor circuitry across the IGBT. While the snubber capacitor is still needed in its initial role to absorb the continuing current from the coil after the IGBT is opened, in the original configuration, the snubber discharges very quickly when the IGBT is

closed. With the new thyristor gate drive system, this has the side effect of closing any thyristor attached to the IGBT, and subsequently all coils fire simultaneously, rather than individually. To rectify this, a resistor is added to the snubber capacitor, along with a diode network, such that the resistance is low when charging the capacitor and it can effectively absorb the current pulse from the coil, but the discharging resistance is raised to $4\,\Omega$, slowing the snubber capacitor discharge. This slower discharge profile of the snubber capacitor does not inadvertently close the thyristors, which enables their continued use as isolating switches between the individual coil channels. Similar to the circuitry used in the proof-of-principle experiment, there are still eight total IGBTs, and the digital timing pulses are sent by the same FPGA board. However, with 64 total coils, each IGBT now switches eight individual coil channels.

The final alteration of the drive circuitry is the loss of the .25 Ω resistor in series with the coil. The resistor limited the current and with the new higher blocking voltage components is not needed in the improved circuit. This does mean that there is no convenient way to directly measure the current in the coil, as there is no resistor in series with the coil whose voltage drop can be measured. The total resistance of the circuit is measured to be $\approx 0.34\,\Omega$, leading to a calculated current of $\approx 750\,\mathrm{A}$, but this measurement is accurate to 10 % at best, limiting the accuracy of the estimations of the field strengths.

While a pickup coil is used to measure the temporal profile of the magnetic field in the proof-of-principle experiment, this is not an good method for determining the overall field strength. The problem is that the pickup coil measures $\frac{\mathrm{d}B}{\mathrm{d}t}$, and so must be integrated to determine the absolute field. Small errors can accumulate and lead to a large uncertainty when determining the absolute field in the coil using this method. What is needed is a direct method of measuring the field, which would be accurate on the timescale of the pulse.

The method chosen for measuring the field is the Faraday effect. Some materials (called Faraday rotators) rotate the plane of polarization of light when a magnetic field in the direction of propagation of the light is present. The angle of rotation of linearly polarized light passing through a Faraday rotator in the presence of a magnetic field in the propagation direction is

$$\psi = \int_{z_1}^{z_2} vB(z)\mathrm{d}z, \tag{4.31}$$

where v is the Verdet constant of the material, B is the field in the direction of propagation, and x_1 and x_2 are the coordinates of material boundaries. For a constant field this simplifies to

$$\psi = vBd, \tag{4.32}$$

where d is the length of the Faraday rotator. Importantly, since this is an optical process, the effect takes place on time scale much shorter than the coil switching time, and has even been resolved in the femtosecond domain [85]. The Faraday rotator used to measure the field in the coilgun coils is a terbium gallium garnet (TGG) crystal, which has a Verdet constant at room temperature of $-134\,\mathrm{rad/T/m}$ for 633 nm wavelength light, where the sign indicates the direction of rotation. For a negative Verdet constant, the rotation will be in the direction opposite the direction of current in a coil around the crystal creating a magnetic field. This is independent of the direction of propagation of the light.

A schematic drawing of the Faraday rotation measurement setup is shown in figure 4.25. Polarized light from a Helium-Neon laser passes through a polarizing beamsplitter cube, which sets the polarization angle of the light. This angle can be further adjusted using a $\lambda/2$ plate, which is the next element in the setup. The light is then focused by an 83 mm focal length lens down to a waist of 40 μm in the center of the coil. The TGG crystal is inserted into the bore of the coil, and care is

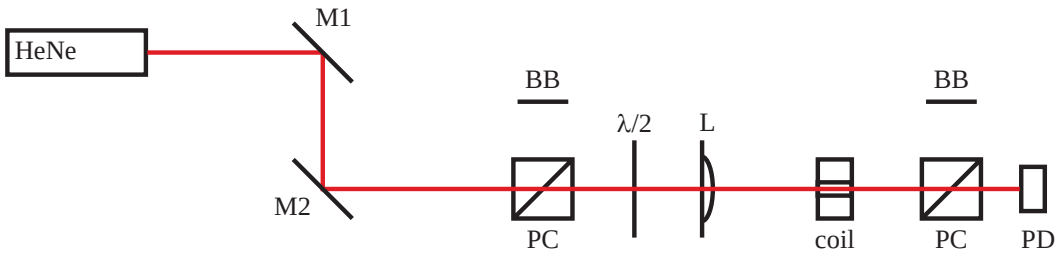

Figure 4.25: A schematic drawing of the Faraday rotation setup. M1 and M2 are beam steering mirrors which are used to align the light to the remainder of the apparatus. The beam passes though a polarizing beamsplitter cube (PC) and a $\lambda/2$ plate before being focused by an 83 mm focal length lens (L) to a 40 μm waist in the center of the coil. The TGG crystal is inserted into the coil, and the light is rotated in proportion to the magnetic field. The light passes though another polarizing cube before being detected on a photodiode (PD). The stray beams coming out of the polarizing cubes are blocked with beam blocks (BB). Figure Courtesy Christian Parthey.

taken to ensure that the crystal is centered axially in the coil (the crystal used has a 3 mm diameter, and thus fills the coil radially). The light then passes though another polarizing cube before being detected on a photodiode.

Because the magnetic field changes the angle of rotation, but the photodiode measures the power of the light hitting its surface, it is necessary to convert between the two. This gives the relation

$$\frac{I}{I_0} = \mathrm{Cos}^2\left[\int_{z_1}^{z_2} vB(z)\mathrm{d}z\right], \tag{4.33}$$

where I is the detected intensity on the photo diode, and I_0 is the maximum intensity corresponding to no light being split off by the second polarizing cube. As the field in the coil varies along the axis of the coil, the measured change in intensity corresponds to the integral of the field along the length of the crystal. The measured field values on the axis of the coil are used to calibrate the current used in the finite element calculation of the spatial field profile.

When actually measuring the field in the coil, crystals of several lengths are

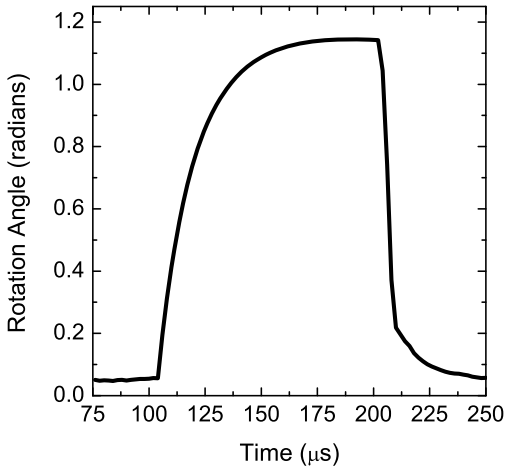

Figure 4.26: The temporal profile of the Faraday rotation in one of the slowing coils used in the 64 stage slower. This is measured using a 1.4 mm TGG crystal, and the maximal rotation observed corresponds to a 5.4 T peak magnetic field. The curve shown is an average of 200 individual measurements.

used to reduce the dependence on the particulars of the integral described above. The temporal profile for the 1.4 mm crystal is shown in figure 4.26. Pulse lengths of 100 μs are used, instead of the 80 μs pulse lengths used in the proof-of-principle experiment. The field follows an initial rise with a time constant of 19 μs. At switch-off, the field first falls linearly to 20 % of its initial value over a 6 μs, and after this the field decays exponentially with a time constant of 17 μs. This exponential decay after the current has been switched off is most likely due to eddy currents in the Permendur surrounding the coils.

Measurements are taken with crystals of 1.40, 2.35, and 5.00 mm length. Table 4.1 shows the calculated peak field for each crystal length. These measurements give a peak field on axis of 5.21 ± 0.20 T, which is extremely close to the field found by finite element calculation assuming a current found by examining the total resistance of the circuit (5.56 T). The discrepancy is probably the result of imperfect windings, where

Table 4.1: Peak magnetic fields measured for different lengths of TGG crystal

crystal thickness	measured peak B-field
1.40 mm	5.43 T
2.35 mm	5.05 T
5.00 mm	5.14 T

the true volume occupied by the wire is greater than in the finite element calculation, leading to a lower real current density.

The numerical calculations of the trajectory of a synchronous atom, as well as Monte-Carlo simulations of the entire beam being slowed by the coilgun, use the finite element calculations for the spatial profile of the fields in the coils. The temporal profile is provided by the measurements taken using Faraday rotation.

4.5.3 64 Stage Apparatus and Vacuum Chamber

The finished coils are mounted on an 89.77 cm long aluminum support, which holds the coils in vacuum. Since the support is machined from a single piece of aluminum, the coils are aligned to each other to better than 50 μm. The center-to-center distance between the coils is 14.1 mm. Each coil is mounted in a semi-circular groove in the support and the coils are clamped in place individually to allow for easy replacement of a single coil. The system is designed this way because of the problems encountered with the proof-of-principle experiment, where the coils were not easily replaceable, and where several coils failed. Improved coil testing techniques have made individual coil replacement unnecessary to this point, but the system has the capability for quick repair should it be required.

Vented screws that clamp the coils individually to the aluminum support are also used in conjunction with a folded piece of Kapton to secure the wires coming out of the coil. This ensures that the wires are securely held a way that reduces the ease

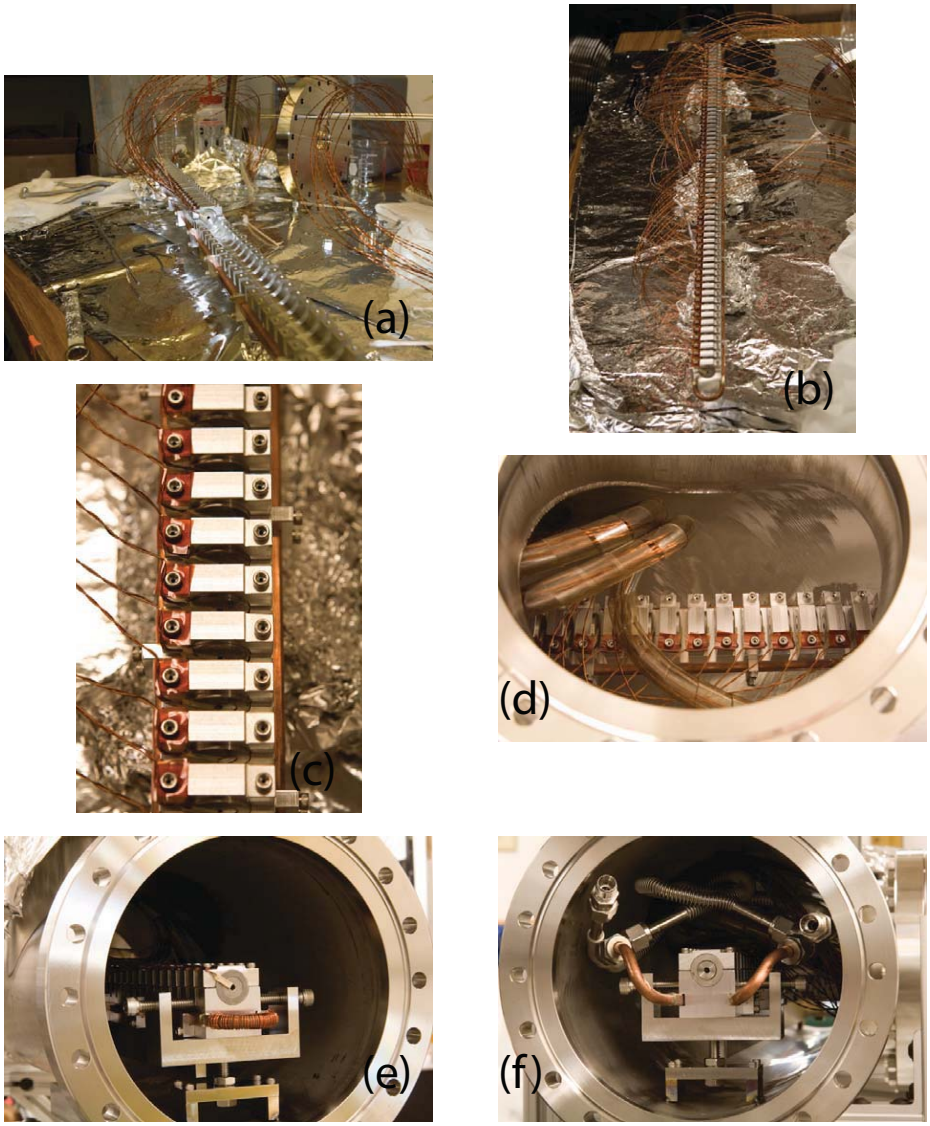

Figure 4.27: Photographs of the 64 stage aluminum coil support structure. The structure with half the coils is shown in (a), and with all coils in (b). A closeup of the coil mounting is shown in (c), and a similar closeup in the chamber illustrating the PEEK wire bundling is shown in (d). Images (e) and (f) show the translation stage mounting of the MCP and valve side of the coilgun respectively.

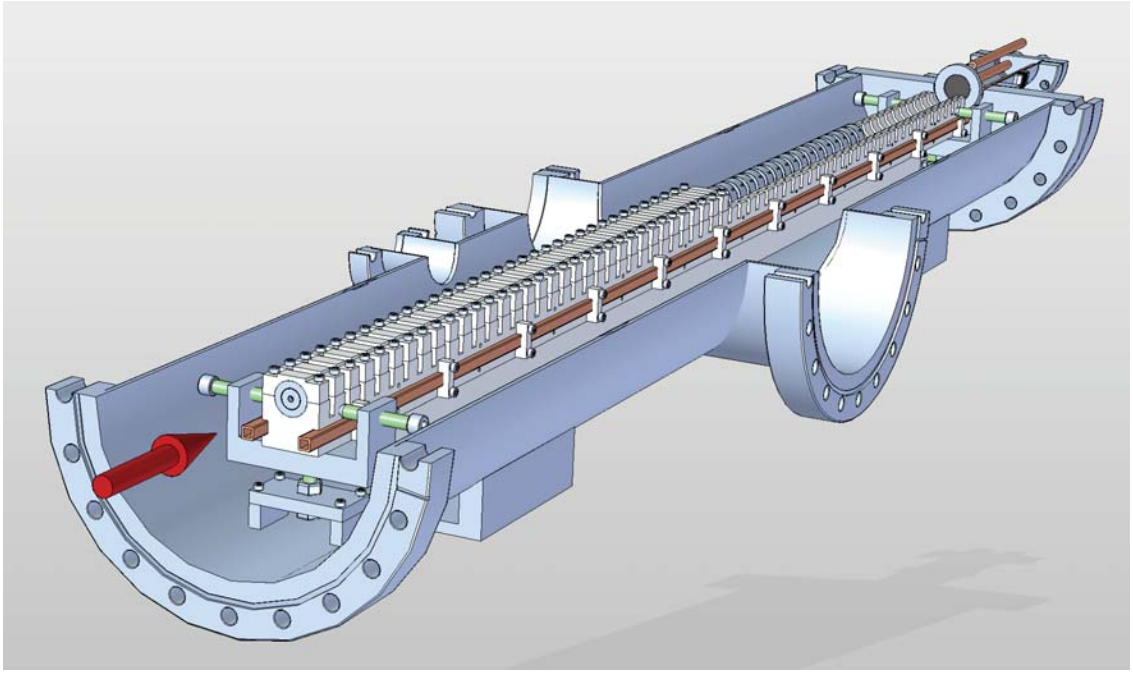

Figure 4.28: A cut-away CAD image of the 64 stage coilgun and vacuum chamber. The red arrow indicates the supersonic beam trajectory. The first 49 of 64 coils are shown, and only the first 32 of these are shown as clamped in place. The 89.77 cm long water cooled aluminum support is held in place by an x-y translation stage at either end of the vacuum chamber. The MCP used to detect the beam is shown at the end of the slower, and is mounted on copper rods so that the MCP is closer to the end of the slower that it would be if mounted directly to the translation stage mounted feedthrough flange.

of scratching their insulation. Since the coils and the aluminum structure they are mounted to are heavy, the support has a stiffening fin along its bottom to ensure that the total deviation of the coils from a straight line is less than 100 μm. In addition, the coil support is water cooled via a 1/4 inch square copper tube that is clamped in place running down either side of the support. This provides a source of cooling for the coils. Pictures of the aluminum coil support can be seen in figure 4.27.

The aluminum support structure is held in the vacuum chamber by two stages which each allow the support to be moved in x and y. The stages are placed at

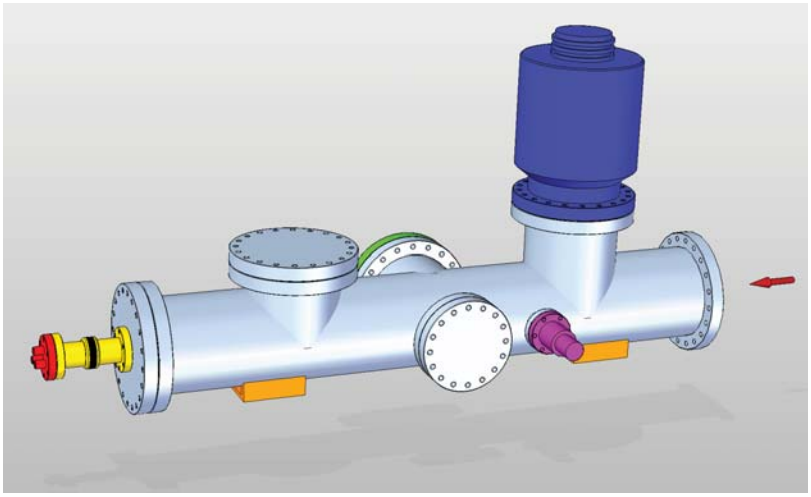

Figure 4.29: A CAD image of the vacuum system of the 64 stage coilgun. The supersonic beam path is indicated by the red arrow. The chamber itself is a 6 inch OD stainless steel tube, configured with 8 inch Conflat flanges on either end. The chamber is pumped by a Varian 500 l/s turbomolecular pump, shown in blue, and the pressure is measured by a cold cathode ion gauge, shown in pink. The MCP feedthrough mounting flange is shown at the end of the coilgun in red, while the edge welded bellows translation stage is shown in yellow (without the aluminum translation stage support structure). The 3 50-pin D-Sub feedthroughs for the coil wires are mounted on an 8 inch blank, shown in green. The chamber itself is mounted to an aluminum support structure using the feet welded to the chamber shown in orange. Figure Courtesy Christian Parthey.

either end of the vacuum system, allowing the support to be adjusted in 4 axes: x, y, pitch, and yaw. There is no need for the support mounting to have z or roll degrees of freedom. The coil support is aligned to the axis of the vacuum chamber using a telescope, and the same telescope is used to align the nozzle and skimmer to the coils. A cut-away CAD image of the support mounted in the vacuum chamber is shown in figure 4.28.

The vacuum system of the coilgun chamber is shown in figure 4.29. The beam creation chamber which produces the supersonic beam of metastable neon atoms is unchanged from the proof-of-principle experiment, as are the nozzle operation pa-

rameters. The coilgun chamber consists of a 6 inch OD cylindrical stainless steel chamber, with 8 inch Conflat flanges on either end. The chamber is pumped by a Varian 551 l/s turbomolecular pump, shown in blue. The pump keeps the chamber at pressures below 1×10^{-8} Torr, as measured by a cold cathode vacuum gauge (pink). As in the proof-of-principle experiment, the beam is detected after slowing by an MCP (feedthrough mount flange shown in red), which is mounted on a 50.8 mm translation stage (shown in yellow, without the translation stage support structure). The center of the first coil of the coilgun is located 555.7 mm from the exit of the nozzle, and ≈ 250 mm from the base of the skimmer. The MCP is located 40.7 mm (91.5 mm) from the center of the of the last coil in the coilgun with the translation stage retracted (extended), and the distance from the nozzle to the MCP is 147.98 cm (153.06 cm). The beam is slowed enough to be resolved from the main peak directly at the end of the coilgun, so it is advantageous to have the detector as close as possible to the exit of the coils to avoid losing flux transversely.

The electrical connections to the coils are made using an 8 inch Conflat blank, with 3 50-pin D-Sub feedthroughs from Accu-Glass welded in. The in-vacuum connectors are made from PEEK to ensure vacuum compatibility. Additionally, the wires running from the feedthroughs to the coils are bundled together and wrapped in a PEEK spiral to protect them. This also makes organizing the 64 pairs of wires somewhat simpler. While the feedthroughs worked without a problem for several tens of thousands of pulses, they did eventually fail catastrophically (electrical failure, not loss of vacuum) without warning. The data shown in the following section is taken with the first set of feedthroughs. A higher isolation voltage set of feedthroughs was purchased (1 kV instead of 600 V), but these too eventually failed after several tens of thousands of shots, again without warning. Moving the coils out of vacuum became a priority for the 3rd generation coilgun described in chapter 5.

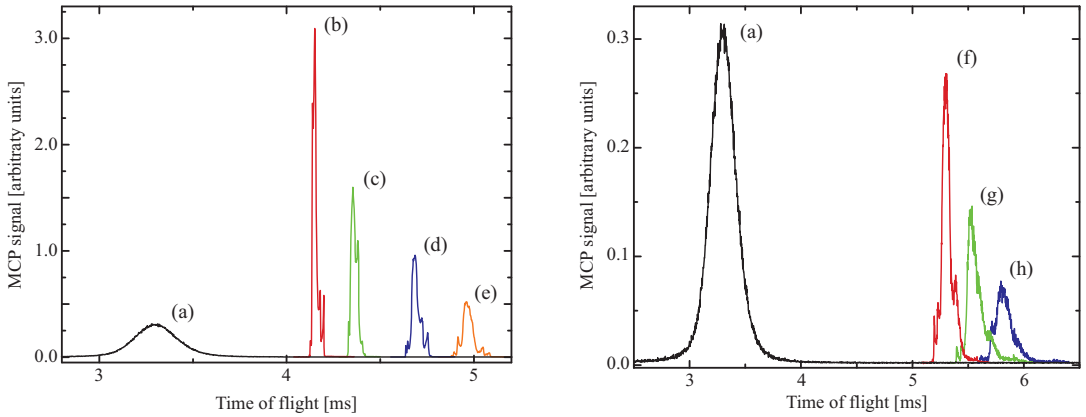

Figure 4.30: Metastable neon slowing data from the 64 stage coilgun. Each time of flight curve is the average of 20 individual measurements. The slowed curves only show the slowed portion of the beam for clarity. The time of flight signal for a beam without firing the coils is shown in (a). The velocities of the beams are (a) 446.5 m/s, (b) 222 m/s, (c) 184.7 m/s, (d) 142.7 m/s, (e) 109.9 m/s, (f) 84.1 m/s, (g) 70.3 m/s, and (h) 55.8 m/s. These results are summarized in table 4.2.

4.5.4 64 Stage Data and Results

Results of using the coilgun to slow metastable neon are now presented. The MCP is used to record time-of-flight signals for a given coil timing configuration. The timing configuration is used to record two sets of data, one with the MCP in the retracted position, and one with the MCP in the extended position. In this manner, the velocity can be measured by observing the time it takes the slowed beam to propagate the extra distance to the MCP. Similarly, the temperature of the beam can be measured by observing how much the beam spreads in this time.

Time-of-flight measurements of the beam are presented in figure 4.30. The initial beam without firing the coils is shown in curve (a), and this beam has a velocity of 446.5 ± 2.5 m/s and a temperature of 525 ± 10 mK. This beam is slowed in the coilgun using phases between 36° and 44°, where the phase used is held constant for each coil in the coilgun (constant phase for the entire slowing sequence). The slowed beams have velocities between 222 ± 11 m/s on the high side, and as slow as

Table 4.2: Final velocities (v_f), simulated final velocities of the synchronous atom (v_s), coil phases, and temperatures (T) of the beams shown in figure 4.30.

	v_f [m/s]	v_s [m/s]	phase [degrees]	T [mK]
(a)	446.5 ± 2.5	–	–	525 ± 10
(b)	222 ± 11	220	36.4°	108 ± 22
(c)	184.7 ± 7.6	185	38.7°	184 ± 39
(d)	142.7 ± 9.1	139	41.0°	117 ± 32
(e)	109.9 ± 5.4	110	42.1°	147 ± 34
(f)	84.1 ± 3.1	83	43.0°	79 ± 20
(g)	70.3 ± 7.4	67	43.4°	92 ± 57
(h)	55.8 ± 4.7	53	43.7°	106 ± 59

55.8 ± 4.7 m/s. The full results of the slowing are summarized in table 4.2.

The time-of-flight profiles of the slowed beams show internal structure, which changes with the coil phase used. This is likely due to the anharmonicity of the magnetic potential of a coil, which causes the atoms to oscillate at different frequencies longitudinally as they are slowed. Hence, the velocity distribution of the slowed peak is not expected uniform, as reflected in the features in the slow beam time-of-flight signal. This phenomenon is also observed and characterized in the pulsed electric field decelerator [78].

It should also be noted how well the predicted velocity of the synchronous atom (from the numerical simulation that generates the timing parameters of the coil for a particular phase) matches with the measured final velocity of the slowed beam. This is likely due to the phase stability of the coilgun, which makes the system robust against small errors in coil position or coil timing. It also indicates that the finite element calculations of the field profiles of the coil are reasonably accurate.

An estimate of the slowed flux can be done assuming that each metastable neon atom hitting the surface creates exactly one free electron at the front of the MCP. After accounting for the gain of the MCP ($\approx 10^5$), and the gain of the current amplifier

($\approx 10^6$) which amplifies the current at the anode of the MCP, the integral under the curve of the detected signal should correspond to the number of particles detected. Since the gains of the MCP and amplifier are known to an order of magnitude at best, the best estimate number of particles slowed per shot is $10^3 - 10^5$, where the beam slowed to 222 m/s has approximately six times more atoms than the beam slowed to 55.8 m/s. Simulations indicate that the 64 stage slower has an efficiency of a few percent for the 222 m/s beam to a few tenths of a percent for the 55.8 m/s beam, where efficiency is defined as the slowed fraction of the metastable beam entering the coilgun. Most of the metastable atoms are not in the $m_J = 2$ magnetic sublevel that is slowed, reducing the efficiency achieved. There is also another metastable state of neon, $^3P_0$, which is probably produced by the discharge, and which is unslowed, again reducing the efficiency.

Rather than keep the phase of the coils constant, a variable phase can be used as well. Figure 4.31 shows how the slowed peak changes when the phase of a coil is tuned in proportion to the velocity of the atoms entering the coil. As the beam is slowed, the coil is switched with the atoms closer to the center of the coil, with phase angle related to the velocity by the relation $\phi_n = \frac{\phi_i v_n}{v_i}$, where ϕ_n is the phase of the nth coil, v_n is the velocity of the synchronous atom entering the nth coil, and ϕ_i and v_i are the initial phase and velocity respectively. Letting the phase be adjusted in this manner leads to a decrease in slowed flux, probably due the atoms being too close to the center of the coil when the coil is switched, reducing the longitudinal phase stability. However, if the phase adjustment is applied, with a cutoff maximum phase allowed, then the slowed beam actually has greater flux. All three slowed peaks in figure 4.31 have the same target velocity at the end of the slower, 136 m/s, however the variable phase peaks have a greater slowed flux.

Simulations indicate that the reason the flux increases for variable phases is

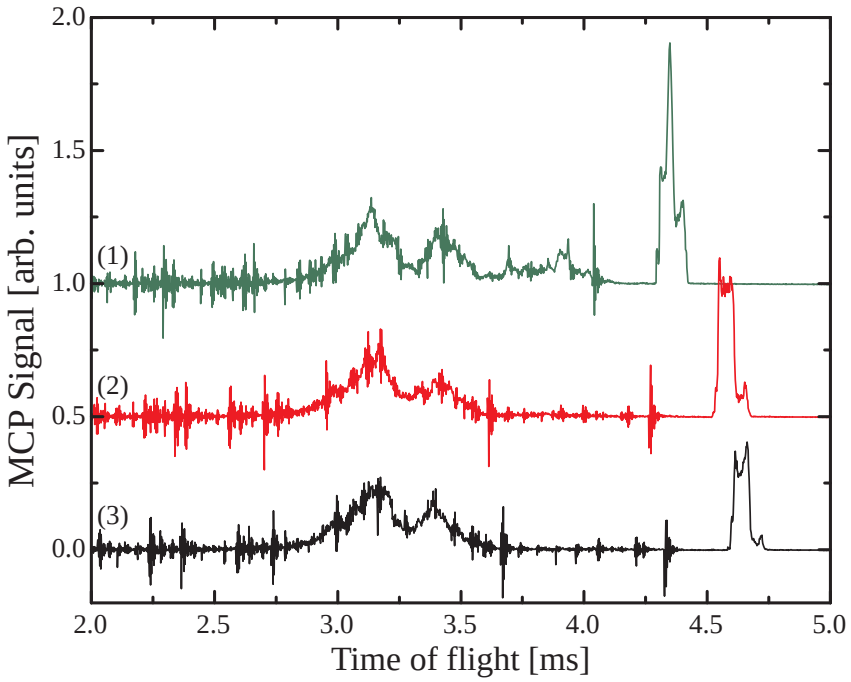

Figure 4.31: Time of flight profiles where a variable phase is used in the coilgun All three coils have the same predicted final velocity, 136 m/s. The phases used are, (1) $\phi_i = 32.9°$ and a cutoff phase $\phi_f = 53.6°$ which is reached after 20 coils, and a measured velocity $v_f = 130$ m/s, (2) $\phi_i = 38.6°$ and a cutoff phase $\phi_f = 46.3°$ reached after 47 coils with v_f 136 m/s, and (3) has constant phase $\phi = 44.3°$ and 132 m/s final velocity. The slow peak of (1) has increased flux by a factor of 1.18 over (2) and a factor of 1.85 over (3). The noise that appears on and around the unslowed portion of the beam is electronic noise due to the pulsing of the coils. This data was taken before the coil electronics was optimized, which is why the phases are higher than for the data described in figure 4.30 and table 4.2.

actually a transverse rather than longitudinal effect. With so many coils, the beam can be brought to low velocity using small phase angles. While the beam is still moving quickly the atoms propagate well into the coil before it is completely off, allowing the focusing and transverse guiding properties of the field inside the coil to provide transverse stability. However, by continuing to use a large phase even when the atoms are slow, they do not get into the coil before the field is off. This means that the antiguiding force applied outside of the coil dominates and the atoms are de-focused, leading to a loss of flux as atoms hit the walls of the coilgun and are lost.

The metastable neon beam is not slowed below 55.8 m/s because the beam flux decreases rapidly for lower velocities. Looking at how the flux decreases at both the extended and retracted positions of the MCP shows that the decrease in flux is primarily due to the transverse velocity of the atoms in the beam. However, this velocity is sufficient for most purposes, since the beam is slow enough that it could be loaded into a stationary magnetic trap.

4.6 The Molecular Coilgun: Stopping Molecular Oxygen

Having shown that the coilgun method gives full control of the velocity of a beam of atoms, the next step is to show the generality of the coilgun method. Slowing of metastable neon demonstrates control over atomic velocities, but slowing of a molecule demonstrates the broad applicability of the coilgun method, and shows that it is a general method of producing cold and slow samples. To this end, the coilgun is used to slow a beam of molecular oxygen. Oxygen is chosen because it has a permanent magnetic moment in the ground electronic state, and because of its importance in several chemical processes. Being a molecule, oxygen has a more complicated internal energy level structure with avoided level crossings between different rotational states. This makes an experimental demonstration of the coilgun method's

viability for molecules essential.

In the experiment described here, oxygen is slowed from an initial velocity of 389 m/s to a 83 m/s, removing more than 95% of the initial kinetic energy. The apparatus used for this is almost identical to the apparatus used to stop metastable neon, and the coilgun used is exactly the same. As such, only the differences between the experiments (in beam creation and detection) are discussed.

4.6.1 Oxygen's Magnetic Levels

At this point, it is useful to consider how the rotational spin states of oxygen affect its interaction with a magnetic field. The state that is slowed is the $^3\Sigma_g^-$ ground state. Nuclear statistics forbid the existence of $K = 0, 2, 4, \ldots$ rotational levels of $^{16}O_2$. The two unpaired electrons of the oxygen molecule lead to an electron spin angular momentum $S = 1$. The sum of spin angular momentum, S, and the rotational angular momentum, K, gives three possible total angular momentum states $J = 0, 1, 2$ for the $K = 1$ ground state ($J = K + S$, electron orbital angular momentum in the Σ state is zero). In the weak-field Zeeman regime, the three total angular momentum states each split into 3 sub-levels with different total angular momentum projections in the direction of the magnetic field. At high fields (above 2.5 T) the electron spin decouples from the rotational angular momentum and forms three different sets of three levels, where each set is almost-degenerate. The $J = 2$, $M_J = 2$ sub-level has the highest magnetic moment both in the low and high field regions, which leads to the magnetic moment being approximately the same in both the low and high field regime. The magnetic moment of this sublevel is approximately equal to 1.8 Bohr magnetons [86]. This magnetic sublevel has an avoided level crossing with the $J = 2$, $M_J = 2$ sub-level from the $K = 3$ manifold at about 8 T, changing its behavior from low-field to high-field seeking, which is incompatible with the coilgun as it is currently

operated. More detail on the magnetic and rotational states of molecular oxygen can be found in [86].

Since high-field-seeking states are incompatible with the coilgun in its current method of operation, the avoided crossing sets a limit on the maximal value of the magnetic fields that can be used to slow oxygen. The measured peak value of the magnetic fields created on the axis of the coilgun solenoids is 5.2 T, and finite element calculations indicate that the maximal magnetic field adjacent to the coil windings is just below 6 T. This maximal field which an atom can encounter in the coilgun is well below the limit of 8 T. Thus, the avoided crossing is not expected to play a large role in the operation of the coilgun. In order to achieve the highest slowing efficiency, coil timing is tuned for molecules in the ground rotational $J = 2$, $M_J = 2$ sub-level.

4.6.2 Oxygen Beam Creation and Detection

Since the ground state of oxygen is magnetic, a discharge is not required to excite the molecule to a metastable state to enable its slowing by the coilgun. However, because the boiling point of oxygen is ≈ 90 K, the nozzle cannot be cooled to 77 K by liquid nitrogen. Additionally, the mass of molecular oxygen is greater than that of metastable neon, and so oxygen has a greater kinetic energy at the same velocity. The magnetic moment of oxygen is less than that of metastable neon as well, which means that less energy will be removed per coil. These two issues make it highly advantageous to start with as slow an initial beam as possible. For this reason, rather than continue to use neon as a carrier gas, krypton is used instead as its higher mass leads to a slower initial beam. Oxygen is mixed with krypton at a 15:75 ratio, and the valve is run with a backing pressure of 34 psi.

Oxygen is a liquid at 77 K (krypton is too, its boiling point is $\approx$ 120 K), so cooling the nozzle with liquid nitrogen is not a possibility. However, it is still

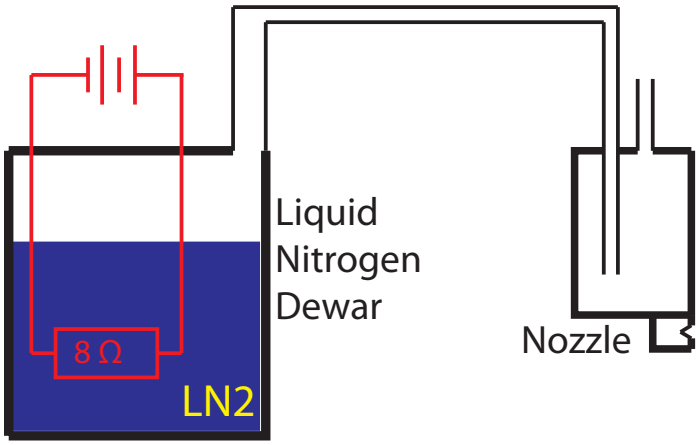

Figure 4.32: A schematic summary of the system used to cool the nozzle to 150 K. A 50 W 8Ω resistor provides a variable heat source in a liquid nitrogen dewar, which controls the rate of the boiloff. The cold gaseous nitrogen is forced to flow through the cryostat of the nozzle, cooling it. Increasing the current in the resistor increases the boil off rate, and cools the nozzle more.

advantageous to cool the nozzle significantly below room temperature to reduce the initial velocity of the beam. With a lower temperature limit of 120 K from the krypton carrier gas, and due to a desire to avoid excessive clustering and heating of the beam, the nozzle is run at a temperature of 150 K. To run the nozzle at this temperature, a gas line is run from a tank of liquid nitrogen to the pool type cryostat of the nozzle. This directs the cold boil off gas from the liquid nitrogen dewar into the cryostat, cooling the nozzle. The flow rate of the boil off is controlled by inserting a 50 W 8 Ω resistor into the nitrogen dewar, with greater current leading to more power in the resistor and faster boil off. Currents of up to 3 A are run through the resistor, which exceeds the power rating and is only feasible because the resistor is immersed in liquid nitrogen. A schematic illustration of this method of cooling the nozzle is shown in figure 4.32.

The final change in the apparatus concerns the detection method. Since the oxygen is not metastable, it will not create a free electron when it hits the surface

of an MCP, and thus will not be detected. However, the RGA used to detect helium in chapter 3 can also be used to detect molecular oxygen. The oxygen is ionized by electron impact and accelerated into a quadrupole which provides mass filtering. The ions are then detected with an electron multiplier. The RGA is mounted on the same translation stage as the MCP, which again permits measurement of the velocity of the beam. However, the ionization region of the RGA has some depth, as compared to the flat MCPm, which creates inherent uncertainty in the position of an oxygen molecule when it is ionized and subsequently detected. The distance from the end of the coilgun to the ionizer with the translation stage in the retracted(extended) position is 100 mm(150.8 mm).

4.6.3 Oxygen Data and Results

The results of the oxygen slowing are now presented. Using the oxygen-krypton mixture described above, and with the nozzle cooled to 150 K, the oxygen beam without any slowing has a velocity of 389 ± 5 m/s with a standard deviation of 28 m/s. This corresponds to a temperature of 3.0 K. This is significantly hotter that the previous beams described here. Clustering in the beam is the likely explanation of the increased beam temperature, as the valve is operated close to the condensation temperature of the krypton carrier gas.

Time-of-flight profiles of the oxygen beam are presented in figure 4.33. The coilgun is operated at constant phase, where the phases used are between 47.8° and 63.2°. This corresponds to slowing of the beam to 242 m/s at the lowest phase, and to 83 m/s at the highest phase. For the 83 m/s beam, over 95% of the initial kinetic energy is removed. A full summary of the slowing results can be found in table 4.3.

The time-of-flight profile in the graph on the right in figure 4.33 shows the entire beam, including the perturbed fast beam. As is clear from the figure, the

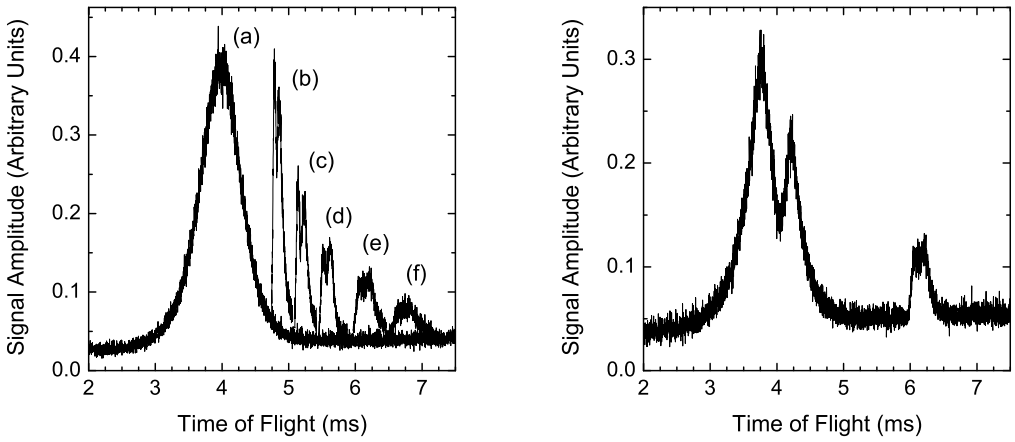

Figure 4.33: Molecular oxygen slowing data from the 64 stage coilgun. Each time-of-flight curve is the average of 200 individual measurements. The slowed curves only show the slowed portion of the beam for clarity. The time of flight signal for a beam without firing the coils is shown in (a). The velocities of the beams shown on the left are (a) 389 m/s, (b) 242 m/s, (c) 195 m/s, (d) 155 m/s, (e) 114 m/s, and (f) 83 m/s. These results are summarized in table 4.3. The figure on the right shows the entire time of flight profile of the beam slowed to 114 m/s, including the perturbed initial beam. This shows how a portion of the initial beam is slowed, while the rest is relatively unchanged.

Table 4.3: Final velocities (v_f), simulated final velocities of the synchronous atom (v_s), and coil phases, of the beams shown in figure 4.33.

	v_f [m/s]	v_s [m/s]	phase [degrees]
(a)	389 ± 5	–	–
(b)	242 ± 13	243	47.8°
(c)	195 ± 8	195	53.7°
(d)	155 ± 5	156	57.7°
(e)	114 ± 3	114	61.2°
(f)	83 ± 3	84	63.2°

coilgun is only able to slow a portion of initial phase space volume. Part of this is due to the phase space acceptance of the coilgun as defined by its volume of phase stability. It is also partially due to the fact that not all of the oxygen molecules in the initial beam are in the correct angular momentum state to be slowed by the coilgun. Comparison with simulation indicates that the oxygen coilgun has an efficiency which ranges from around a percent for the least slowing, to about a tenth of a percent for maximal slowing.

Slowing of a molecule using the coilgun method demonstrates the broad applicability of pulsed magnetic slowing of supersonic beams. The supersonic source is a general method of producing cold but fast samples and, as demonstrated here, the coilgun is a general method of reducing the velocity of these cold samples to trappable velocities. These techniques combine to provide a general method of controlling atomic motion.

Chapter 5

Towards Trapping and Cooling of Atomic Hydrogen Isotopes

Having verified that the coilgun functions as a general method of producing slow and cold samples, the next step is to use the samples it produces to make a physical measurement. For many measurements, such as precision spectroscopy, a long interrogation time is highly desirable. Even with a slow beam, the observation time is limited because the atoms are not confined to the interrogation region. However, by integrating a trap with the end of the coilgun, the atoms can be observed for several seconds at a time (depending on the lifetime of the trap).

This chapter describes the ongoing experiment to trap, cool, and study atomic hydrogen isotopes. The coilgun is modified to a hydrogen specific device and these modifications, along with the reasons behind them, are described. Hydrogen plays a particularly important role in physics and the rational behind the desire use the coilgun to study hydrogen is explained. An anti-Helmholtz magnetic trap is added to the end of the coilgun to trap the slowed atoms, and this trap and its switching electronics are detailed. The entire device is characterized and is believed to be working properly. While the experiment has not yet succeeded in detecting trapped hydrogen, for reasons which are explained, detailed simulations have been performed. These simulations suggest that the hydrogen coilgun and trap will operate with high efficiency, trapping a significant quantity of the atomic hydrogen entering the coilgun. The apparatus described in this chapter is also described in [81]. Finally, the future goals of the experiment, including a proposed method for cooling the trapped atoms

and possible future improvements to the coilgun, are described.

5.1 Hydrogen Motivation and Structure

As the simplest atom, and the atom best understood by theorists, hydrogen serves as an important test system for fundamental physics. Precision measurements of the atomic structure of hydrogen, such as the determination of the Lamb-shift [87], provide stringent tests of theory. The frequency of the $1S - 2S$ transition is one of the most precise atomic measurements to date [88–91], and is important in the determination of fundamental constants, such as the Rydberg constant. Measurements on muonic hydrogen have contributed to the determination of the size of the proton [92], and studies of the isotopic shift of deuterium have helped to determine the size and structure of the deuteron [93, 94]. Finally, work with anti-hydrogen [95–97] promises to shed light on important fundamental symmetries.

Due to the importance of hydrogen, considerable effort has been spent on trapping and cooling of hydrogen over the years. Hydrogen has been magnetically trapped at high phase space density [98] and has been evaporatively cooled to quantum degeneracy in a cryogenic apparatus [99, 100]. Unlike many other atoms, the internal structure of hydrogen is amenable to laser cooling, but the 121 nm wavelength needed to excite the cycling transition is difficult to produce with sufficient power for laser cooling. Nevertheless, hydrogen has been laser cooled, though at very low cooling rates [101]. A pulsed Sisyphus cooling scheme has been proposed [102], as has a scheme using ultrafast pulses [103]. The difficulties that present themselves in all of the above experiments and proposals, along with the canonical importance of hydrogen, motivate the search for a method of trapping and cooling of hydrogen isotopes in a room temperature apparatus.

As a gas, hydrogen is simple to seed into a supersonic beam, and the cooling

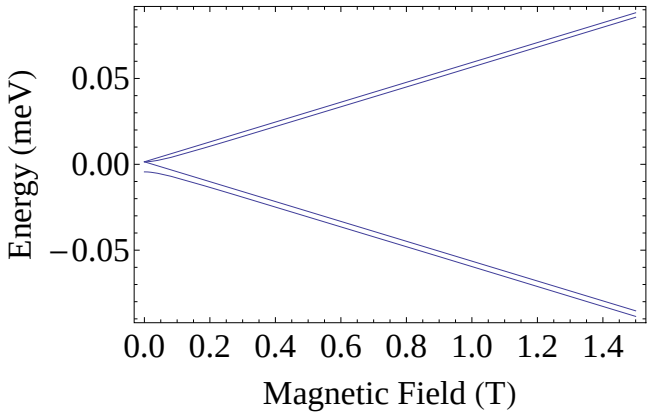

Figure 5.1: The hyperfine structure of the ground state of hydrogen. The transition between the low field and high field regimes takes place around .01 T, and at the fields used in the experiment described below the structure can be approximated as two states, with magnetic moments of $\pm$ 1 Bohr magneton.

provided by the expansion is sufficient to trap a significant portion of the beam. Furthermore, the magnetic properties of atomic hydrogen are ideal for use in the coilgun. Hydrogen is the lightest of the atoms with a mass of 1 amu, and in the high field limit, has a magnetic moment of $\pm$ 1 Bohr magneton. The hyperfine structure of the ground state of hydrogen as a function of magnetic field is shown in figure 5.1. Because the transition between the high field and low field limits takes place around 0.01 T, the behavior of hydrogen in the fields of the coilgun and trap is nearly always that of the high field limit. The large magnetic moment to mass ratio (the highest for ground state atoms), makes hydrogen ideal for use in the coilgun. The importance of precision measurements on hydrogen isotopes, coupled with the excellent compatibility of hydrogen with the coilgun method, motivates the work described here. It should also be noted that the coilgun method has been pursued independently and in parallel to the one described here [21, 22], and been used to confine hydrogen and deuterium for a few milliseconds [23, 24].

5.2 The Hydrogen Apparatus

The apparatus constructed to slow and trap a supersonic beam of hydrogen is different from the coilgun used to slow neon and oxygen. The differences are primarily due to the increased magnetic moment to mass ratio of hydrogen, which significantly lowers the magnetic field strength requirements while still permitting slowing in a short distance. The lower field requirements permit the coils to be increased in size. This increases number of atoms slowed due to the increased aperture and allows the coils to be removed from vacuum, reducing the experimental complexity. The new coil geometry and electronics are described, along with the advantages of this change in geometry. The experiment also integrates a magnetic trap at the end of the coilgun, and the method by which the atoms can be trapped is explained, along with the trap design and electronics. Finally, since the beam consists of hydrogen, the beam creation and detection are changed.

5.2.1 Slowing Coils and Electronics

The greatly increased magnetic moment to mass ratio of hydrogen ($\frac{1\,\mu_B}{1\text{ amu}}$), compared with metastable neon ($\frac{3\,\mu_B}{20\text{ amu}}$) or molecular oxygen ($\frac{1.8\,\mu_B}{32\text{ amu}}$), permits smaller fields to be used in a convenient laboratory scale coilgun. An increase in the coil size is thus feasible, which results in several advantages. The most significant advantage of using larger coils is that the coils can be moved outside of the vacuum envelope, greatly reducing the experimental complexity. Outside of vacuum, the need for feedthroughs is eliminated and the coils can be cooled much more easily, allowing the repetition rate of the experiment to be increased and reducing the time required to take data. The larger coils require a modification to the switching electronics since the inductance of a coil increases with its size.

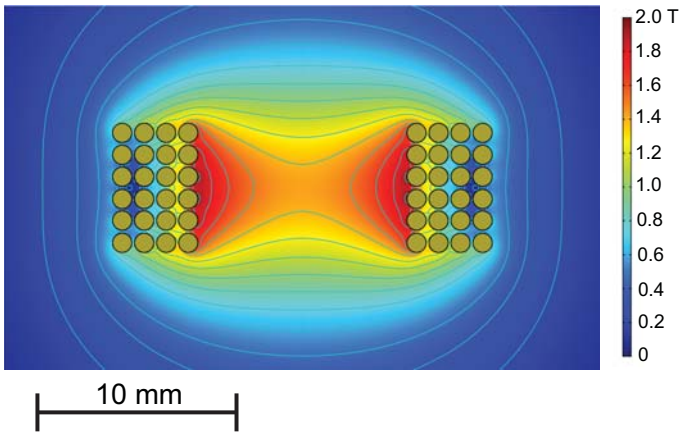

Figure 5.2: The spatial profile of the magnetic field produced in the hydrogen coilgun coils at 825 A. The fields are calculated numerically using finite element analysis.

5.2.1.1 Slowing Coil Geometry

Larger coils have a bigger physical aperture, allowing more particles to enter the coilgun and increasing the number of particles slowed. Larger coils also create a smaller field gradient. This means that more of the atoms in the beam experience approximately equal fields, again increasing the phase space volume which can be slowed and thereby improving the slowed flux. The reduced field gradient is especially important at lower velocities. Small differences in the energy removed result in a larger spread in velocities for lower speeds, due the quadratic relationship between energy and velocity. In addition, the same velocity spread in a beam will have a much larger effect at lower velocity. For example, a Δv of 5 m/s on a beam moving at 500 m/s results in the beam spreading by less that .1 mm between coil, but a Δv of 5 m/s on a beam moving at 50 m/s results in a spread of more than 2 mm over the length of a single coil. Larger coils and lower field gradients are therefore particularly beneficial, in terms of final slowed flux, for low final beam velocities.

The coils consist of 24 windings (4 layers of 6 turns) of 1 mm Kapton insulated copper magnet wire. The inner diameter of the coil is 11.5 mm and the outer diameter

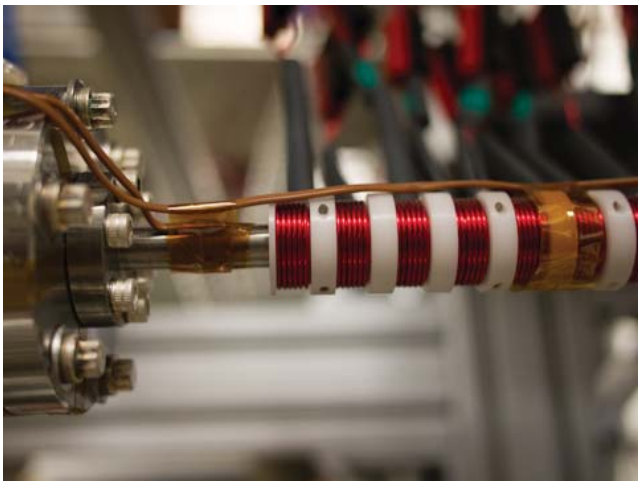

Figure 5.3: Photograph of the hydrogen coilgun coil wound around the slowing chamber. The coils are held in place by Delrin clamps, which are clamped onto the chamber before the coil is wound.

is 19.5 mm. The overall length of the coil is 7.26 mm. When the coil is pulsed, a peak current of approximately 825 A flows through the coil, producing peak on axis fields of 1.4 T. The finite element numerical calculation of the field profile can be seen in figure 5.2.

The coils are wound around the vacuum chamber and held in place by Delrin spacers that are clamped to the chamber, as shown in figure 5.3. This allows the coils to be as close as possible to the atoms without being in vacuum. The Delrin spacers place the coils 12.45 mm apart center-to-center. This spacing is set by the wire thickness, as the wire from the inner layer of a coil needs a path to the outside of the coils. Making the coils much closer together would make it difficult to insert the inner winding if the adjacent coil is already wound.

5.2.1.2 Slowing Coil Electronics and Switching

Due to the increased size of the coils, the inductance of the hydrogen coilgun coils is larger than the coils used in the previous coilgun iterations (10 μH instead

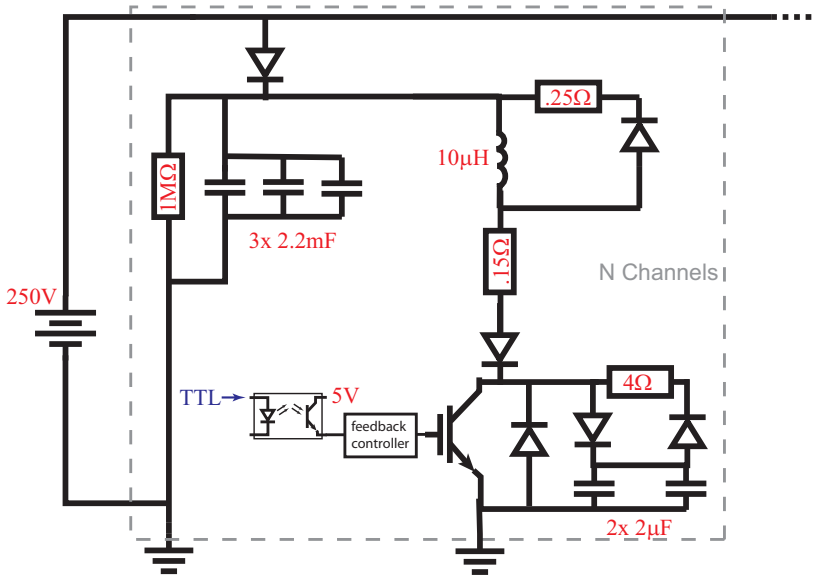

Figure 5.4: A schematic of the hydrogen coilgun coil switching circuit. Each coil has its own IGBT. The only shared circuitry is the power supply circuit which charges the capacitors. This change eliminates the need for a thyristor, and is feasible because only 18 coils are required to slow hydrogen, reducing the cost of not sharing IGBTs.

of the previous 7 μH). In addition, the effective resistance of the coil is significantly reduced due to the increased wire size and the lack of vacuum feedthroughs. Since the magnetic moment to mass ratio of hydrogen is so high, only 18 coilgun coils are needed, reducing the need to switch multiple channels with the same circuitry. Finally, the current being run in the coil has increased from 750 A to 825 A. All of these changes result in the need to tune the switching circuit to better suit the needs of the hydrogen specific coilgun.

A schematic of the modified switching circuit can be seen in figure 5.4 and a photograph of the circuit can be seen in figure 5.5. This circuit differs from that described in sections 4.4.3 and 4.5.2 in several ways, though the overall form remains broadly similar. The largest change from the previous circuits is the removal of the thyristor. This significantly increases the reliability of the circuit, as the thyristors

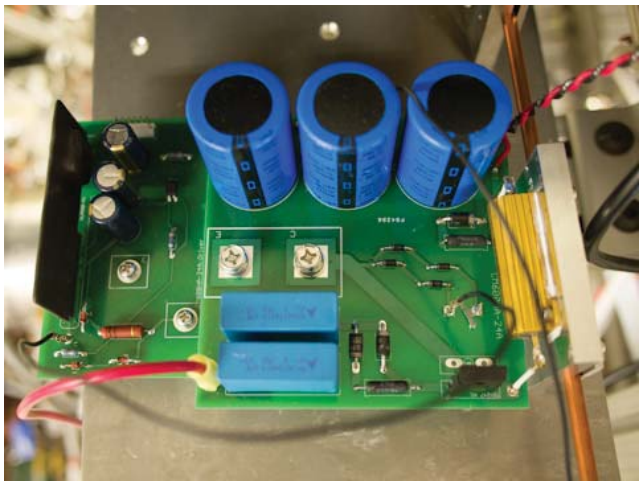

Figure 5.5: Photograph of the hydrogen coilgun coil driver circuit. The board on the right is the circuit which drives the coils, while the board on the left is the circuit which drives the gate of the IGBT.

are the component that failed at the highest rate. The down side of removing the thyristor though is that each coil requires its own IGBT (changed to a Powerex CM600HA-24H 1200 V 600 A IGBT for higher current carrying capability), which is feasible (each IGBT and gate driver costs $250) as only 18 coils are needed because of the high magnetic moment to mass ratio of hydrogen.

There are a few other changes made to the circuit, most notably the return to a resistor in series with the coil, and a freewheel diode and resistor to lower the voltage spikes produced by switching the coil. The series resistor is required because the new coil has significantly lower resistance, leading to too great a current for the IGBT without this resistor. Similarly, the greater current and increased inductance of the coil produces a damaging voltage spike on the IGBT without the freewheel diode. Even with these additions, the snubber capacitor must be doubled in size to limit the switching voltage.

The final change in the coilgun electronics is the addition of two more discharge

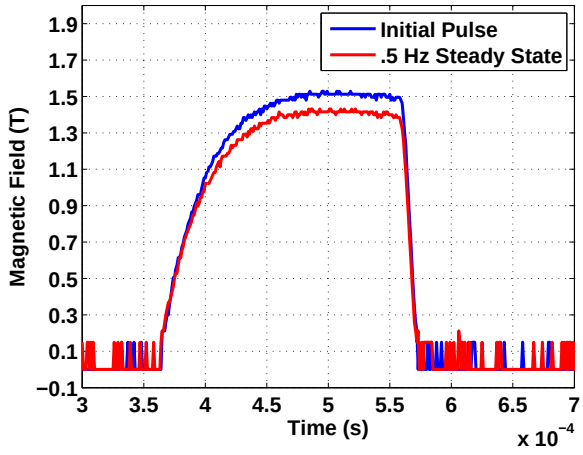

Figure 5.6: The time profile of the magnetic field in the hydrogen coilgun coils as measured by Faraday rotation. The blue curve shows the field measured for the initial current pulse in the coil, while the red curve shows the field after the coil and resistors in the switching circuit have warmed up by being pulsed at .5 Hz. The noise in the measurements that appears when field is near zero is due to the resolution of the digital oscilloscope used to record the signal on the photodiode.

capacitors to drive each coil. These are needed for a few reasons. First, the lack of Permendur around the coil means that the field extends much farther from each coil than in the previous coilgun iterations. This means that in order for the atom to feel the full effect of the field, the coil must be turned on earlier. This is especially true for the slow atoms at the end of the coilgun. Hence, the pulse length is increased to 200 μs. Additionally, the increased current in the coil means that the capacitors discharge faster. A longer pulse of greater current requires larger capacitance to keep the current nearly constant through the end of the pulse. Since the coils no longer share IGBTs, the pulses can overlap at the beginning of the coilgun without causing any problems.

The field in the coils and the temporal switching profile of the magnetic field is measured using the Faraday effect. The field turns on exponentially with a time constant of 30 μs, and turns off linearly in 10 μs. The time profile of the magnetic

field can be seen in figure 5.6. There are two curves shown here, one for the initial pulse, and one after the wires and resistors in the circuit have heated up by pulsing the coil at a .5 Hz repetition rate.

One other change made from the coilgun used to slow neon and oxygen is in the control electronics. The larger coils of the hydrogen slower mean that the timing resolution needed in the control circuitry is reduced. Simulations show that the timing resolution of 100 ns used in the 64 stage coilgun does not have any benefit over a much easier to implement 1 μs resolution. Because the experiment is controlled by LabVIEW, it is simpler to use LabVIEW and National Instruments data acquisition cards (NI PCIe-6259) to control the coil timing, rather than using LabVIEW for data acquisition and an FPGA for control.

5.2.2 Trapping Coils and Electronics

The experimental goals require trapping the slowed hydrogen atoms at the end of the coilgun, and so the experiment must incorporate a trapping potential and a method of bringing the atoms to rest at the center of that potential. Because the coilgun operates on low-field-seekers, the slowed atoms are already in a magnetically trappable state, and state preparation between the coilgun and the trap is not necessary. A static magnetic trap is a conservative potential, and the atoms must first be brought to rest in the trapping volume. The proposed trapping sequence this is explained in this section. The details of the anti-Helmholtz magnetic trap and the switching circuitry built to control the magnetic potentials created by the trapping coils are described.

5.2.2.1 Principle of Operation of the Trap

To maximize the number of atoms trapped at the end of the coilgun, the coilgun and trap need to be mode-matched. What this means is that the stable phase space volume of the trap should closely match with the phase space volume occupied by the slowed atoms at the end of the coilgun. Using an anti-Helmholtz trap configuration, with appropriate coil dimensions and spacing, enables this. The anti-Helmholtz trap is placed as close as possible to the end of the coilgun, and the coils are co-axial with the coilgun coils.

To aid in the mode-matching of the atoms to the trap, the first trapping coil the atoms encounter (this is referred to from now on as the front trapping coil) will be used as a final slowing coil, and will be operated in the same manner as the coilgun coils. Initially both trap coils are turned on with the same polarity (the polarity of the coilgun coils). As the atoms enter the front trapping coil, the coil is switched off, in an exact analog to the coilgun coils used to slow the atoms. Using the front trapping coil in this manner serves several purposes. First, since the coil is used as the final slowing coil, the atoms can be extracted from the coilgun at a slightly higher velocity, enabling a slightly greater stable phase space volume in the coilgun. Additionally, using the front trapping coil to do the final slowing means that the atoms do not have much distance to propagate between the last slowing coil and the trap center, limiting the degree to which the packet of atoms will spread out during this propagation. Finally, the phase stability and waveguide effects of the coil serve to preserve the phase space density of the slowed packet and aid in the mode matching of the slowed atoms to the trap.

After the atoms have passed through the front trapping coil, they enter the trapping region and will be brought to rest by the magnetic potential of the rear trapping coil. The phases of the coils in the coilgun can be tuned such that the

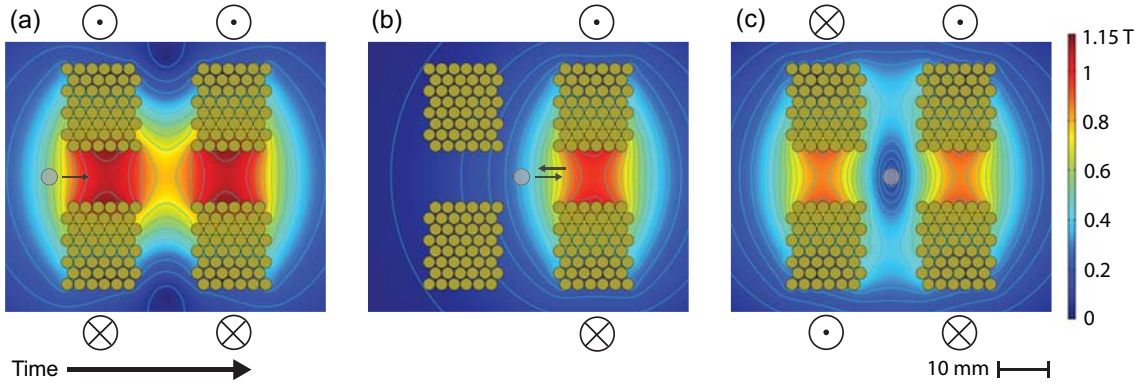

Figure 5.7: This figure shows numerically calculated magnetic field profiles at different points in the trapping sequence. In (a), the slowed bunch is approaching the trap from the coilgun and both coils are turned on with the same polarity. The front trapping coil is used as the final slowing coil. Once it is turned off, the atoms are brought to rest at the center of the trapping volume by the rear trapping coil, as illustrated in (b). At this point the front trapping coil is quickly switched back on with the opposite polarity, creating a 100 mK deep trap quadrupole trap for the cold atoms, which is shown in (c).

slowed packet has the proper velocity for it to come to rest in the center of the trapping volume. At this point in the trapping sequence, the front trapping coil will be quickly switched back on, but with the opposite polarity. This forms an anti-Helmholtz quadrupole trap. The entire trapping sequence is shown in figure 5.7.

It may also be necessary to remove the trapped atoms from the trap in order to detect them. To do this, the rear trapping coil can be turned off, leaving the front trapping coil on. The field from the front trapping coil will repel the low-field-seeking atoms out of the trap through the bore of the rear trapping coil. The rear trapping coil can be switched back on with the opposite polarity once the atoms have passed through the bore, which will continue accelerating the atoms into a detection surface or volume.

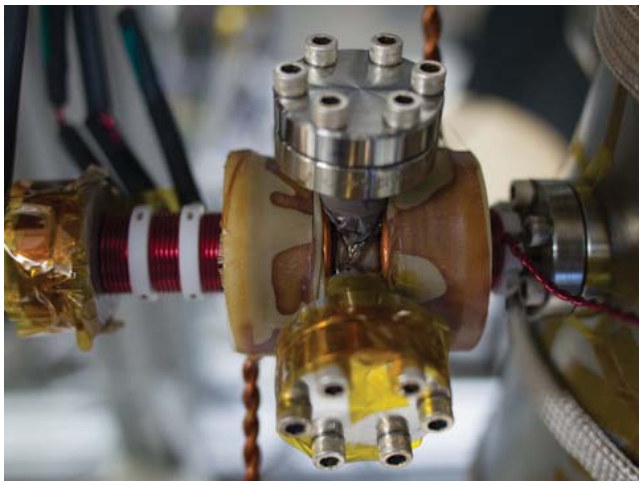

Figure 5.8: Photograph of the trapping coils wound around the chamber. The coils are held in place with epoxy, and wrapped with an epoxied nylon cloth. Delrin clamps on either side of the trap hold the coils in position.

5.2.2.2 Trapping Coil Geometry and Fields

Each trapping coil consists of 48 turns (8 layers of 6 turns) of Kapton insulated copper tubing. Each coil is 1.57 cm long and has a bore of 1.05 cm. The outer diameter of each trapping coil is 4.44 cm. The wires have a 2.4 mm outer diameter and an inner diameter of 1.2 mm. Using hollow wires allows the coils to be cooled by running distilled water through the core of the wires themselves. The water flow rate through the coils is 5 ml/s. This amount of cooling allows the coils to be run with 200 A of steady state current. Each coil has a measured resistance of 0.028 $m\Omega$, meaning that the water cooling is capable of removing a little more than 1100 W. Given that this much power can be removed from each coil, at a repetition rate of 0.5 Hz (given by the coilgun coils) the trap coils can be operated at their full current of 450 A for up to 300 ms without over-heating.

The coils are wound around in place around the trapping chamber, since the bore of the coil is too small for them to slide over the flanges of the chamber. They

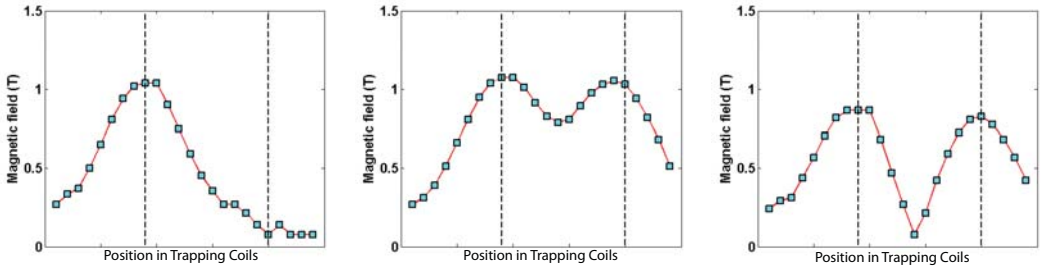

Figure 5.9: Magnetic field values on the axis of the trapping coils measured by Faraday rotation. The dashed lines indicate the center of the two trapping coils. The field measurements are taken 2.5 mm apart. The plot on the left shows the field for a single coil, while the plot in the center and on the right show the fields with the coils running in Helmholtz and anti-Helmholtz configurations respectively. In all three configurations, the current in the coils is 450 A.

are held in position on the chamber by Delrin clamps, which also hold the final three slowing coils of the coilgun. Once wound and clamped in position, the coils are epoxied (Epo-Tek 353nd) and wrapped with epoxy impregnated nylon cloth. This ensures that the coils do not shift or move under the exerted magnetic forces. The center of the trapping volume is located 2.99 cm from the center of the last coilgun coil, and the coils are 2.7 cm apart, measured center-to-center. A photo of the trap coils on the chamber is shown in figure 5.8.

The magnetic field values on the axis of the trapping coils, measured by Faraday rotation, is shown in figure 5.9. When 450 A is run through a single coil, the peak field on axis is 1.05 T. The measured field at the center of a coil is higher when the coils are run in Helmholtz configuration, with a measured peak field of 1.11 T. When the coils are run in anti-Helmholtz configuration, the field value on axis is lower, 0.85 T. In anti-Helmholtz configureation, the field is at a minimum at the center of the trap, creating a 100 mK deep trapping potential. The oscillation period for a hydrogen atom in this potential is approximately 1 ms.

Atoms in the trap need to adiabatically follow the field, as discussed in section

4.1.3. In an anti-Helmholtz trap, the trapping field vanishes at the center of the trap. This means that the Larmor frequency also goes to zero and thus the atoms cannot follow changes in the field direction. The radius from the trap center where atoms of a particular velocity cannot follow the field and undergo spin flips is derived in [104]. For the trapping fields described above, the calculated radius of spin flits is $\approx 3\,\mu$m for 100 m/s hydrogen atoms. Most of the atoms in the trap are slower than 100 m/s, and the radius for spin flips will be even smaller for slower atoms. Such a small radius, even for the fastest atoms, means that spin flips are not expected to be a significant source of trap loss.

5.2.2.3 Trapping Electronics and Switching

Due to the requirements of the experiment, each trapping coil must be able to have current flow in either direction. In addition, the polarity of the current in one coil must be independent of the polarity or state of the other coil. Each coil needs to be able to be switched on or off independent of the state of the other coil as well. Fast switching times are required, because the loading of the trap depends on using the front coil as the final coilgun coil and on switching this coil back on quickly to confine the atoms in the trap. It is desireable to have the coils operating in series, rather than in parallel to ensure that both coils operate at the same current, locating the trap center in the middle of the trapping volume. It also means that any current noise will appear in both coils, rather than just one. This avoids heating of the atoms due to an induced motion of the trap center.

A schematic of the circuit constructed to fulfill these requirements is shown in figure 5.10 and a photograph of the circuit is presented in figure 5.11. Each coil is switched using an H-bridge configuration of discrete metal-oxide-semiconductor field-effect transistors (MOSFET). Each MOSFET shown in the figure represents a

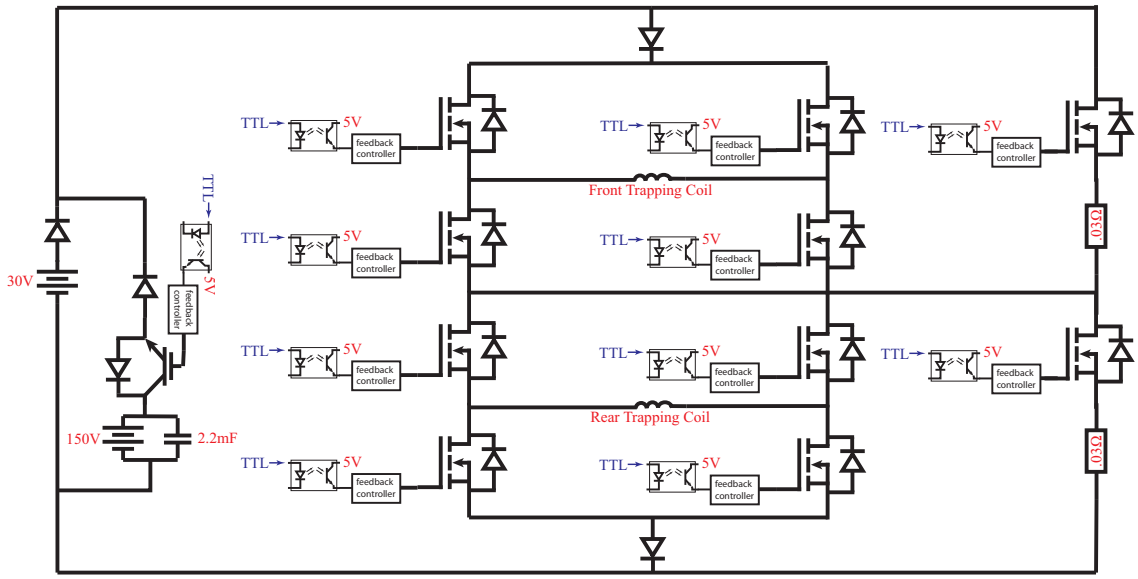

Figure 5.10: A schematic of the trapping coil switching circuit. This circuit allows both the front and rear coils to be switched on or off independently, due to the bypass paths on the right side of the diagram. The current direction can be changed independently, since each coil is set-up with an H-bridge of MOSFETs (each MOSFET shown represents a bank of 8 individual FETs). The IGBT, high voltage supply, and capacitor on the left side of the circuit allow the current to be turned on quickly using the high voltage when needed.

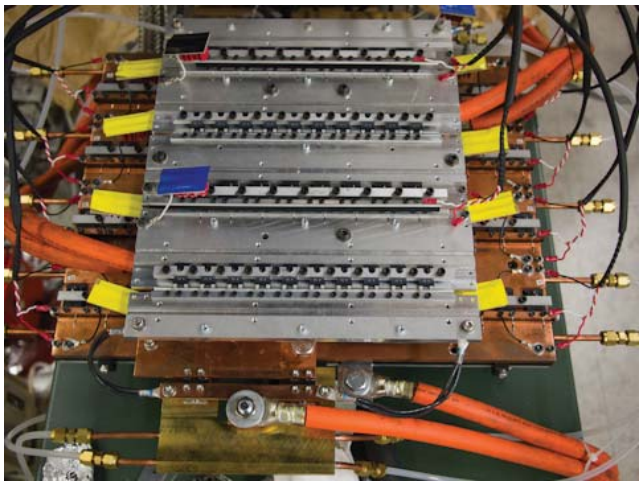

Figure 5.11: Photograph of the hydrogen trap coil driving circuit. The bypass circuit is the top layer, with the two H-bridges in series in copper on the bottom layer. Varistor banks are placed across each MOSFET bank to absorb any voltage spikes.

bank of eight individual MOSFETs (IXYS IXFX230N20T) connected in parallel (a MOSFET bank). The H-bridges are wired in series, ensuring that the current is the same in both coils. Each H-bridge also has a bypass circuit, which is switched by a separate MOSFET bank. This allows one coil to be switched off without turning off the current to the other coil. The current in the coils is supplied by a 30 V 500 A programmable power supply (Lambda ESS 30-500) with excellent current stability (0.1% regulation). The current control for the circuit is provided by the power supply itself rather than by a seperate current control feedback loop. Schottky diodes (Microsemi APT2X101S20J) are used between the bypass circuits and the coils to keep the body diode of the MOSFET from operating as a freewheel for the trapping coils and increasing the switching time. The circuit is designed to limit stray inductances by keeping gaps between current paths as small as possible. Varistor banks are placed across each MOSFET bank. Combined these measures limit the voltage spikes in the circuit.

The reason that MOSFET banks are used in the trap switching circuit, instead

of the IGBTs used for the coilgun coil switching circuit, has to do with the "on" state behavior of these devices. In their "on" state, IGBTs behave like diodes with a relatively fixed voltage drop as a function of current. They are also susceptible to thermal runaway, as the forward voltage drop decreases with increasing temperature. This makes it difficult to place multiple IGBTs in parallel. However, the forward voltage drop will even then lead to significant losses. While MOSFETs do not have blocking voltages as high as an IGBT, they behave like a resistive load in their "on' state, and their positive temperature co-efficient stops thermal runaway, enabling simple parallelling of multiple devices. Parallelling MOSFETs reduces the "on" state losses since the current is distributed across what is effectively many resistors in parallel, allowing the effective resistance to be reduced. For the low voltage but high current switching needed in the magnetic trapping coils, using MOSFETs enables lower losses in the switches.

The final element of the circuit is a high voltage capacitor and an IGBT (Powerex CM600HA-24H) in parallel with the power supply. This part of the circuit is used to turn the coils on rapidly by briefly increasing the voltage across the circuit. This voltage boost is particularly necessary when the front coil is turned on to create an anti-Helmholtz trap. The front coil must switch on quickly to capture the atoms once they have been stopped.

As the switching of the front trap coil is particularly important in the trapping sequence, the temporal switching profile is shown in figure 5.12. The profile is measured by Faraday rotation with the crystal located in the middle of the front trapping coil. This shows that the front coil switches off and on linearly in 150 μs. The IGBT provides a voltage boost to shorten the switch on time (it is over 600 μs without the IGBT in the circuit).

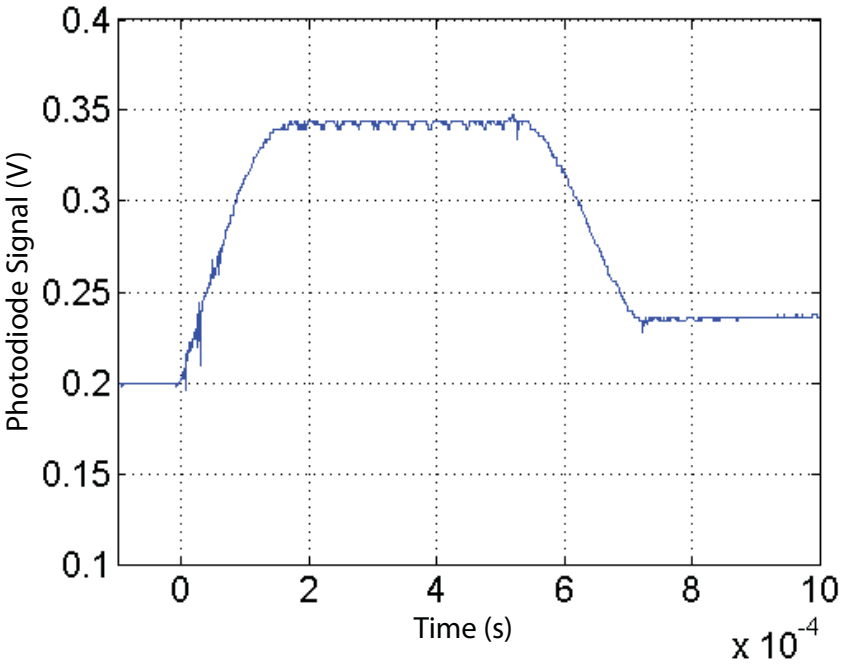

Figure 5.12: The Faraday rotation signal for a crystal located in the front trapping coil during the trapping sequence. Initially, both coils are on in a Helmholtz configuration (1.11T. The front coil is then switched off (without switching the rear coil) at $t = 0\mu$s. The field falls linearly over 150 μs and the coil remains off for 400 μs. At this time, the front coil is switched back on with the opposite polarity (using the IGBT boost) and the field in the center of the coil rises linearly over 150 μs to .85 T. The TGG crystal used to take this measurement is 5 mm long.

5.2.3 Beam Creation and Detection

Creating and detecting a supersonic beam of hydrogen is more challenging than for the neon and oxygen slowed in previous generations of the coilgun. Creating metastable neon requires a discharge, and the efficiency of metastable creation is likely around $10^{-4} - 10^{-5}$. However, detecting of the metastable atoms is possible with near unit efficiency using an MCP. Molecular oxygen is simple to seed into a beam and no state preparation is required. Detecting oxygen is more difficult however, and in the experiments described in chapter 4, oxygen is detected by an RGA, which has a similar efficiency to the discharge. Hydrogen is particularly difficult to work with because it requires a state preparation (cracking molecular hydrogen into atomic hydrogen) similar to metastable neon, and hydrogen must also be excited or ionized in order to detect it.

In this experiment, atomic hydrogen is created using the same discharge apparatus that is used to create metastable neon, described in section 4.4.1. Molecular hydrogen is mixed into a neon carrier at a ratio of 1:4, and the nozzle is run with a backing pressure of 50 psi. The near resonance between the neon metastable energy and the cracking energy of hydrogen plus the excitation energy of one hydrogen atom increases the efficiency of cracking with a discharge in the presence of neon [105].

Detection is accomplished by a custom designed jumbo ionizer and quadrupole mass spectrometer (Ardara Technologies). Hydrogen is ionized by electron bombardment and guided by ion optics into the quadrupole, which provides mass filtering. The electron energy can be tuned to optimize the ionization process. Unfortunatly, the detected atomic hydrogen ions are produced via two channels, and only one of these is desired. The first, and desired channel, is that atomic hydrogen produced by the discharge at the nozzle is ionized and detected. In the second channel, molecular hydrogen in the beam is cracked and ionized in the ionizer, creating a hydrogen ion

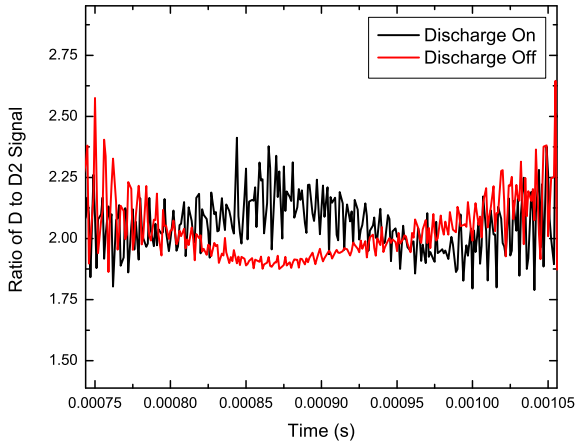

Figure 5.13: Time of flight plot of the ratio of D($\times$100) to D_2 with and without the discharge firing. The noise on either side of this signal is meaningless as the signals go to 0 and the increased noise is a result of dividing by nearly 0. The greater ratio of D to D_2 with the discharge indicates that the discharge is working properly and producing atomic hydrogen, but that it is covered by the background signal due to molecular hydrogen.

which is then detected. This second channel is an undesired background which causes a significant problem.

The background signal due to the cracking of hydrogen molecules in the ionizer causes atomic hydrogen to be observed even without firing the discharge. In fact, the observed atomic hydrogen signal actually decreases when the discharge is fired, since the beam is heated by the discharge and less of it arrives at the detector. It is thus difficult to determine the quality of the beam and discharge. However, one method of examining the beam did produce interesting results, and this was to look at the ratio of atomic deuterium to molecular deuterium, with and without the discharge (deuterium is chosen for its heavier mass, ensuring more atoms stay on the centerline of the beam). Doing this allows the effect of the discharge to be seen more clearly and shows that atomic deuterium is being produced by the discharge, but that the

signal is masked by the background due to molecules. Figure 5.13 shows the ratio of D to D_2 (multiplied by 100) with and without the discharge.

The background signal due to molecular hydrogen left in the beam presents a major problem for the experiment. This signal obscures the atomic signal and the experiment has been unable to use the quadrupole as originally intended. Other detection methods are being investigated, and their use is more fully explored in section 5.4.1.

5.2.4 Vacuum Chambers

The vacuum chambers for the hydrogen experiment can be divided into three sections, the beam creation and differential pumping chambers, the slowing and trapping chamber, and the detection chamber. While the beam creation chamber is quite similar to the design used for the previous experiments, the beam slowing and trapping chamber is quite different from the one used in previous generations of the coilgun since the coils are outside of the vacuum envelope. Also, in the previous coilgun experiments there was no dedicated detection chamber.

Previously the supersonic nozzle was in the center of a 6-way 6 inch cross, but in the hydrogen experiment the nozzle is in the center of a 6-way 8 inch cross. This change was made to allow for more space around the nozzle, as placing the filament for the discharge was difficult in the previous setup. The cross is pumped by a 500l/s Varian turbomolecular pump. The base of the skimmer (5mm diameter) now sits 142 mm from the exit of the nozzle. From the skimmer, the beam goes though a $2-3/4$ inch bellows to a 4-way $2-3/4$ inch cross which is pumped by a 70 l/s Varian turbo pump. Another skimmer (3 mm diameter) sits 380 mm from the nozzle. This second chamber provides differential pumping. Differential pumping is needed because of the small diameter of the slowing and trapping chamber. The conductance

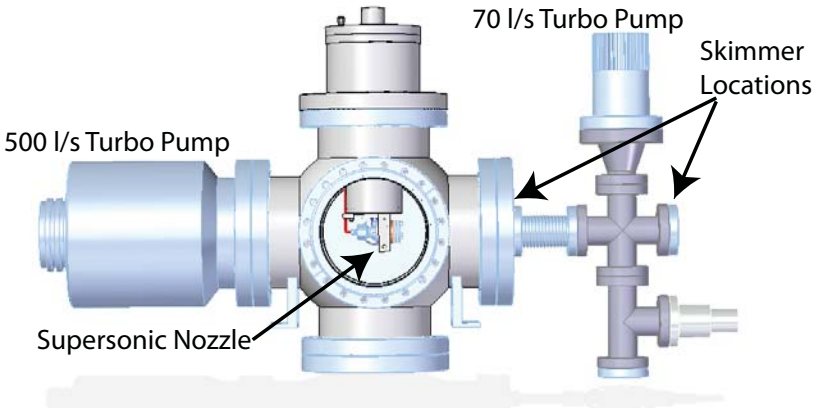

Figure 5.14: A CAD image of the hydrogen valve chamber and differential pumping chamber. The nozzle hangs from the pool type cryostat in the center of the 8 inch cross. The chamber on the right provides differential pumping.

of these latter chambers is low enough that without differential pumping, the gas load from the supersonic beam would build up in these chambers, decreasing the mean free path and trap lifetime. A CAD drawing of the valve chamber and differential pumping chamber can be seen in figure 5.14.

The slowing and trapping chamber consists of two parts, the coilgun tube and the trapping cross. The coilgun tube has two titanium 1.33 inch conflat flanges separated by a 22.0 cm grade 5 titanium tube. The tube has a wall thickness of .38 mm and an outer diameter of 9.52 mm. The reason this titanium alloy is chosen is its high electrical resistance which reduces the eddy currents caused by the coil, coupled with its lack of magnetism and good vacuum properties. The coils are wound around the tube, with the center of the first coil sitting 43.44 cm from the nozzle. This chamber has 15 of the 18 coilgun coils wound around it (the last three coils are wound around the trapping chamber). A CAD image of the slowing and trapping chamber, with the coils, is shown in figure 5.15 and a photograph of the assembled chamber with the coils is presented in figure 5.16.

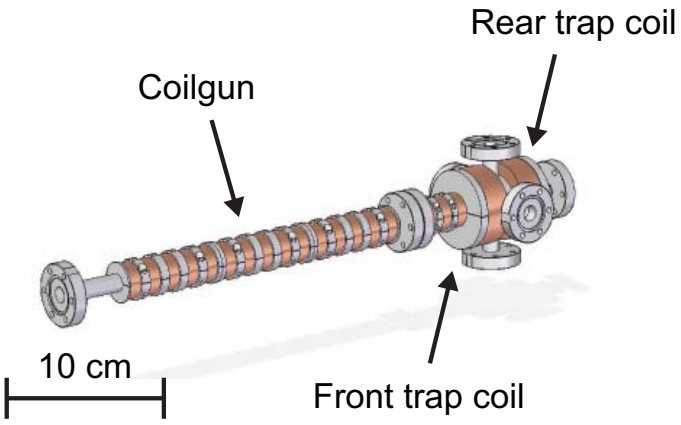

Figure 5.15: A CAD image of the hydrogen coilgun and trap along with the slowing and trapping vacuum chambers respectively. There are 18 coilgun coils, with 12.45mm spacing center to center (except where the slowing chamber and trapping chamber meet, where the coils are 28.65 mm apart).

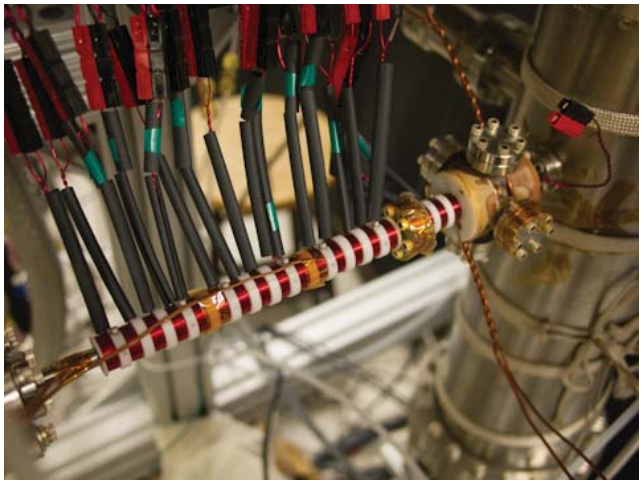

Figure 5.16: Photograph of the assembled slowing and trapping chambers, with the coilgun coils and trapping coils wound around the chambers themselves.

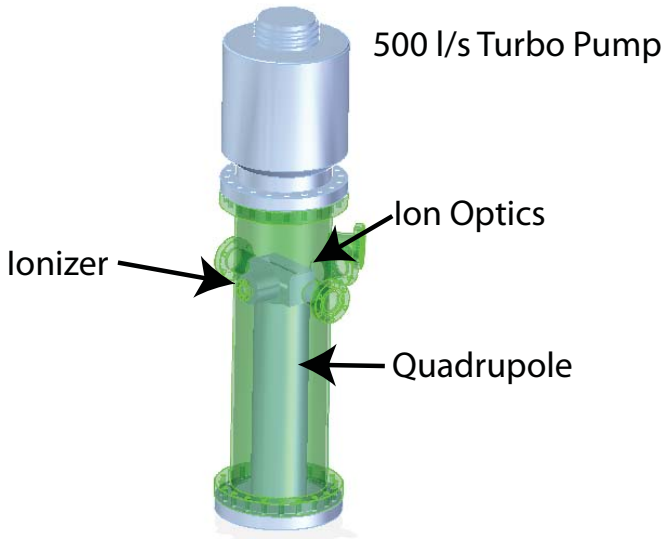

Figure 5.17: A CAD image of the hydrogen detection chamber with a mockup of the Ardara Technologies ionizer and quadrupole. The front of the ionizer sits 6.04 cm from the center of the trapping volume.

The trapping chamber is a 6-way 1.33 inch conflat cross with tubes made of the same dimensions as the coilgun chamber's tube. This chamber is made from 304 stainless steel, since it would be difficult to construct this chamber from titanium. The center of the trapping volume is 6.30 cm from the start of the chamber. The final three coilgun coils are wound around this chamber, and the junction between the coilgun and trapping chambers leads to a center-to-center distance of 28.65 mm between coils 15 and 16, instead of the 12.45 mm between the other coilgun coils. The cross arms of the chamber allow for optical access to the trapped atoms, which will be important for laser detection and spectroscopy (see section 5.4.1). Finally, the chamber extends past the trap to another 1.33 inch conflat 5.03 cm from the center of the trapping volume, which connects to the detection chamber.

The detection chamber serves two purposes in the experiment. First, it houses the Ardara Technologies ionizer and quadrupole mass spectrometer used to for detec-

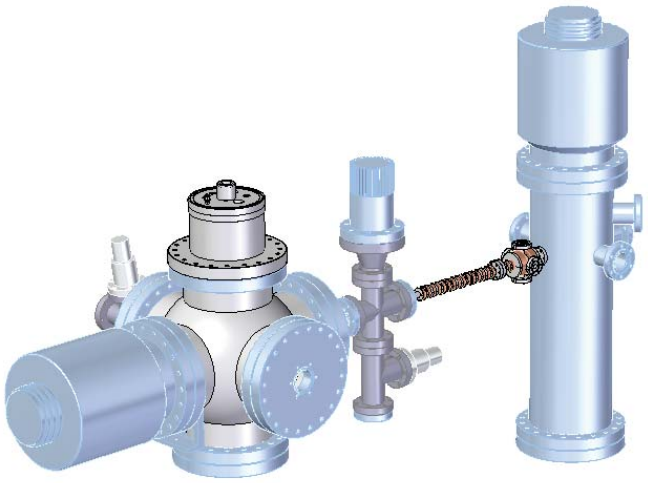

Figure 5.18: A CAD image of the complete hydrogen apparatus.

tion. The front of the ionizer sits 6.04 cm from the center of the trapping volume, and only 1.01 cm from the beam's entrance to the detection chamber. Atoms are ionized, and steered by ion optics though a 90° turn into the quadrupole mass filter. They are detected by an electron multiplier at the base of the quadrupole. The detection chamber also provides pumping for the unslowed portion of the supersonic beam. If the trapping chamber were terminated without the detection chamber, the gas in the beam would quickly fill the small volume of the slowing and trapping chamber, raising the vacuum pressure significantly. Having the detection chamber (with $\approx 10\,\mathrm{L}$ volume) which is pumped by a 500 l/s as a beam dump volume allows the pressure in the trapping volume to stay below 2×10^{-9} torr. A CAD image of the detection chamber is shown in figure 5.17, and the complete apparatus is shown in figure 5.18.

5.3 Simulations of Hydrogen Trapping

The device described above is designed to maximize the phase space density of the trapped hydrogen. Numerical simulations of the slowing and trapping process

guided the design of the apparatus. These simulations indicated the loss mechanisms that needed to be addressed when moving to a hydrogen specific coilgun. They also helped to determine the required field and switching parameters needed to create an effective magnetic trap. The simulation uses a variable step Runge-Kutta algorithm, coded in Visual C++, which numerically calculates the trajectory of one atom at a time. Monte-Carlo techniques allow the behavior of the beam to be determined by randomly initializing each atom according to the expected distribution of the beam entering the coilgun.

Once the design was finalized, simulations were run to optimize the switching sequence, and to determine what performance could be expected. Data from these simulations is shown below, indicating a trapping efficiency of 4.8%, at a temperature of 62 mK. As The detection of the trapped atoms may require ejecting them from the trap onto a detector, this process is also simulated, showing an extraction efficiency of 13%. Once the experimental parameters were finalized the measurements of the fields in the slowing and trapping coils, as well as their switching paraments, are inserted into the simulations. Furthermore, the dimensions used in the simulations are matched with those in the physical apparatus.

5.3.1 Slowing and Trapping

The first simulations presented are only of the slowing process, without current in the trapping coils. The simulated hydrogen beam has the same temperature (525 mK) as the metastable neon beam measured in chapter 4. This temperature value is used because the hydrogen is seeded into a neon carrier gas, which dominates the behavior of the resulting beam. The initial velocity used is 527 m/s which matches the measured velocities of the beam. Variations in the real velocity of the initial beam can be accounted for by varying the phase used in the coilgun so that the beam exits

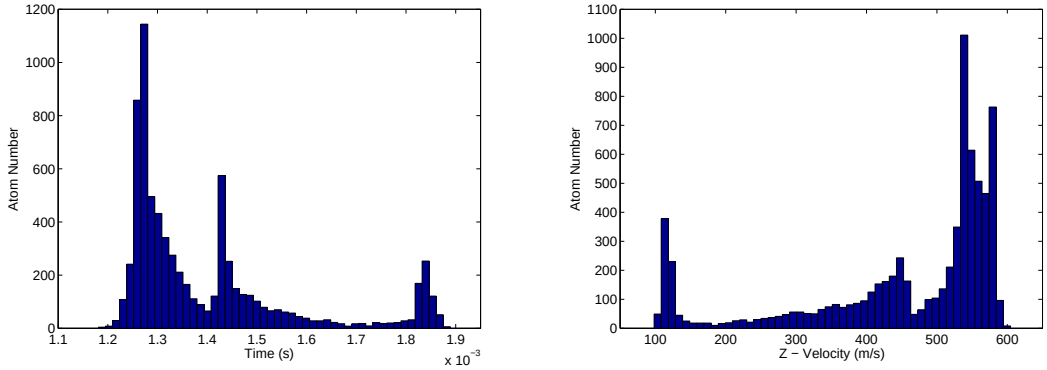

Figure 5.19: Simulated time-of-flight and corresponding velocity distribution for slowed hydrogen. The simulated initial beam parameters were for a 527 m/s velocity and 525 mK temperature. The phase of the simulated coilgun switching is 45.2°. Of the 10,000 simulated atoms entering the coilgun, 579 reach the trap center with a velocity within 10 m/s of the 119 m/s target velocity.

with the correct velocity to be brought to rest at the center of the trap. Figure 5.19 shows the simulated time of flight and velocity distribution of a beam of atoms slowed by the coilgun to a target velocity of 119 m/s, using a coilgun phase of 45.2°. Of the 10,000 simulated atoms that enter the coilgun, 668 arrive at the trap location with a z-velocity less than 130 m/s, and 579 atoms have a final velocity within 10 m/s of the target velocity, giving a slowing efficiency of about 6%.

Trapping of the atoms is simulated with the same initial beam conditions and parameters. The coilgun is used in an identical configuration, and with the same phase as above. The front trapping coil is used as the final slowing coil in this simulation. Since there is no coil spacing which can be used in calculating the phase angle, the position of the synchronous atom relative to the center of the coil when the front coil is switched is used instead. A distance of 4 mm is used as the front coil switching parameter in this simulation. The other parameter that needs to be determined is the time when the front coil is switched back on in anti-Helmholtz orientation. While small variations in this parameter do not greatly change the number of atoms trapped,

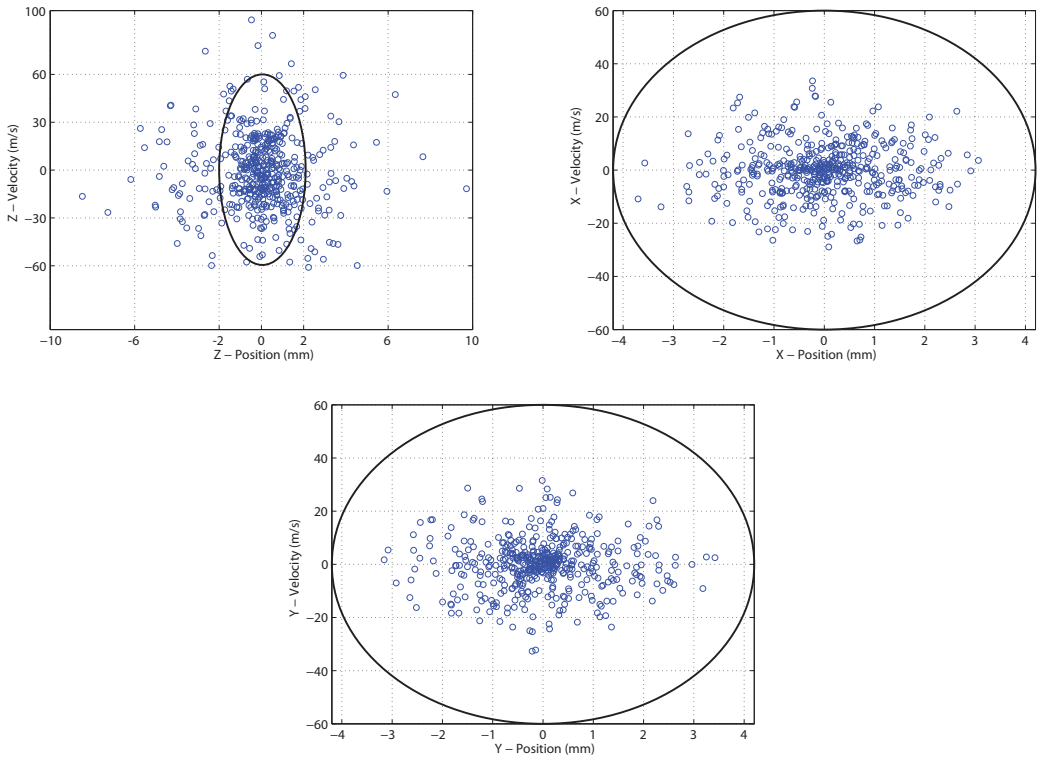

Figure 5.20: Simulated phase space distributions of the trapped hydrogen atoms. Of the 10,000 simulated atoms that enter the coilgun, 481 remain in the trapping volume for 10 ms (around 10 oscillation periods) after the trap is turned on. The temperature of the trapped atoms is calculated to be 62 mK. Included in the figures is an ellipse which encloses the region of phase space which is lower than the overall trap depth. While the X and Y phase space distributions are contained in this region, the Z distribution is not. Atoms with energies in the Z direction that are greater than the overall trap depth can remain in the trap because the trap is significantly deeper in the Z direction, and the velocity components are not well mixed.

the temperature of these atoms is significally affected. Turning the trap on too soon results in increased temperature because the atoms are not yet at rest in the center of the trap. Similarly, turning the trap on too late results in the trapping of atoms that have been reflected from the magnetic field of the rear coil, again raising the temperature of the trapped atoms.

Simulation results presented in figure 5.20 show the phase space distributions of the trapped atoms along each axis of the trap. Of the 10,000 simulated atoms that enter the coilgun, 481 remain in the trap 10 ms after the trap is turned on. This corresponds to about 10 trap oscillation periods, meaning that most of that atoms with unstable trajectories have left the trapping region. The calculated temperature of the trapped atoms is 62 mK. While most of the atoms have energies that are below the overall trap depth (indicated by the ellipses in figure 5.20), some atoms have a larger energy. This is possible because the trap is significantly deeper in the Z direction than in the radial direction, and the velocity components are not well mixed. The trapping efficiency shown in these simulations is significantly higher than the slowing efficieny of neon and oxygen at the lowest velocities, which indicates the benefits of the refinements made to the coilgun.

5.3.2 Ejection Simulations

Detecting trapped hydrogen atoms will probably require them to be ejected from the trap and onto a detector. For atoms which have been excited to a metastable state this could be an MCP; for ground state atoms some kind of ionization could be used to detect the atoms. Given that the atoms must exit the trapping volume (and most likely the trapping chamber itself) for this to work, it is important to consider the best way to get the atoms from the trap to a detector. For the simulations shown here, the detector is assumed to be the Ardara Technologies ionizer and quadrupole,

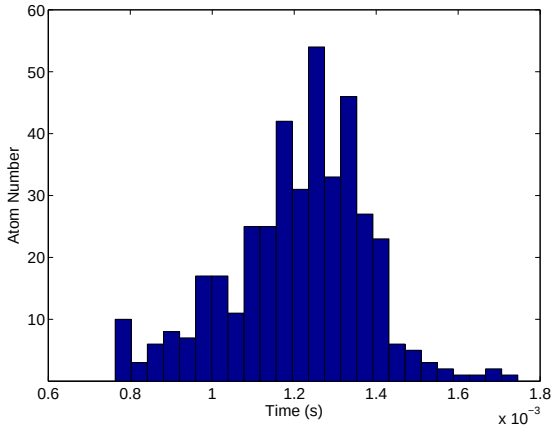

Figure 5.21: Simulated time-of-flight profile for atoms arriving at the Ardara Technologies ionizer after being ejected from the trap. The atoms are ejected by turning off the rear trapping coil, waiting until they have passed through the coil, and then turning the rear trapping coil back on with the opposite polarity. Of the 3,000 simulated trapped atoms, 406 atoms entered the detection region, an efficiency of 13%, which compares favorably with the 0.1% efficiency expected from turning off the trap.

though the calculations described are accurate for any detector at the same location in the experiment.

The simplest method of getting atoms from the trap into the detector is to turn off the trap. The atoms are no longer confined and will thus continue to fly with the velocity they have when they are released. The problem with this method is that since the atoms are moving in all directions in the trap, releasing the atoms will distribute them in all directions. Geometric considerations give a good estimate of the number of atoms that propagate to the detector, approximately 0.1% with the described experimental apparatus. A better method of removing the atoms is to switch off the rear trapping coil, while leaving the front trapping coil on. The atoms are still repelled by the front coil, pushing them through the bore of the rear trapping coil. Once they have passed through the rear coil, it is switched back on with the opposite polarity, continuing to push the atoms towards the detection region. This

sequence was simulated for 3,000 trapped atoms (their trajectories in the trap are determined by previous simulations), and a simulated time-of-flight profile of atoms arriving at the detector is presented in figure 5.21. Of the 3,000 simulated atoms, 406 propagate into the detection volume, giving an efficiency of 13%, which is a significant improvement over releasing the trap.

5.4 Future Directions

While the experiment described above has been constructed and characterized, it has so far been unable to detect trapped hydrogen using the Ardara ionizer and quadrupole. The primary reason for this is thought to be the large background signal due to hydrogen gas in the detection volume. An alternative detection method is to use a laser to excite the trapped hydrogen atoms to a metastable state, and to detect the metastable atoms with an MCP, which should be background free. This detection method and the progress made to date on implementing laser detection in the experiment are described in this section. Once trapped hydrogen has been detected, the initial goal of the experiment is to investigate the isotopic shift of atomic tritium, which should provide information on the charge radius and structure of the triton.

In addition to spectroscopy, research will continue on improving the coilgun method of trapping particles cooled by a supersonic beam. In particular, a method of creating a moving magnetic trap which is decelerated should increase the phase space density of the trapped sample significantly. A moving trap confines the particles in all three dimensions throughout the deceleration and trapping process, thus preserving the initial phase space denisty of the beam at the enterance of the slower.

Finally, a cooling method for magnetically trapped hydrogen isotopes is discussed. The proposed scheme is based on single-photon cooling, which uses a single

scattered photon to extract information and entropy from the trapped ensemble. This method should enable cooling to the recoil limit, increasing the phase space density of the trapped sample. The principle of operation of single-photon cooling is described, as is the proposed implementation in hydrogen. Cooling below the recoil limit would require evaporative cooling, and exploration of a degenerate gas of atomic tritium is a possible research avenue which is also explored.

5.4.1 Laser Detection and Spectroscopy of Hydrogen Isotopes

Detecting trapped hydrogen using laser excitiation should allow for a background free detection method by eliminating the signal from molecular hydrogen. The process which will be used is to excite the trapped hydrogen to the $2S$ metastable state, and then to eject the atoms onto an MCP. Since the atoms start in the $1S$ state, the transition does not change the orbital angular momentum and is forbidden by parity conservation for a single photon. This leads to a long lifetime in the excited state, though it makes the excitation process more complicated.

While one photon cannot drive the $1S-2S$ transition at 121 nm, two photons can excite the atom through a virtual P state. This is illustrated in a simplified energy level diagram shown in figure 5.22. In this case, the first photon excites the atom from the ground state to a virtual state which is a linear combination of all nP states as well as the continuum, and the second photon excites the atom from there to the $2S$ state. This process is discussed in detail elsewhere [106, 107], and so only a brief discussion is given here.

5.4.1.1 Two-Photon Excitation of Hydrogen

While the two photons which drive the $1S-2S$ transition do not need to be of the same wavelength, for reasons which will become clear shortly, it is advantageous

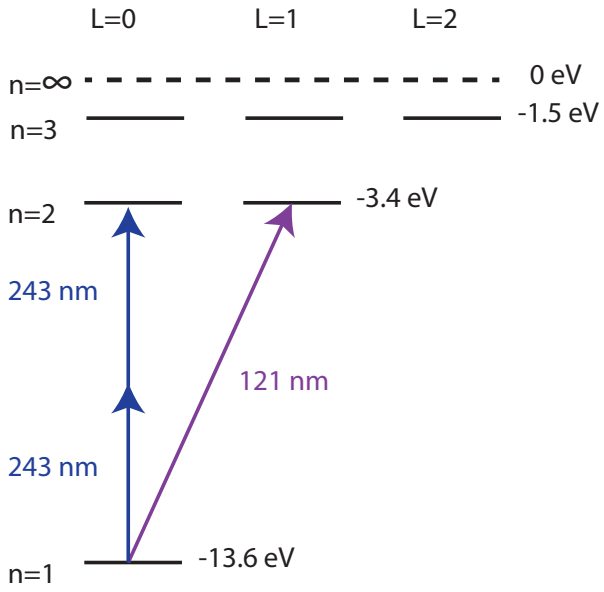

Figure 5.22: A simplified energy level diagram of atomic hydrogen, with the transition wavelengths illustrating the dipole forbidden two-photon $1S-2S$ transition, as well as the dipole allowed $1S-2P$ transition.

both practically (only one laser is needed) and physically (elimination of the Doppler shift) to use a single wavelength. Energy and momentum must be conserved in the process, and so for a resonant transition from a single beam the laser frequency must be

$$2\hbar\omega_{laser} = E_{2S_{\vec{p}_f}} - E_{1S_{\vec{p}_i}} = \hbar\omega_{1S-2S} + \frac{2\hbar\vec{k}\cdot\vec{p}_i}{m} + \frac{2\hbar^2k^2}{m}, \tag{5.1}$$

where $\vec{p}_i$ and $\vec{p}_f$ are the initial and final momentum states of the atom, $\vec{k}$ is the wave vector of the beam, and m is the mass of the hydrogen atom. The term $\frac{2\hbar\vec{k}\cdot\vec{p}_i}{m}$ is referred to as the Doppler shift of the transition, and the $\frac{2\hbar^2k^2}{m}$ term is the recoil shift. The Doppler shift due to the momentum distribution in the trap leads to a broad excitation spectrum and thus a low transition rate, decreasing the number of atoms which will be excited and detected.

Instead of a single beam, it is relatively simple to use two counterpropagating beams from the same laser. Here an atom can absorb one photon from each beam,

where $\vec{k}_1 \cdot \vec{p}_i = -\vec{k}_2 \cdot \vec{p}_i$, where $\vec{k}_1$ and $\vec{k}_2$ are the wave vectors of photons from the two beams. In this scenario, the Doppler shifts cancel each other out. Since the photons are counterpropagating, there is no recoil either. This is known as Doppler free two-photon excitation, and in this regime the laser must satisfy

$$2\hbar\omega_{laser} = E_{2S_{\vec{p}_f}} - E_{1S_{\vec{p}_i}} = \hbar\omega_{1S-2S}. \tag{5.2}$$

This means that a laser at $\lambda = 243$nm is needed to excite the atoms. Assuming linearly polarized light of equal intensity I in the counter propogating beams, and a laser frequency $\omega_{laser} = \omega_{1S-2S}$, the Rabi frequency for Doppler free excitation is [107]

$$\Omega_0(r) = 2M^{12}_{2S,1S}\left(\frac{\alpha}{2R_\infty}\right)^3 \frac{1}{3\pi^2\hbar c} I(r) = 9.264 I(r)\frac{\mathrm{cm}^2}{\mathrm{Ws}}, \tag{5.3}$$

where $M^{12}_{2S,1S}$ is the sum over the two-photon dipole matrix elements, α is the fine structure constant, R_∞ is the Rydberg constant, and $I(r)$ is the spatial intensity distribution. For a weak resonant excitation the transition rate is

$$R = \frac{\Omega_0(r)^2}{\Gamma} = \frac{85.8}{\Gamma} I(r)^2 \frac{\mathrm{cm}^4}{\mathrm{W}^2\mathrm{s}}, \tag{5.4}$$

where Γ is the linewidth of the dominant homogeneous broadening mechanism. In the absence of other broadening mechanisms, the natural linewidth of the transition determined by the lifetime of the $2S$ state, sets $\Gamma = 8.2\mathrm{s}^{-1}$. Thus, without another source of broadening, the saturation intensity of the transition is $I = 0.89\frac{\mathrm{W}}{\mathrm{cm}^2}$. The most important aspect of equation 5.4 is that the rate scales as the square of the intensity of the light, rather than linearly as would be the case for a one-photon transition. This means that the transition rate is very sensitive to the intensity, and that two-photon transitions typically require much higher powers to achieve the desired transition rates.

The limiting homogeneous broadening mechanism for the experiment is likely to be the laser linewidth. This source of broadening is due to the inherent properties

of the laser itself, and the limits in reducing this broadening are discussed along with the planned laser to produce 243 nm light. There are also two important sources of inhomogeneous broadening which will play an important role, both in detecting the atoms, and in setting resolution limits for spectroscopy of hydrogen isotopes in the magnetic trap.

The first source of inhomogeneous broadening is transit time broadening, which results from the time energy uncertainty relationship $\Delta E \Delta t > \hbar$. An atom moving through a laser beam will see the beam for a finite period of time, placing limits on the energy resolution that can be observed. This broadening mechanism is inhomogeneous because the amount of time an atom spends in the beam depends on each atom's trajectory, and so only the average broadening, related to the temperature of the atoms in the sample, can be calculated. The width of the transit time broadening for a sample at a specific temperature will vary as $1/\omega_0$, where ω_0 is the beam waist. This means that decreasing the beam waist will increase the transit time broadening, lowering the transition rate.

Given a fixed power in the laser, decreasing the waist will increase intensity, making the transition rate go as ω_0^{-4}, while transit time broadening affects the rate as ω_0. The number of trapped atoms which pass through the beam also goes as ω_0, leading to an overall dependance of ω_0^{-2}. One technique for obtaining maximum energy resolution for spectroscopy is to balance the laser linewidth broadening with the transit time broading by picking a beam waist that makes their contributions equal.

The final source of broadening that must be considered is the inhomogeneous broadening due to the magnetic field of the trapping potential. While the Zeeman shift of the $1S$ and $2S$ states both scale linearly with the magnetic field, there are relativistic effects which introduce broadening. Specifically, there is a relativistic correction to

the electron g-factor which causes the resonance energy of the the $1S-2S$ transition to vary with magnetic field. The g-factor varies with the principle quantum number n as [108]

$$g_e(n) = g_e\left(1 - \frac{\alpha^2}{3n^2}\right). \tag{5.5}$$

This leads to a frequency shift of [109]

$$\delta\nu = 186\text{kHz/T}. \tag{5.6}$$

In the quadrupole trap, this corresponds to shift of up to 65 kHz, which will limit the resolution of any spectroscopy performed in the trap.

5.4.1.2 The Hydrogen Laser

The laser which will be used to excite the $1S-2S$ transition is currently under construction. The original design was based on the laser described in [110]. To create light at 243 nm, a diode laser at 972 nm was amplified and frequency doubled twice. This laser was found not to produce enough power at 243 nm to efficiently drive the transition, and the front end of the laser is currently being replaced with an optically pumped semiconductor laser (OPSL) at 972 nm which should enable more UV power to be generated.

A schematic of the original laser system constructed by Travis Bannerman is shown in figure 5.23. A tunable extended cavity diode laser (ECDL) is used to produce single-mode light at 972 nm. Because the laser frequency must be quadrupled, and the eventual excitation is a two-photon process, phase noise in the laser must be reduced as much as possible. A diffraction grating mounted 27 cm from the diode creates a long cavity with acts as a flywheel, reducing high-frequency phase noise. The frequency of the laser can be tuned by changing the angle of the diffraction grating with a piezo stack, and the frequency and single mode operation of the laser

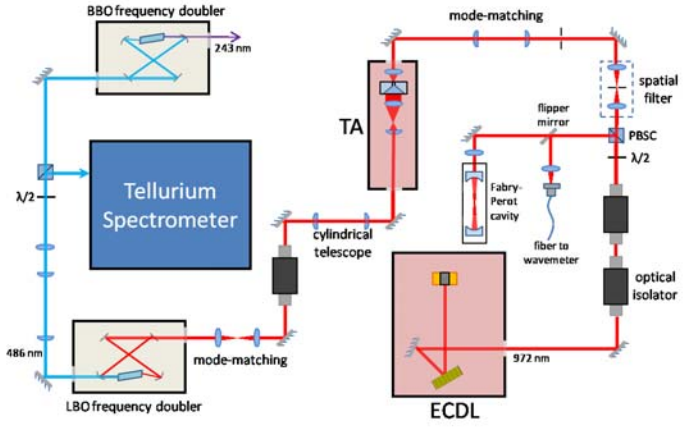

Figure 5.23: A schematic of the original hydrogen laser. A tunable extended cavity diode laser (ECDL) produces single-mode light at 972 nm which is verified by a Fabry-Perot cavity and wavemeter. The light from the diode is amplified in a tapered amplifier (TA) and injected into a doubling cavity. After doubling the light is at 486 nm, where Tellurium spectroscopy is performed to lock the laser. The blue light is injected into another doubling cavity, which produces 243nm light. Figure Courtesy Travis Bannerman.

are verified by a wavemeter and Fabry-Perot cavity respectively. Since the output of the diode laser (Eagleyard Photonics RWE-0980-08020-1500-SOT02) is limited to about 70 mW and the doubling processes are nonlinear, where the efficiency of the doubling goes as the square of the power, it is necessary to amplify the laser before injecting it into the first doubling cavity. This is done using a tapered amplifier (TA). The TA (Toptica Photonics BoosTA-L-980) takes approximately 35 mW of 972 nm light (after isolation, diagnostics, spatial filtering, and mode-matching), and amplifies it to approximately 1 W, which results in approximately 900 mW of power which can be injected into the first doubling cavity after isolation and mode-matching.

The doubling cavity is a bow-tie cavity locked to the frequency of the diode laser. A nonlinear lithium triborate (LBO) crystal doubles the frequency of the light to 486 nm, which is coupled out of the cavity by a dichroic mirror. The doubling cavity (Coherent MBD-200) produces around 100 mW of light at 486 nm. This light

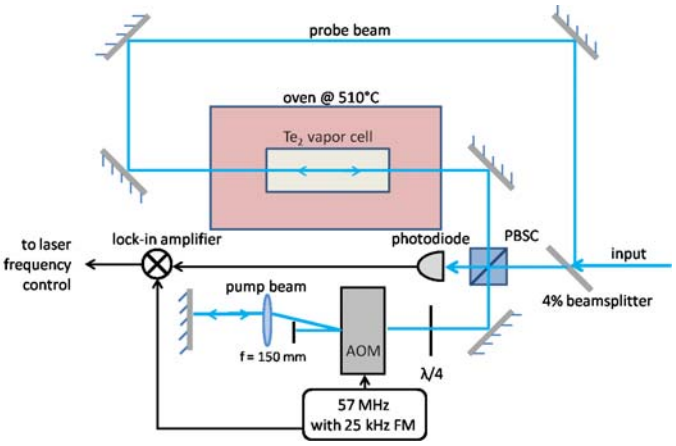

Figure 5.24: A schematic of the tellurium saturated absorption spectroscopy setup. The frequency of the pump beam is offset by double passing through a 57 MHz AOM that compensates for the frequency difference between the line in tellurium and the 1/4 harmonic of the hydrogen line. Figure Courtesy Travis Bannerman.

is mode-matched to a second doubling cavity, which is identical to the first except that it uses a beta barium borate (BBO) crystal as the nonlinear medium instead of LBO. The BBO crystal is used because of its optical properties in the UV. The second doubling cavity produces approxiately 2 mW of light at 243 nm.

Between the first and second doublers, approximately 10 mW of 486 nm light is picked out of the beam and used in a saturated absorption spectroscopy setup to lock the diode laser frequency. More detail on saturated absorption spectroscopy can be found in [111]. Molecular tellurium is used as a frequency reference for the diode since the lines around 486 nm have been mapped to better than 1 MHz [112–114], and the i_2 line has been found to be approximately 57 MHz above the 1/4 harmonic of the $1S-2S$ transition [115]. This frequency difference can be compensated for by double passing the pump beam of the saturated absorption setup through an acousto-optic modulator (AOM), as shown in figure 5.24. The spectroscopy signal is fed back to the piezo on the diffraction grating, locking the laser frequency. More detail on the laser setup can be found in [116].

Since the output power of the laser described above is only 2 mW, the $1S-2S$ transition cannot be driven effectively. Two steps are being taken to increase the power. The first step is the addition of a build up cavity around the atoms. This should allow the UV power to be increased by a factor of around 50. The second step is to replace the ECDL and TA of the current system with an OPSL laser [117, 118]. The OPSL has an output power of over 1.3 W and the mode quality is much better than that of the TA. This should enable more light to be coupled into the first doubler, hopefully increasing the output power at 486 nm to over 500 mW. This much blue power should increase the efficiency of the second doubler, with the goal of having 50 mW of power at 243 nm. When used in conjunction with the build up cavity, the atoms should see 2.5 W of 243 nm light. Calculations performed by Robert Clark indicate that this much light should be sufficient to excite a few percent of the atoms in the trap, and enabling the laser to be used as a background free detection method. Finally, linewidth of the OPSL (and thus the laser linewidth broadening) can be reduced as needed by locking the laser to an external cavity.

5.4.1.3 Spectroscopic Goals

While the laser will initially be used to detect the atoms in the trap, it can also be used to perform spectroscopy on the trapped hydrogen. Of particular importance is the isotopic shift of the $1S-2S$ transition in tritium. The current uncertainty on this measurement is several MHz [119]. By lowering the uncertainty to a few hundred kHz, the size of the triton charge radius can be measured, which is of great interest to nuclear theorists [120]. There is no reason why the linewidth of the OPSL should not be able to narrowed to below this uncertainty, and the waist of the beam in the cavity can be set so that the time of flight broadening is equivalent to the laser linewidth. Since the Zeeman broadening in the trap is only a few tens of kHz, this should not

prevent the determination of the isotopic shift of the $1S-2S$ transition in tritium to around 100 kHz with the apparatus described above.

5.4.2 An Adiabatic Coilgun

Beyond the problem of detection, research has continued on possible improvements to the methods used to bring the cold samples produced by supersonic expansion to rest in the lab frame. The most promising avenue for increasing the flux and phase space density of the final sample is to trap the sample and then decelerate it, as opposed to the current design which first slows and then traps the atoms. By trapping the atoms in a moving magnetic trap, they are confined in all three dimensions even though they are still moving quickly in the lab frame. Adiabatically decelerating the moving trap will then bring the trapped particles to rest without decreasing the phase space density, thereby allowing the high phase space densities produced by the initial supersonic expansion to be maintained through the deceleration process. This idea is proposed in [68].

There are several approaches to creating a moving trap which are appropriate for the use in an adiabatic decelerator. Moving magnetic traps have been used to transport atoms in cold atom experiments, both on microchips [121] and using macroscopic coils [122], inspiring the concept of a moving magnetic trap decelerator, though far greater velocities are needed to trap a supersonic beam in the co-moving frame. The moving and decelerated trap concept has also been applied to moving electrostatic traps for chips [123, 124] and for macroscopic electrodes [125], and these designs have been able to capture and slow molecules in a supersonic beam. A moving, though not decelerated magnetic trap has also been used to capture and translate metastable argon atoms in a supersonic beam [126]. An adiabatic magnetic decelerator has been used to stop metastable neon atoms [127, 128]. These last experiments

show that the acceptance volume in phase space of a magnetic adiabatic decelerator is well matched to the phase space volume of a supersonic beam after the skimmer. This means that the adiabatic decelerator should capture the greatest possible flux at the highest possible phase space density. The configuration used in these last experiments is the one which will be implemented here.

5.4.3 Single-Photon Cooling of Hydrogen

The coilgun method can bring samples of atoms to rest in the laboratory frame, where they can then be trapped at temperatures in the tens of millikelvin. For some applications though, the trapped atoms need to be colder than this. For true precision spectroscopy, trapping the atoms in a magic wavelength optical trap would be ideal [129], and this requires cooling the atoms to load them into the shallower optical potential. Studies of degenerate gases also require further cooling, since efficient evaporative cooling to degeneracy requires a high initial phase space density to enable collisions and short rethermalization times. This section presents a general cooling technique called single-photon cooling, which is capable of cooling nearly any species to the recoil limit [130–134]. A proposed scheme for applying single-photon cooling to atomic hydrogen isotopes [81, 135] is discussed.

The principle of single-photon cooling is the creation of a one-way wall, which is permeable to atoms approaching the barrier from one direction but not from the other. A one-way wall permits a sample to be compressed without doing work on it, thus increasing the phase space density. This is illustrated for atoms trapped in a 1D potential in figure 5.25(a). While the wall increases the phase space density, the temperature of the sample is not reduced. Figure 5.25(b) illustrates how starting the wall at the edge of the trap and sweeping it slowly through the potential allows the wall to both compress the sample and reduce its temperature. In this example, the

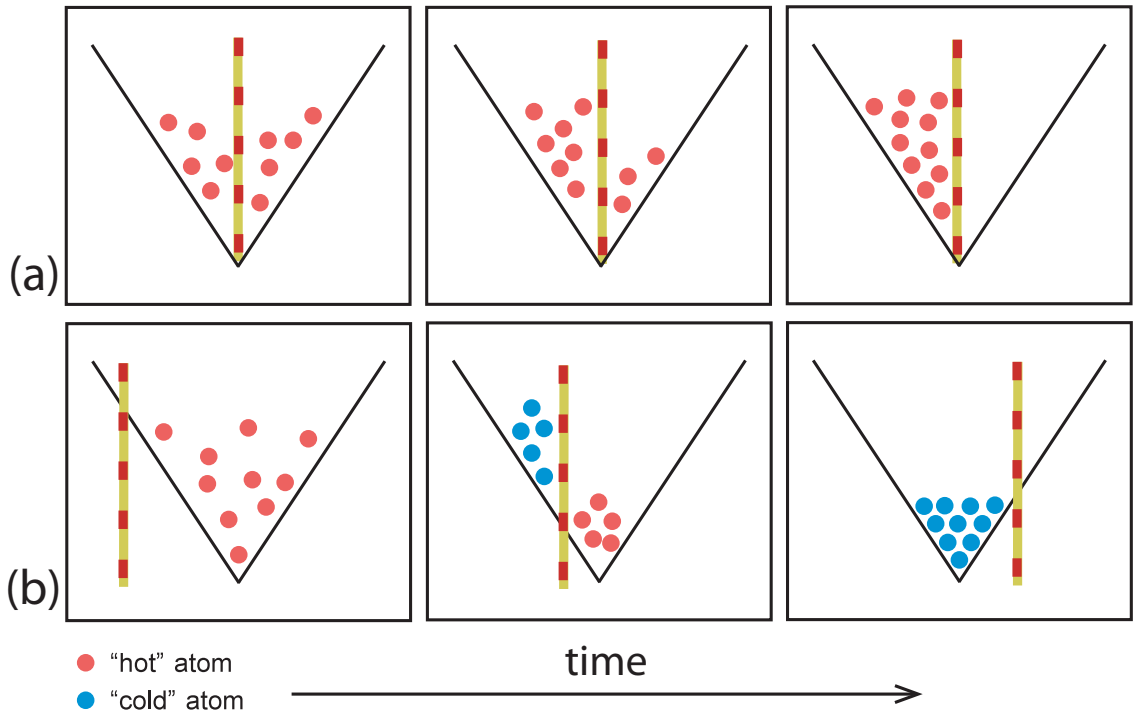

Figure 5.25: A 1D illustration of how a one-way call can be used to cool and compress phase space. In (a) the one-way wall is placed at the center of the trap. Atoms can pass through the wall in one direction, but not the other, compressing the atoms and increasing the density without doing work or heating the atoms. In (b) the wall is swept through the 1D trapping potential, starting from the edge of the trap. By sweeping the wall slowly, each atom encounters the wall near its classical turning point, resulting in their having low kinetic energy when they pass through the wall and are trapped by the wall and the original potential. By sweeping the wall though the entire sample, each atom has its energy reduced and the entire trap is cooled.

wall starts at the edge of the trap, and the atoms do not encounter the wall during their oscillations. To cool the atoms, the wall is slowly swept towards the center of the trap. By moving the wall slowly, each atom in the trap will encounter the one-way wall close to the classical limit of its trajectory in the trap. Since the atoms are near their classical turning points, they have very little kinetic energy when they interact with the barrier. Encountering the one-way barrier, the atoms pass through the wall and are trapped near the minimum of the potential formed by the one-way wall and the original trap. In this manner the atoms have been transferred from a high potential energy state in one trap, to a low potential energy state in a different trap. By sweeping the wall adiabatically, atoms trapped by the one-way barrier are not heated by the sweep, and the entire 1D trap can be cooled. The one-way wall only cools the atoms in one dimension, even when placed in a three dimensional trap, unless trap ergodicity mixes the degrees of freedom sufficiently for each atom to encounter the wall at nearly zero kinetic energy.

The experimental technique used to create a one-way wall is a laser that switches the internal state of atoms that encounter it, such that the trapping potential in the new state is changed. One important point though is that the state change must involve spontaneous emission to make the process irreversible. Single-photon cooling takes its name from the fact that changing the state of the atoms, and thus the potential landscape they see, requires each atom to scatter just one photon. A closed two level system, as is required for traditional laser cooling, is not required for single-photon cooling, making the technique broadly applicable. In fact, single-photon cooling will not work in a true two level system, since the atom will return to its orignal state after scattering the photon, and the process will not be irreversible.

Single-photon cooling is a physical realization of the "Maxwell's Demon" thought

experiment, where a being is capable of observing the motion of each particle in an ensemble and, using these observations, is able to reduce the entropy of the ensemble with doing work on it, apparently violating the Second Law of Thermodynamics. The classic example of this is a trap door, operated by the observer, which separates two chambers of gas. If the observer opens the trap door when a particle approaches from the left side, but not the right, then the observer can create a one-way wall using the door, and atoms will accumulate in the chamber on the right. This increases the phase space density without doing work. The resolution to this apparent paradox lies in the fact that information carries entropy. To operate the trap door, the observer must store information about the ensemble, which must eventually be erased, increasing the entropy of the universe in accordance with Landauer's erasure principle [136]. For single-photon cooling, the photon scattered each time an atom changes state provides information about that particular atom's position in the trap, making a measurement of the original state of the ensemble possible. In this manner, the entropy removed from the atomic ensemble is accounted for in the increased entropy of the radiation field [133].

A one-way wall for atoms has been demonstrated for optically trapped rubidium, in analog to the situation illustrated in 5.25(a) [137]. This experiment compressed the sample, but due to the heating from spontaneous scattering, the phase space density of the sample was only increased by a factor of 1.07. A different scheme, in which rubidium atoms are transferred from a magnetic trap to a gravito-optical trap has also been used for single-photon cooling [138, 139]. Using laser excitation and spontaneous emission to affect the irreversible state change that creates the effect of a one-way wall, this method succeeded in increasing the phase space density of the trapped sample by a factor of 350. More details on these experiments can be found in [116, 140, 141].

The one drawback of the above implementations of single-photon cooling is that the cooling only operates in one dimension, and cooling in more than one dimension requires trap ergodicity or collisions between the atoms to provide mixing of the degrees of freedom. Cooling in three dimensions requires the one-way wall and the new trapping potential to form a shell around the trap the atoms are initially held in. While this is not practical with an optical potential, it is possible to create these potentials using static and radio frequency (RF) or microwave magnetic fields. The RF field couples the internal states of the atom and shifts their energies [142, 143], which can create a conservative potential. The atoms are considered to be "dressed" by the RF photons, and the resulting states are referred to as "RF-dressed states". RF-dressed potentials have been proposed [144] and demonstrated [145] for trapping of neutral atoms, and have permitted the creation of novel trap geometries [146–148]. The effect on the trapped atoms of the interaction between the RF field and the trap has been analyzed in detail [149–151]. Single-photon cooling of atoms and molecules using population transfer between RF-dressed states of a magnetic trap was first considered and proposed in [152]. It is this technique which will be adapted for use with hydrogen isotopes, though the current trap will need to be changed to one with a bias field, such as a Ioffe-Pritchard magnetic trap [153].

An illustration of the proposed implementation of RF-dressed single-photon cooling for atomic hydrogen isotopes is shown in figure 5.26. For simplicity, only two states of the $1S$ manifold are considered in the following discussion. Atoms which are in the lower dressed state (labeled $|-\rangle$ in figure 5.26) are low-field-seeking at low magnetic field, and high-field-seeking at high magnetic field, with the field strength of the cross-over in behavior defined by the frequency of the RF magnetic field. Similarly, atoms in the upper dressed state(labeled $|+\rangle$) are high-field-seeking at low field, and low-field-seeking at high field, with the shift in behavior occuring

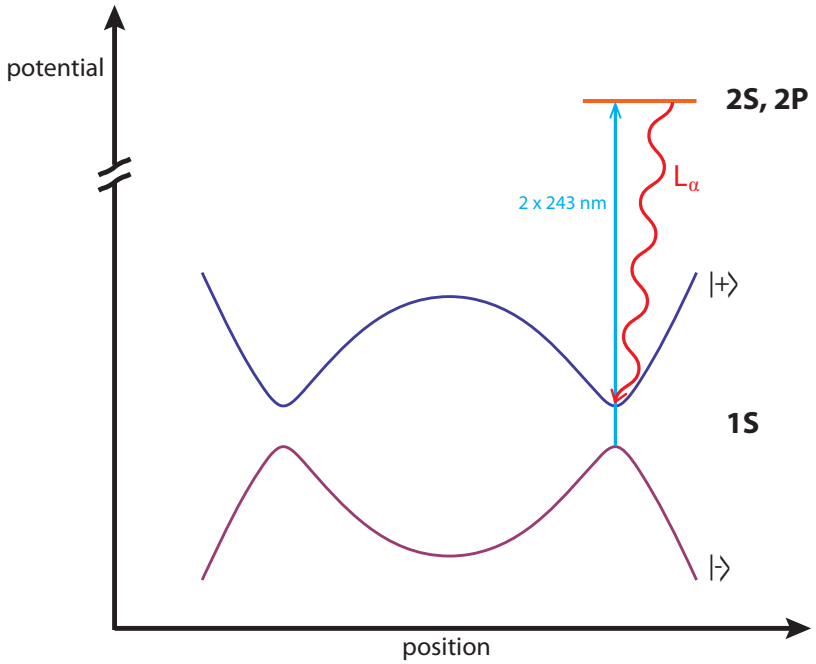

Figure 5.26: A scheme for implementing RF-dressed single-photon cooling of atomic hydrogen isotopes in the high field limit. Two of the four levels of the $1S$ manifold are depicted, labeled $|+\rangle$ and $|-\rangle$. These levels are coupled by RF magnetic fields, which shift the "bare" states and create avoided crossings. Atoms which are originally trapped in the $|-\rangle$ state are excited to the $2S$ level at their classical turning point using a two photon transition at 243 nm. Stark mixing of the $2S$ and $2P$ states causes the atoms to quickly decay via spontaneous emission of a Lyman α photon. The atoms that decay into the $|+\rangle$ state are trapped with nearly zero kinetic energy, with the probablitiy of reaching the $|+\rangle$ state determined by the branching ratio.

at the same field strength. This change in behavior is due to an avoided crossing of the levels. By sweeping the RF frequency, hydrogen atoms can be pumped at their classical turning point from the $|-\rangle$ dressed state to the $2S$ manifold by 243 nm light. The lifetime of the $2S$ state is 0.122 s, which is too long for single photon cooling to operate effectively. The $2S$ state can, however, be Stark mixed with the $2P$ state, resulting in rapid decay back to the $1S$ manifold via spontaneous emission of a Lyman α photon [154]. The branching ratio determines the probablity of the atom decaying to the $|+\rangle$ dressed state where it is trapped with nearly zero kinetic energy. The process is irreversible because the 243 nm light will not be on resonance for an atom in the $|+\rangle$ dressed state. The scattered Lyman α photon sets the fundamental limit of the cooling, since this scattered photon gives a random momentum kick. With a mass of 1.67×10^{-27} kg for hydrogen, and the Lyman α photon having 5.5×10^{-27} kg m/s momentum, the recoil limit works out to be 1.3 mK.

Once the atoms have been cooled to the recoil limit, there are several avenues of research which can be pursued. Spectroscopy of the $1S - 2S$ transition in tritium would provide information on the charge radius of the triton [120]. Cold trapped tritium has been proposed as a source for studies of bound state beta-decay [155, 156], and possible determination of the neutrino rest mass [157]. It has also been predicted that a degenerate gas of tritium may exhibit novel liquid like behavior [158–161]. Finally, cold trapped samples at high phase space density may enable the study of the quantum nature of chemical reactions at cold temperatures [162].

Appendix

Appendix A

Reflection From a Spinning Rotor

The equations of motion for an atom reflecting from the rotor are more complicated than the 1D simplification given in chapter 3. In the experiment, the atoms can hit the crystal at an angle, and the velocity of the crystal may not be square to the normal of the crystal face.

The key to finding the equations that define the reflection of an atom from a moving crystal is to move into the frame of the crystal. In this reference frame, the speed of the atom is unchanged by the reflection, and the angle of incidence is equal to the angle of reflection. The solution to the problem can thus be found finding the velocity of the atom in the frame of the crystal, reflecting the atom from the apparently stationary crystal, and transforming the velocity of the atom back into the laboratory frame.

The geometry used in deriving these equation is shown in figure A.1. The atom starts with an initial velocity

$$V_i = [V_{x_i}, V_{x_i}] \tag{A.1}$$

and in the geometry chosen, the crystal has velocity

$$V_c = \left[r\dot{\theta}\sin\theta, r\dot{\theta}\cos\theta \right], \tag{A.2}$$

where r is the radius of the rotor. Thus in the frame of the rotor, the atom has a velocity of

$$V_{i-c} = \left[V_{x_i} - r\dot{\theta}\sin\theta, V_{x_i} - r\dot{\theta}\cos\theta \right]. \tag{A.3}$$

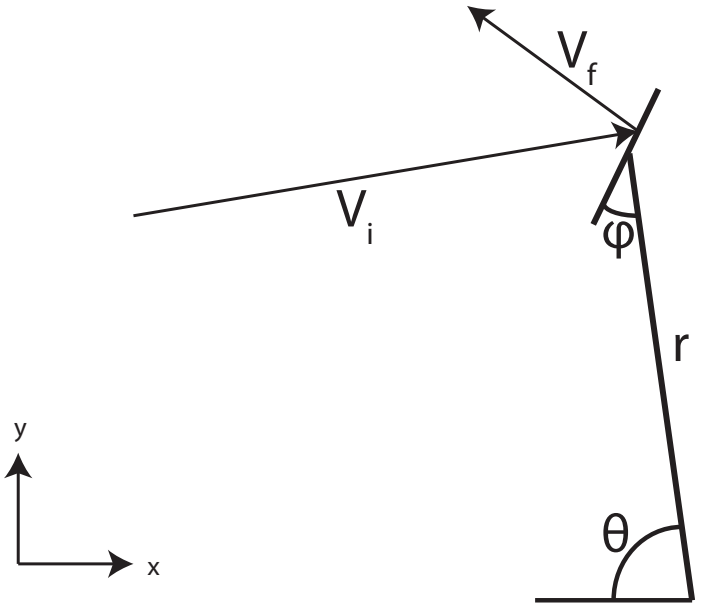

Figure A.1: The geometry used for calculating the reflection of an atom from the tip of the rotor.

Since the situation is now one of an atom reflecting from an apparently stationary object, the initial and final speeds of the atom will be equal in this frame, and the trajectory of the atom will be reflected about the normal vector of the crystal $\hat{n}$. This calculation can be done quite simply in vector form and the final velocity is

$$V_{out} = V_{in} - 2\left(V_{in} \cdot \hat{n}\right)\hat{n}. \tag{A.4}$$

In the chosen geometry, the normal vector of the crystal is

$$\hat{n} = \left[\sin\left(\theta + \phi\right), \cos\left(\theta + \phi\right)\right]. \tag{A.5}$$

Inserting equations A.3 and A.5 into equation A.4 and simplifying gives the reflected velocity in the frame of the crystal,

$$\begin{aligned} V_{f-c} = \quad & [V_{x_i}\cos\left(2(\theta+\phi)\right) - V_{y_i}\sin\left(2(\theta+\phi)\right) + r\dot{\theta}\sin\left(\theta+2\phi\right), \\ & -V_{y_i}\cos\left(2(\theta+\phi)\right) - V_{x_i}\sin\left(2(\theta+\phi)\right) + r\dot{\theta}\cos\left(\theta+2\phi\right)]. \end{aligned} \tag{A.6}$$

Returning to the laboratory frame gives the final velocity of the atom,

$$\begin{aligned} V_f = \quad & [V_{x_i} \cos\left(2(\theta+\phi)\right) + 2\left(r\dot{\theta}\cos\phi - V_{y_i}\cos\left(\theta+\phi\right)\right)\sin\left(\theta+\phi\right), \\ & 2r\dot{\theta}\cos\phi\cos\left(\theta+\phi\right) - V_{y_i}\cos\left(2(\theta+\phi)\right) - V_{x_i}\sin\left(2(\theta+\phi)\right)]. \end{aligned} \tag{A.7}$$

Bibliography

[1] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer-Verlag, New York, 1999.

[2] William D. Phillips and Harold Metcalf. Laser deceleration of an atomic beam. *Phys. Rev. Lett.*, 48(9):596–599, Mar 1982.

[3] E. L. Raab, M. Prentiss, Alex Cable, Steven Chu, and D. E. Pritchard. Trapping of neutral sodium atoms with radiation pressure. *Phys. Rev. Lett.*, 59:2631–2634, Dec 1987.

[4] John M. Doyle, Bretislav Friedrich, Jinha Kim, and David Patterson. Buffer-gas loading of atoms and molecules into a magnetic trap. *Phys. Rev. A*, 52(4):R2515–R2518, Oct 1995.

[5] Jonathan D. Weinstein, Robert deCarvalho, Thierry Guillet, Bretislav Friedrich, and John M. Doyle. Magnetic trapping of calcium monohydride molecules at millikelvin temperatures. *Nature*, 395:148–150, Sep 1998.

[6] Cindy I. Hancox, S. Charles Doret, Matthew T. Hummon, Linjiao Luo, and John M. Doyle. Magnetic trapping of rare-earth atoms at millikelvin temperatures. *Nature*, 431:281–284, Sep 2004.

[7] Wesley C. Campbell, Edem Tsikata, Hsin-I Lu, Laurens D. van Buuren, and John M. Doyle. Magnetic trapping and zeeman relaxation of nh ($x\,^3\sigma^-$). *Phys. Rev. Lett.*, 98(21):213001, May 2007.

[8] H. Pauly. *Atom, Molecule and Clusterbeams I: Basic Theory, Production and Detection of Thermal Energy Beams*. Springer-Verlag, Berlin, 2000.

[9] R. Campargue, editor. *Atom and Molecular Beams: The State of the Art 2000*. Springer-Verlag, Berlin, 2001.

[10] G. Scoles, editor. *Atomic and Molecular Beam Methods*. Oxford University Press, New York, 2000.

[11] Manish Gupta and Dudley Herschbach. Slowing and speeding molecular beams by means of a rapidly rotating source. *The Journal of Physical Chemistry A*, 105(9):1626–1637, 2001.

[12] Bodil Holst and William Allison. An atom-focusing mirror. *Nature*, 390:244, Nov 1997.

[13] Michael S. Elioff, James J. Valentini, and David W. Chandler. Subkelvin cooling no molecules via "billiard-like" collisions with argon. *Science*, 302(5652):1940–1943, 2003.

[14] R. Fulton, A. I. Bishop, and P. F. Barker. Optical stark decelerator for molecules. *Phys. Rev. Lett.*, 93(24):243004, Dec 2004.

[15] Hendrick L. Bethlem, Giel Berden, and Gerard Meijer. Decelerating neutral dipolar molecules. *Phys. Rev. Lett.*, 83(8):1558–1561, Aug 1999.

[16] Sebastiaan Y. T. van de Meerakker, Paul H. M. Smeets, Nicolas Vanhaecke, Rienk T. Jongma, and Gerard Meijer. Deceleration and electrostatic trapping of oh radicals. *Phys. Rev. Lett.*, 94(2):023004, Jan 2005.

[17] Brian C. Sawyer, Benjamin L. Lev, Eric R. Hudson, Benjamin K. Stuhl, Manuel Lara, John L. Bohn, and Jun Ye. Magnetoelectrostatic trapping of ground state oh molecules. *Phys. Rev. Lett.*, 98(25):253002, Jun 2007.

[18] E. Vliegen, H. J. Wörner, T. P. Softley, and F. Merkt. Nonhydrogenic effects in the deceleration of rydberg atoms in inhomogeneous electric fields. *Phys. Rev. Lett.*, 92:033005, Jan 2004.

[19] E. Vliegen, S. D. Hogan, H. Schmutz, and F. Merkt. Stark deceleration and trapping of hydrogen rydberg atoms. *Phys. Rev. A*, 76:023405, Aug 2007.

[20] S. D. Hogan, P. Allmendinger, H. Saßmannshausen, H. Schmutz, and F. Merkt. Surface-electrode rydberg-stark decelerator. *Phys. Rev. Lett.*, 108:063008, Feb 2012.

[21] Nicolas Vanhaecke, Urban Meier, Markus Andrist, Beat H. Meier, and Frédéric Merkt. Multistage zeeman deceleration of hydrogen atoms. *Phys. Rev. A*, 75(3):031402, Mar 2007.

[22] S. D. Hogan, D. Sprecher, M. Andrist, N. Vanhaecke, and F. Merkt. Zeeman deceleration of h and d. *Phys. Rev. A*, 76(2):023412, Aug 2007.

[23] S. D. Hogan, A. W. Wiederkehr, H. Schmutz, and F. Merkt. Magnetic trapping of hydrogen after multistage zeeman deceleration. *Phys. Rev. Lett.*, 101(14):143001, Sep 2008.

[24] A. W. Wiederkehr, S. D. Hogan, B. Lambillotte, M. Andrist, H. Schmutz, J. Agner, Y. Salathé, and F. Merkt. Trapping deuterium atoms. *Phys. Rev. A*, 81(2):021402, Feb 2010.

[25] A. W. Wiederkehr, S. D. Hogan, and F. Merkt. Phase stability in a multistage zeeman decelerator. *Phys. Rev. A*, 82(4):043428, Oct 2010.

[26] Alex W. Wiederkehr, Michael Motsch, Stephen D. Hogan, Markus Andrist, Hansjürg Schmutz, Bruno Lambillotte, Josef A. Agner, and Frédéric Merkt. Multistage zeeman deceleration of metastable neon. *The Journal of Chemical Physics*, 135(21):214202, 2011.

[27] J.S. Wu, S.Y. Chou, U.M. Lee, Y.L. Shao, and Y.Y. Lian. Parallel dsmc simulation of a single under-expanded free orifice jet from transition to near-continuum regime. *Journal of fluids engineering*, 127:1161, 2005.

[28] Bronshtein Gonen. Limit of relative velocities for cluster onset in intense supersonic beams. Master's thesis, Tel Aviv University, Tel Aviv, Israel, 2006.

[29] Otto F. Hagena. Nucleation and growth of clusters in expanding nozzle flows. *Surface Science*, 106(13):101 – 116, 1981.

[30] Otto F. Hagena. Cluster ion sources (invited). *Review of Scientific Instruments*, 63(4):2374–2379, 1992.

[31] U. Even, J. Jortner, D. Noy, N. Lavie, and C. Cossart-Magos. Cooling of large molecules below 1 k and he clusters formation. *The Journal of Chemical Physics*, 112(18):8068–8071, 2000.

[32] Matthias Hillenkamp, Sharon Keinan, and Uzi Even. Condensation limited cooling in supersonic expansions. *The Journal of Chemical Physics*, 118(19):8699–8705, 2003.

[33] A Libson, M Riedel, G Bronshtein, E Narevicius, U Even, and M G Raizen. Towards coherent control of supersonic beams: a new approach to atom optics. *New Journal of Physics*, 8(5):77, 2006.

[34] R. Campargue. Progress in overexpanded supersonic jets and skimmed molecular beams in free-jet zones of silence. *The Journal of Physical Chemistry*, 88(20):4466–4474, 1984.

[35] Wieland Schollkopf and J. Peter Toennies. The nondestructive detection of the helium dimer and trimer. *The Journal of Chemical Physics*, 104(3):1155–1158, 1996.

[36] L. Pedemonte and G. Bracco. Study of he flow properties to test he dimer potentials. *The Journal of chemical physics*, 119:1433, 2003.

[37] Uzi Even. Private communication, 2006.

[38] M.-O. Mewes, M. R. Andrews, D. M. Kurn, D. S. Durfee, C. G. Townsend, and W. Ketterle. Output coupler for bose-einstein condensed atoms. *Phys. Rev. Lett.*, 78:582–585, Jan 1997.

[39] Immanuel Bloch, Theodor W. Hänsch, and Tilman Esslinger. Atom laser with a cw output coupler. *Phys. Rev. Lett.*, 82:3008–3011, Apr 1999.

[40] J.R. Buckland, B. Holst, and W. Allison. Helium reflectivity of the si(111)-(11) h surface for use in atom optical elements. *Chemical Physics Letters*, 303(12):107 – 110, 1999.

[41] R.B. Doak, K. Kevern, A. Chizmeshya, R. David, and G. Comsa. Dynamic atom optics to produce ultraslow, ultracold helium atoms: design study and possible applications. In *Proceedings of SPIE*, volume 2995, page 146, 1997.

[42] E. Narevicius, A. Libson, M. F. Riedel, C. G. Parthey, I. Chavez, U. Even, and M. G. Raizen. Coherent slowing of a supersonic beam with an atomic paddle. *Phys. Rev. Lett.*, 98(10):103201, Mar 2007.

[43] Max Riedel. Elastic slowing of supersonic beams. Master's thesis, The University of Texas at Austin, Austin, TX, 2006.

[44] U. Bonse and H. Rauch. *Neutron Interferometry*. Oxford Science Publications, Oxford, 1979.

[45] Justin Rorke Buckland. *New Methods in Helium Atom Scattering*. PhD thesis, Cambridge University, Cambridge, England, 1998.

[46] I. Estermann and O. Stern. Beugung von molekularstrahlen. *Zeitschrift fr Physik A Hadrons and Nuclei*, 61:95–125, 1930. 10.1007/BF01340293.

[47] D. J. Riley, A. P. Jardine, S. Dworski, G. Alexandrowicz, P. Fouquet, J. Ellis, and W. Allison. A refined he–lif(001) potential from selective adsorption resonances measured with high-resolution helium spin-echo spectroscopy. *The Journal of Chemical Physics*, 126(10):104702, 2007.

[48] D. W. Keith, M. L. Schattenburg, Henry I. Smith, and D. E. Pritchard. Diffraction of atoms by a transmission grating. *Phys. Rev. Lett.*, 61:1580–1583, Oct 1988.

[49] David W. Keith, Christopher R. Ekstrom, Quentin A. Turchette, and David E. Pritchard. An interferometer for atoms. *Phys. Rev. Lett.*, 66:2693–2696, May 1991.

[50] Michael S. Chapman, Christopher R. Ekstrom, Troy D. Hammond, Richard A. Rubenstein, Jörg Schmiedmayer, Stefan Wehinger, and David E. Pritchard. Optics and interferometry with na_2 molecules. *Phys. Rev. Lett.*, 74(24):4783–4786, Jun 1995.

[51] O. Carnal, M. Sigel, T. Sleator, H. Takuma, and J. Mlynek. Imaging and focusing of atoms by a fresnel zone plate. *Phys. Rev. Lett.*, 67:3231–3234, Dec 1991.

[52] R. B. Doak, R. E. Grisenti, S. Rehbein, G. Schmahl, J. P. Toennies, and Ch. Wöll. Towards realization of an atomic de broglie microscope: Helium atom focusing using fresnel zone plates. *Phys. Rev. Lett.*, 83(21):4229–4232, Nov 1999.

[53] JR Buckland and W. Allison. Determination of the helium/si (111)–(1× 1) h potential. *The Journal of Chemical Physics*, 112:970, 2000.

[54] D. Farias and K.H. Rieder. Atomic beam diffraction from solid surfaces. *Reports on Progress in Physics*, 61:1575, 1998.

[55] Christoph J. Schaeff. Elastic slowing of supersonic beams with an atomic mirror. Master's thesis, The University of Texas at Austin, Austin, TX, 2008.

[56] D.A. MacLaren, N.J. Curson, P. Atkinson, and W. Allison. An afm study of the processing of hydrogen passivated silicon(1x1) of a low miscut angle. *Surface Science*, 490(3):285 – 295, 2001.

[57] Donald Angus MacLaren. *Development of a Single Crystal Mirror for Scanning Helium Microscopy.* PhD thesis, Cambridge University, Cambridge, England, 2002.

[58] Directed Light Inc., 633 River Oaks Parkway, San Jose, CA 95134. *www.directedlight.com.*

[59] MKS Instruments, 2 Tech Drive, Suite 201, Andover, MA 01810. *www.mksinst.com.*

[60] Ferrotec (USA) Corporation, 33 Constitution Drive, Bedford, NH 03110. *www.ferrotec.com.*

[61] Stanford Research Systems Inc., 1290-D Reamwood Avenue, Sunnyvale, CA 94089. *www.thinksrs.com.*

[62] K. Kuhnke, K. Kern, R. David, and G. Comsa. High efficiency molecular-beam ionization detector with short ionization region. *Review of Scientific Instruments*, 65(11):3458–3465, nov 1994.

[63] Mark Kasevich and Steven Chu. Atomic interferometry using stimulated raman transitions. *Phys. Rev. Lett.*, 67(2):181–184, Jul 1991.

[64] David M. Giltner, Roger W. McGowan, and Siu Au Lee. Atom interferometer based on bragg scattering from standing light waves. *Phys. Rev. Lett.*, 75(14):2638–2641, Oct 1995.

[65] P. Berman, editor. *Atom Interferometry.* Academic Press, San Diego, 1997.

[66] John F. Clauser. Ultra-high sensitivity accelerometers and gyroscopes using neutral atom matter-wave interferometry. *Physica B+C*, 151(1):262 – 272, 1988.

[67] R.J. Kaye, E.C. Cnare, M. Cowan, B.W. Duggin, R.J. Lipinski, B.M. Marder, G.M. Douglas, and K.J. Shimp. Design and performance of sandia's contactless coilgun for 50 mm projectiles. *Magnetics, IEEE Transactions on*, 29(1):680 – 685, jan 1993.

[68] E Narevicius, C G Parthey, A Libson, M F Riedel, U Even, and M G Raizen. Towards magnetic slowing of atoms and molecules. *New Journal of Physics*, 9(4):96, 2007.

[69] E Narevicius, C G Parthey, A Libson, J Narevicius, I Chavez, U Even, and M G Raizen. An atomic coilgun: using pulsed magnetic fields to slow a supersonic beam. *New Journal of Physics*, 9(10):358, 2007.

[70] Edvardas Narevicius, Adam Libson, Christian G. Parthey, Isaac Chavez, Julia Narevicius, Uzi Even, and Mark G. Raizen. Stopping supersonic beams with a series of pulsed electromagnetic coils: An atomic coilgun. *Phys. Rev. Lett.*, 100(9):093003, Mar 2008.

[71] Edvardas Narevicius, Adam Libson, Christian G. Parthey, Isaac Chavez, Julia Narevicius, Uzi Even, and Mark G. Raizen. Stopping supersonic oxygen with a series of pulsed electromagnetic coils: A molecular coilgun. *Phys. Rev. A*, 77(5):051401, May 2008.

[72] Christian G. Parthey. Magnetic control of supersonic beams. Master's thesis, The University of Texas at Austin, Austin, TX, 2007.

[73] Isaac Chavez. Magnetic control of supersonic beams: Magnetic slowing to isotope separation. Master's thesis, The University of Texas at Austin, Austin, TX, 2010.

[74] John David Jackson. *Classical Electrodynamics 3rd ed.* Wiley, New York, 1999.

[75] T. Bergeman, Gidon Erez, and Harold J. Metcalf. Magnetostatic trapping fields for neutral atoms. *Phys. Rev. A*, 35:1535–1546, Feb 1987.

[76] Edwin M. McMillan. The synchrotron¯a proposed high energy particle accelerator. *Phys. Rev.*, 68:143–144, Sep 1945.

[77] S.Y. Lee. *Accelerator Physics.* World Scientific, Singapore, 1999.

[78] Hendrick L. Bethlem, Giel Berden, André J. A. van Roij, Floris M. H. Crompvoets, and Gerard Meijer. Trapping neutral molecules in a traveling potential well. *Phys. Rev. Lett.*, 84(25):5744–5747, Jun 2000.

[79] Sebastiaan Y. T. van de Meerakker, Nicolas Vanhaecke, Hendrick L. Bethlem, and Gerard Meijer. Transverse stability in a stark decelerator. *Phys. Rev. A*, 73:023401, Feb 2006.

[80] B. C. Sawyer, B. K. Stuhl, B. L. Lev, J. Ye, and E. R. Hudson. Mitigation of loss within a molecular stark decelerator. *Eur. Phys. J. D*, 48:197–209, 2008.

[81] Adam Libson, Stephen Travis Bannerman, Robert J. Clark, Thomas R. Mazur, and Mark G. Raizen. The atomic coilgun and single-photon cooling. *Hyperfine Interactions*, pages 1–10, 2012. 10.1007/s10751-012-0586-7.

[82] R. Clark, T. Mazur, A. Libson, and M. Raizen. Nanofabrication by magnetic focusing of supersonic beams. *Applied Physics B: Lasers and Optics*, 103:547–551, 2011. 10.1007/s00340-010-4229-x.

[83] Martin Zinner, Peter Spoden, Tobias Kraemer, Gerhard Birkl, and Wolfgang Ertmer. Precision measurement of the metastable $^3P_2$ lifetime of neon. *Phys. Rev. A*, 67:010501, Jan 2003.

[84] A. Lurio, G. Weinreich, C. W. Drake, V. W. Hughes, and J. A. White. Atomic g_J values for neon and argon in their metastable $^3p_2$ states; evidence for zero spin of $_{10}\mathrm{Ne}^{20}$. *Phys. Rev.*, 120:153–157, Oct 1960.

[85] A. V. Chetvertukhin, M. I. Sharipova, A. G. Zhdanov, T. B. Shapaeva, T. V. Dolgova, and A. A. Fedyanin. Femtosecond time-resolved faraday rotation in thin magnetic films and magnetophotonic crystals. *Journal of Applied Physics*, 111(7):07A944, 2012.

[86] Giel Berden, Richard Engeln, Peter C. M. Christianen, Jan C. Maan, and Gerard Meijer. Cavity-ring-down spectroscopy on the oxygen A band in magnetic fields up to 20 t. *Phys. Rev. A*, 58:3114–3123, Oct 1998.

[87] Willis E. Lamb and Robert C. Retherford. Fine structure of the hydrogen atom by a microwave method. *Phys. Rev.*, 72:241–243, Aug 1947.

[88] Thomas C. Killian, Dale G. Fried, Lorenz Willmann, David Landhuis, Stephen C. Moss, Thomas J. Greytak, and Daniel Kleppner. Cold collision frequency shift of the $1s$- $2s$ transition in hydrogen. *Phys. Rev. Lett.*, 81(18):3807–3810, Nov 1998.

[89] M. Niering, R. Holzwarth, J. Reichert, P. Pokasov, Th. Udem, M. Weitz, T. W. Hänsch, P. Lemonde, G. Santarelli, M. Abgrall, P. Laurent, C. Salomon, and A. Clairon. Measurement of the hydrogen $1s$- $2s$ transition frequency by phase coherent comparison with a microwave cesium fountain clock. *Phys. Rev. Lett.*, 84(24):5496–5499, Jun 2000.

[90] F. Biraben, T. W. Hänsch, M. Fischer, M. Niering, R. Holzwarth, J. Reichert, T. Udem, M. Weitz, B. de Beauvoir, C. Schwob, L. Jozefowski, L. Hilico, F. Nez, L. Julien, O. Acef, J.-J. Zondy, and A. Clairon. Precision spectroscopy of atomic hydrogen. In S. G. Karshenboim, F. S. Pavone, G. F. Bassani, M. Inguscio, and T. W. Hänsch, editors, *The Hydrogen Atom: Precision Physics of Simple Atomic Systems*, volume 570 of *Lecture Notes in Physics, Berlin Springer Verlag*, page 17, 2001.

[91] Christian G. Parthey, Arthur Matveev, Janis Alnis, Birgitta Bernhardt, Axel Beyer, Ronald Holzwarth, Aliaksei Maistrou, Randolf Pohl, Katharina Predehl, Thomas Udem, Tobias Wilken, Nikolai Kolachevsky, Michel Abgrall, Daniele Rovera, Christophe Salomon, Philippe Laurent, and Theodor W. Hänsch. Improved measurement of the hydrogen $1s^{\smile}2s$ transition frequency. *Phys. Rev. Lett.*, 107:203001, Nov 2011.

[92] Randolf Pohl, Aldo Antognini, Franois Nez, Fernando D. Amaro, Franois Biraben, Joo M. R. Cardoso, Daniel S. Covita, Andreas Dax, Satish Dhawan, Luis M. P. Fernandes, Adolf Giesen, Thomas Graf, Theodor W. Hnsch, Paul Indelicato, Lucile Julien, Cheng-Yang Kao, Paul Knowles, Eric-Olivier Le Bigot, Yi-Wei Liu, Jos A. M. Lopes, Livia Ludhova, Cristina M. B. Monteiro, Franoise Mulhauser, Tobias Nebel, Paul Rabinowitz, Joaquim M. F. dos Santos, Lukas A. Schaller, Karsten Schuhmann, Catherine Schwob, David Taqqu, Joo F. C. A. Veloso, and Franz Kottmann. The size of the proton. *Nature*, 466(4703):213–216, July 2010.

[93] Christian G. Parthey, Arthur Matveev, Janis Alnis, Randolf Pohl, Thomas Udem, Ulrich D. Jentschura, Nikolai Kolachevsky, and Theodor W. Hänsch. Precision measurement of the hydrogen-deuterium $1s - 2s$ isotope shift. *Phys. Rev. Lett.*, 104(23):233001, Jun 2010.

[94] U.D. Jentschura, A. Matveev, C.G. Parthey, J. Alnis, R. Pohl, Th. Udem, N. Kolachevsky, and T.W. Hänsch. Hydrogen-deuterium isotope shift: From the $1s-2s$-transition frequency to the proton-deuteron charge-radius difference. *Phys. Rev. A*, 83:042505, Apr 2011.

[95] M. Amoretti, C. Amsler, G. Bonomi, A. Bouchta, P. Bowe, C. Carraro, C. L. Cesar, M. Charlton, M. J. Collier, M. Doser, V. Filippini, K. S. Fine, A. Fontana, M. C. Fujiwara, R. Funakoshi, P. Genova, J. S. Hangst, R. S. Hayano, M. H. Holzscheiter, L. V. Jorgensen, V. Lagomarsino, R. Landua, D. Lindelof, E. Lodi Rizzini, M. Macri, N. Madsen, G. Manuzio, M. Marchesotti, P. Montagna, H. Pruys, C. Regenfus, P. Riedler, J. Rochet, A. Rotondi, G. Rouleau, G. Testera, A. Variola, T. L. Watson, and D. P. van der Werf. Production and detection of cold antihydrogen atoms. *Nature*, 419(6906):456–459, 2002.

[96] G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hänsch, and E. A. Hessels. Background-free observation of cold antihydrogen with field-ionization analysis of its states. *Phys. Rev. Lett.*, 89(21):213401, Oct 2002.

[97] G. B. Andresen, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, P. D. Bowe, E. Butler, C. L. Cesar, S. Chapman, M. Charlton, A. Deller, S. Eriksson, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, A. Gutierrez, J. S. Hangst, W. N. Hardy, M. E. Hayden, A. J. Humphries, R. Hydomako, M. J. Jenkins, S. Jonsell, L. V. Jrgensen, L. Kurchaninov, N. Madsen, S. Menary, P. Nolan, K. Olchanski, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, S. Seif el Nasr, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki. Trapped antihydrogen. *Nature*, 468:673–676, Oct 2010.

[98] Harald F. Hess, Greg P. Kochanski, John M. Doyle, Naoto Masuhara, Daniel Kleppner, and Thomas J. Greytak. Magnetic trapping of spin-polarized atomic hydrogen. *Phys. Rev. Lett.*, 59(6):672–675, Aug 1987.

[99] Dale G. Fried, Thomas C. Killian, Lorenz Willmann, David Landhuis, Stephen C. Moss, Daniel Kleppner, and Thomas J. Greytak. Bose-einstein condensation of atomic hydrogen. *Phys. Rev. Lett.*, 81(18):3811–3814, Nov 1998.

[100] Dale Fried. *Bose-Einstein Condensation of Atomic Hydrogen.* PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1999.

[101] I. D. Setija, H. G. C. Werij, O. J. Luiten, M. W. Reynolds, T. W. Hijmans, and J. T. M. Walraven. Optical cooling of atomic hydrogen in a magnetic trap. *Phys. Rev. Lett.*, 70(15):2257–2260, Apr 1993.

[102] Saijun Wu, Roger C. Brown, William D. Phillips, and J. V. Porto. Pulsed sisyphus scheme for laser cooling of atomic (anti)hydrogen. *Phys. Rev. Lett.*, 106:213001, May 2011.

[103] D. Kielpinski. Laser cooling of atoms and molecules with ultrafast pulses. *Phys. Rev. A*, 73(6):063407, Jun 2006.

[104] Wolfgang Petrich, Michael H. Anderson, Jason R. Ensher, and Eric A. Cornell. Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Phys. Rev. Lett.*, 74:3352–3355, Apr 1995.

[105] Jianxun Yan, Ashraf El-Dakrouri, Mounir Laroussi, and Mool C. Gupta. 121.6 nm radiation source for advanced lithography. *J. Vac. Sci. Technol. B*, 20(6):2574–2577, 2002.

[106] R. G. Beausoleil and T. W. Hansch. Ultrahigh-resolution two-photon optical ramsey spectroscopy of an atomic fountain. *Phys. Rev. A*, 33:1661–1670, Mar 1986.

[107] Thomas Killian. *1S-2S Spectroscopy of Trapped Hydrogen: The Cold Collision Frequency Shift and Studies of BEC.* PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1999.

[108] H.A. Bethe and E.E. Salpeter. Quantum mechanics of one and two-electron atoms. In S. Flügge, editor, *Encyclopedia of physics*, Handbuch der Physik, Berlin, 1957. Springer-Verlag.

[109] Claudio L. Cesar. Zeeman effect on the $1s - 2s$ transition in trapped hydrogen and antihydrogen. *Phys. Rev. A*, 64:023418, Jul 2001.

[110] N. Kolachevsky, J. Alnis, S. D. Bergeson, and T. W. Hänsch. Compact solid-state laser source for $1s - 2s$ spectroscopy in atomic hydrogen. *Phys. Rev. A*, 73(2):021801, Feb 2006.

[111] W. Demtröder. *Laser Spectroscopy: Basic Concepts and Instrumentation*. Advanced Texts in Physics. Springer, Berlin, 2003.

[112] J.R.M. Barr, J.M. Girkin, A.I. Ferguson, G.P. Barwood, P. Gill, W.R.C. Rowley, and R.C. Thompson. Interferometric frequency measurements of 130te2 transitions at 486 nm. *Optics Communications*, 54(4):217 – 221, 1985.

[113] D. H. McIntyre, W. M. Fairbank, S. A. Lee, T. W. Hänsch, and E. Riis. Interferometric frequency measurement of $^{130}Te_2$ reference transitions at 486 nm. *Phys. Rev. A*, 41:4632–4635, May 1990.

[114] J. D. Gillaspy and Craig J. Sansonetti. Absolute wavelength determinations in molecular tellurium: new reference lines for precision laser spectroscopy. *J. Opt. Soc. Am. B*, 8(12):2414–2419, Dec 1991.

[115] J. Cariou and P. Luc. *Atlas du Spectre dAbsorption de la Molécule de Tellure*. Editions du Centre National de la Recherche Scientifique, Paris, 1980.

[116] Stephen Travis Bannerman. *Cooling Atomic Ensembles with Maxwell's Demon*. PhD thesis, The University of Texas at Austin, Austin, TX, 2011.

[117] Yushi Kaneda, J. M. Yarborough, Li Li, N. Peyghambarian, Li Fan, Chris Hessenius, Mahmoud Fallahi, Jörg Hader, Jerome V. Moloney, Yoshiyuki Honda, Masato Nishioka, Youhei Shimizu, Kenshi Miyazono, Hiroya Shimatani, Masashi Yoshimura, Yusuke Mori, Yasuo Kitaoka, and Takatomo Sasaki. Continuous-wave all-solid-state 244 nm deep-ultraviolet laser source by fourth-harmonic

generation of an optically pumped semiconductor laser using cslib6o10 in an external resonator. *Opt. Lett.*, 33(15):1705–1707, Aug 2008.

[118] Justin Paul, Yushi Kaneda, Tsuei-Lian Wang, Christian Lytle, Jerome V. Moloney, and R. Jason Jones. Doppler-free spectroscopy of mercury at 253.7 nm using a high-power, frequency-quadrupled, optically pumped external-cavity semiconductor laser. *Opt. Lett.*, 36(1):61–63, Jan 2011.

[119] A. E. Kramida. A critical compilation of experimental data on spectral lines and energy levels of hydrogen, deuterium, and tritium. *Atomic Data and Nuclear Data Tables*, 96:586–644, Nov 2010.

[120] Friar J.L. The structure of light nuclei and its effect on precise atomic measurements. *Canadian Journal of Physics*, 80(11):1337–1346, 2002.

[121] W. Hänsel, J. Reichel, P. Hommelhoff, and T. W. Hänsch. Magnetic conveyor belt for transporting and merging trapped atom clouds. *Phys. Rev. Lett.*, 86:608–611, Jan 2001.

[122] Markus Greiner, Immanuel Bloch, Theodor W. Hänsch, and Tilman Esslinger. Magnetic transport of trapped cold atoms over a large distance. *Phys. Rev. A*, 63:031401, Feb 2001.

[123] Samuel A. Meek, Hendrick L. Bethlem, Horst Conrad, and Gerard Meijer. Trapping molecules on a chip in traveling potential wells. *Phys. Rev. Lett.*, 100:153003, Apr 2008.

[124] Samuel A. Meek, Horst Conrad, and Gerard Meijer. Trapping molecules on a chip. *Science*, 324(5935):1699–1702, 2009.

[125] Andreas Osterwalder, Samuel A. Meek, Georg Hammer, Henrik Haak, and Gerard Meijer. Deceleration of neutral molecules in macroscopic traveling traps. *Phys. Rev. A*, 81:051401, May 2010.

[126] A. Trimeche, M.N. Bera, J.-P. Cromières, J. Robert, and N. Vanhaecke. Trapping of a supersonic beam in a traveling magnetic wave. *Eur. Phys. J. D*, 65(1-2):263–271, Nov 2011.

[127] Etay Lavert-Ofir, Liron David, Alon B. Henson, Sasha Gersten, Julia Narevicius, and Edvardas Narevicius. Stopping paramagnetic supersonic beams: the advantage of a co-moving magnetic trap decelerator. *Phys. Chem. Chem. Phys.*, 13:18948–18953, 2011.

[128] Etay Lavert-Ofir, Sasha Gersten, Alon B Henson, Itamar Shani, Liron David, Julia Narevicius, and Edvardas Narevicius. A moving magnetic trap decelerator: a new source of cold atoms and molecules. *New Journal of Physics*, 13(10):103030, 2011.

[129] Masao Takamoto and Hidetoshi Katori. Spectroscopy of the $^1s_0$-$^3p_0$ clock transition of ^{87}Sr in an optical lattice. *Phys. Rev. Lett.*, 91:223001, Nov 2003.

[130] A. Ruschhaupt and J. G. Muga. Atom diode: A laser device for a unidirectional transmission of ground-state atoms. *Phys. Rev. A*, 70(6):061604, Dec 2004.

[131] M. G. Raizen, A. M. Dudarev, Qian Niu, and N. J. Fisch. Compression of atomic phase space using an asymmetric one-way barrier. *Phys. Rev. Lett.*, 94(5):053003, Feb 2005.

[132] A. M. Dudarev, M. Marder, Q. Niu, N. Fisch, and M. G. Raizen. Statistical mechanics of an optical phase space compressor. *Europhys. Lett.*, 70:761, 2005.

[133] A. Ruschhaupt, J. G. Muga, and M. G. Raizen. One-photon atomic cooling with an optical maxwell's demon valve. *J. Phys. B*, 39:3833, 2006.

[134] G. Price, S. Bannerman, E. Narevicius, and M. Raizen. Single-photon atomic cooling. *Laser Physics*, 17:965–968, 2007. 10.1134/S1054660X07070109.

[135] Mark G. Raizen. Comprehensive control of atomic motion. *Science*, 324(5933):1403–1406, 2009.

[136] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5:183, 1961.

[137] Jeremy J. Thorn, Elizabeth A. Schoene, Tao Li, and Daniel A. Steck. Experimental realization of an optical one-way barrier for neutral atoms. *Phys. Rev. Lett.*, 100(24):240407, Jun 2008.

[138] Gabriel N. Price, S. Travis Bannerman, Kirsten Viering, Edvardas Narevicius, and Mark G. Raizen. Single-photon atomic cooling. *Phys. Rev. Lett.*, 100(9):093004, Mar 2008.

[139] S Travis Bannerman, Gabriel N Price, Kirsten Viering, and Mark G Raizen. Single-photon cooling at the limit of trap dynamics: Maxwell's demon near maximum efficiency. *New Journal of Physics*, 11(6):063044, 2009.

[140] Gabriel Noam Price. *Single-Photon Atomic Cooling*. PhD thesis, The University of Texas at Austin, Austin, TX, 2009.

[141] Kirsten Viering. *Experiments to Control Atom Number and Phase-Space Density in Cold Gases*. PhD thesis, The University of Texas at Austin, Austin, TX, 2012.

[142] C. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum Mechanics*. Hermann, Paris, 1977.

[143] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *Atom-photon interactions: basic processes and applications*. Wiley, New York, 1998.

[144] O. Zobay and B. M. Garraway. Two-dimensional atom trapping in field-induced adiabatic potentials. *Phys. Rev. Lett.*, 86:1195–1198, Feb 2001.

[145] Y. Colombe, E. Knyazchyan, O. Morizot, B. Mercier, V. Lorent, and H. Perrin. Ultracold atoms confined in rf-induced two-dimensional trapping potentials. *Europhysics Letters*, 67(4):593, 2004.

[146] S. Hofferberth, I. Lesanovsky, B. Fischer, J. Verdu, and J. Schmiedmayer. Radiofrequency-dressed-state potentials for neutral atoms. *Nature Physics*, 2:710–716, Sep 2006.

[147] O Morizot, C L Garrido Alzar, P-E Pottie, V Lorent, and H Perrin. Trapping and cooling of rf-dressed atoms in a quadrupole magnetic field. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 40(20):4013, 2007.

[148] T. Fernholz, R. Gerritsma, P. Krüger, and R. J. C. Spreeuw. Dynamically controlled toroidal and ring-shaped magnetic traps. *Phys. Rev. A*, 75:063406, Jun 2007.

[149] I. Lesanovsky, T. Schumm, S. Hofferberth, L. M. Andersson, P. Krüger, and J. Schmiedmayer. Adiabatic radio-frequency potentials for the coherent manipulation of matter waves. *Phys. Rev. A*, 73:033619, Mar 2006.

[150] I. Lesanovsky, S. Hofferberth, J. Schmiedmayer, and P. Schmelcher. Manipulation of ultracold atoms in dressed adiabatic radio-frequency potentials. *Phys. Rev. A*, 74:033619, Sep 2006.

[151] S. Hofferberth, B. Fischer, T. Schumm, J. Schmiedmayer, and I. Lesanovsky. Ultracold atoms in radio-frequency dressed potentials beyond the rotating-wave approximation. *Phys. Rev. A*, 76:013401, Jul 2007.

[152] Edvardas Narevicius, S Travis Bannerman, and Mark G Raizen. Single-photon molecular cooling. *New Journal of Physics*, 11(5):055046, 2009.

[153] David E. Pritchard. Cooling neutral atoms in a magnetic trap for precision spectroscopy. *Phys. Rev. Lett.*, 51:1336–1339, Oct 1983.

[154] Dmitry Budker, Derek F. Kimball, and David P. DeMille. *Atomic Physics*. Oxford University Press, New York, 2008.

[155] John N. Bahcall. Theory of bound-state beta decay. *Phys. Rev.*, 124(2):495–499, Oct 1961.

[156] S. Cohen, D. Murnick, and R. Raghavan. Bound-state beta-decay and kinematic search for neutrino mass. *Hyperfine Interactions*, 33:1–8, 1987. 10.1007/BF02394096.

[157] M Jerkins, J R Klein, J H Majors, F Robicheaux, and M G Raizen. Using cold atoms to measure neutrino mass. *New Journal of Physics*, 12(4):043022, 2010.

[158] Willian C. Stwalley and L. H. Nosanow. Possible "new" quantum systems. *Phys. Rev. Lett.*, 36:910–913, Apr 1976.

[159] D. Blume, B. D. Esry, Chris H. Greene, N. N. Klausen, and G. J. Hanna. Formation of atomic tritium clusters and bose-einstein condensates. *Phys. Rev. Lett.*, 89:163402, Sep 2002.

[160] Aurel Bulgac. Dilute quantum droplets. *Phys. Rev. Lett.*, 89:050402, Jul 2002.

[161] Paulo F. Bedaque, Aurel Bulgac, and Gautam Rupak. Quantum corrections to dilute bose liquids. *Phys. Rev. A*, 68:033606, Sep 2003.

[162] N. Balakrishnan and A. Dalgarno. Chemistry at ultracold temperatures. *Chemical Physics Letters*, 341(5-6):652 – 656, 2001.

Vita

Adam Alexander Libson was born in Philadelphia, Pennsylvania on 9 June 1982, the son of Jeffrey Paul Libson and Dr. Christian Alexander Libson. He attended school at Friends' Central School and graduated in 2000. While in high school, he spent the 1998-1999 academic year in Rennes, France with School Year Abroad. He received the Bachelor of Arts degree in Physics from Carleton College in 2004, graduating Magna Cum Laude with distinction. He spent the fall semester of 2002 in Budapest, Hungary with Budapest Semesters in Mathematics. He applied to the University of Texas at Austin for enrollment in the physics program in 2005, and began his graduate studies that summer.

Permanent address: 4312 Duval St. Apt. 206
Austin, Texas 78751

This dissertation was typeset with LaTeX[†] by the author.

[†]LaTeX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's TeX Program.