

# Toward Antihydrogen Spectroscopy

A thesis presented

by

Nathan Jones

to

The Department of Physics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Physics

Harvard University

Cambridge, Massachusetts

May 2020

©2020 - Nathan Jones

All rights reserved.

Thesis advisor  
**Gerald Gabrielse**

Author  
**Nathan Jones**

# Toward Antihydrogen Spectroscopy

## Abstract

ATRAP's new experimental apparatus is successfully commissioned and used to synthesize and capture antihydrogen atoms. While the number trapped per trial,  $5 \pm 2$ , is about the same as our yield in 2011, the rate of antiproton use is a factor of three lower, and the length of a trial is cut in half. A more robust apparatus, the use of smaller radius plasmas, and plasma imaging all contribute.

Simulations I built and ran on Harvard's Odyssey cluster focus on both Lyman- $\alpha$  laser cooling and 1S-2S spectroscopy of trapped antihydrogen, using the laser systems available at ATRAP. These show how laser cooling can reduce average atomic energies by more than an order of magnitude in tens of minutes. Spectroscopy of the 1S-2S line requires similar amounts of time, and laser cooling preceding spectroscopy narrows the observed spectral width close to that seen in hydrogen beam experiments. The computational work points the way to efficient ATRAP laser experiments. Precise 1S-2S spectroscopy on antihydrogen, an important test of CPT symmetry, is within reach.

# Contents

Title Page . . . . .	i
Abstract . . . . .	iii
Table of Contents . . . . .	iv
Publications . . . . .	v
Acknowledgments . . . . .	vi
<b>1 Introduction</b>	<b>1</b>
1.1 An Abridged History of Antihydrogen Trapping . . . . .	2
1.2 CPT Tests with Antihydrogen . . . . .	3
1.3 Thesis Contents . . . . .	5
<b>2 Apparatus</b>	<b>8</b>
2.1 Superconducting Solenoid . . . . .	8
2.2 Cryogenic and Vacuum System . . . . .	10
2.3 Ioffe Magnet . . . . .	14
2.4 Electrode Stack . . . . .	16
2.5 XY Stage . . . . .	19
2.6 Detectors . . . . .	21
2.7 Laser Paths . . . . .	23
2.8 Positron Source . . . . .	28
2.9 Summary . . . . .	30
<b>3 Ioffe Magnet</b>	<b>32</b>
3.1 Antihydrogen in Magnetic Fields . . . . .	32
3.2 Ideal Ioffe Traps . . . . .	34
3.3 BTRAP's Ioffe Trap . . . . .	36
3.4 CTRAP's Ioffe Trap . . . . .	39
3.5 CTRAP Magnet Enclosure . . . . .	43
3.6 Magnet Operation . . . . .	55
3.7 Summary . . . . .	58

<b>4</b>	<b>Penning Traps in Theory</b>	<b>59</b>
4.1	Fields and Single-Particle Motion . . . . .	59
4.2	Plasmas in Equilibrium . . . . .	63
4.3	Collisions and Cooling . . . . .	67
4.4	Summary . . . . .	72
<b>5</b>	<b>CTRAP's Penning Trap</b>	<b>74</b>
5.1	CTRAP Electrodes and Potentials . . . . .	74
5.2	Electrode Construction . . . . .	79
5.3	Enclosure Construction . . . . .	85
5.4	Summary . . . . .	91
<b>6</b>	<b>Experimental Procedures and Results</b>	<b>92</b>
6.1	Particle Counting . . . . .	92
6.2	Plasma Geometry . . . . .	94
6.3	Rotating Wall . . . . .	101
6.4	Electron Loading . . . . .	102
6.5	Antiproton Loading . . . . .	104
6.6	Positron Loading . . . . .	108
6.7	Antihydrogen Synthesis and Trapping . . . . .	109
6.8	Summary . . . . .	112
<b>7</b>	<b>Atomic Theory</b>	<b>114</b>
7.1	Atomic States . . . . .	114
7.1.1	L = 0 States . . . . .	115
7.1.2	L = 1 States . . . . .	120
7.2	1S-2P Equations of Motion . . . . .	121
7.2.1	Atom-Laser Interaction . . . . .	121
7.2.2	Decay . . . . .	123
7.2.3	Target Transitions . . . . .	126
7.2.4	Density Matrix . . . . .	129
7.3	1S-2S Equations of Motion . . . . .	131
7.3.1	Equations of Motion and Two-Photon Quantities . . . . .	132
7.3.2	Stark Effects . . . . .	136
7.4	Summary . . . . .	138
<b>8</b>	<b>Simulation Code</b>	<b>139</b>
8.1	Foundations . . . . .	140
8.1.1	Filesystem Operations . . . . .	140
8.1.2	Data Arrays . . . . .	141
8.2	Grids of Physical Quantities . . . . .	145
8.2.1	Coordinate Grids . . . . .	145

8.2.2	Data Grids . . . . .	147
8.2.3	Field Definitions . . . . .	149
8.2.4	Derived Grid Quantities . . . . .	152
8.2.5	Interpolation . . . . .	154
8.3	Kinematic and Atomic Evolution . . . . .	161
8.3.1	Kinematics . . . . .	161
8.3.2	Atomic States . . . . .	164
8.3.3	Laser Information . . . . .	165
8.3.4	Atomic Physics . . . . .	167
8.3.5	Internal State Evolution . . . . .	168
8.3.6	Decays . . . . .	171
8.3.7	Evolution and Decay . . . . .	172
8.3.8	State Summary . . . . .	174
8.4	Simulations . . . . .	176
8.5	Cluster Management . . . . .	179
8.5.1	Partitions . . . . .	179
8.5.2	Cluster Management . . . . .	181
8.5.3	Program Options . . . . .	184
8.6	Summary . . . . .	186
<b>9</b>	<b>Laser Cooling Simulations</b>	<b>187</b>
9.1	Methods . . . . .	187
9.2	Laser Cooling Simulation Results . . . . .	189
9.2.1	Spectra . . . . .	190
9.2.2	Asymptotic Behaviour . . . . .	196
9.2.3	Losses and Unintended Transitions . . . . .	199
9.3	Cooling Detection Strategies . . . . .	200
9.3.1	Dumping the Trap . . . . .	202
9.3.2	Laser Removal . . . . .	205
9.4	Summary . . . . .	213
<b>10</b>	<b>1S-2S Spectroscopy Simulations</b>	<b>215</b>
10.1	Methods . . . . .	216
10.1.1	Initial Conditions . . . . .	218
10.1.2	Simplifying the Equations of Motion . . . . .	219
10.1.3	No-Recoil Approximation . . . . .	220
10.2	Spectroscopic Signal . . . . .	224
10.3	Summary . . . . .	231
<b>11</b>	<b>Conclusions and Discussion</b>	<b>232</b>
	<b>Bibliography</b>	<b>234</b>

# PhD Publications

1. *Simulation of Spectroscopy of Trapped Antihydrogen*  
N. Jones and G. Gabrielse, (forthcoming).
2. *Simulation of Laser Cooling of Trapped Antihydrogen*  
N. Jones and G. Gabrielse, (forthcoming).
3. *Two-Symmetry Penning-Ioffe Trap for Antihydrogen Cooling and Spectroscopy*  
E. Tardiff, X. Fan, G. Gabrielse, D. Grzonka, C. Hamley, E.A. Hessels, N. Jones, G. Khatri, S. Kolthammer, D. Martinez Zambrano, C. Meisenhelder, T. Morrison, E. Nottet, E. Novitski, and C. Storry, (forthcoming).
4. *Lyman- $\alpha$  source for laser cooling antihydrogen*  
G. Gabrielse, B. Glowacz, D. Grzonka, C. D. Hamley, E. A. Hessels, N. Jones, G. Khatri, S. A. Lee, C. Meisenhelder, T. Morrison, E. Nottet, C. Rasor, S. Ronald, T. Skinner, C. H. Storry, E. Tardiff, D. Yost, D. Martinez Zambrano, and M. Zielinski, *Optics Letters* **43**, 2905 (2018).
5. *Electron-cooled accumulation of  $4 \times 10^9$  positrons for production and storage of antihydrogen atoms*  
D. W. Fitzakerley, M. C. George, E. A. Hessels, T. D. G. Skinner, C. H. Storry, M. Weel, G. Gabrielse, C. D. Hamley, N. Jones, K. Marable, E. Tardiff, D. Grzonka, W. Oelert, and M. Zielinski (ATRAP Collaboration), *Journal of Physics B* **49**, 064001 (2016).

# Acknowledgments

I am grateful for the opportunity to work as part of the ATRAP collaboration on such an exciting experiment with so many wonderful people. My thesis advisor, Gerald Gabrielse, was integral to my growth as an experimentalist during this project. He was not only willing to discuss any and all aspects of the experiment at length, but he also worked long hours alongside the rest of the group on his frequent trips to CERN. His commitment and enthusiasm for the experiment knew no limits, and the example of perseverance he set through the difficulties we faced was a great encouragement.

Eric Tardiff served as my mentor when I entered the lab, getting me up to speed on cryogenics, vacuum systems, and all things ATRAP. The hours spent working on the trap and speaking in Seinfeld-isms will not soon be forgot. Nor will the high standard he set for himself and others, a quality that raised the performance of everyone he worked with.

I overlapped with Rita Kalra and Chris Hamley near the beginning of my work with the experiment. They each taught me a great deal and were instrumental in our progress toward using the apparatus for the first time. Gunn Khatri, Tharon Morrison, Daniel Zambrano, and Edouard Nottet arrived later, and their efforts allowed us to catch antihydrogen for the first time in the new trap. I could not have asked for a better group to share the late nights at the AD.

The collaborators I worked with also made a deep impact on the project. The group at York, made up of Eric Hessels, Cody Storry, Matt George, Matt Weel, Dan Fitzakerley, and Taylor Skinner, ensured that we had a reliable source of positrons. Our detectors were maintained by Walter Oelert, Dieter Grzonka, and Marcin Zielin-

## *Acknowledgments*

---

ski. A new Lyman- $\alpha$  source was also constructed with the advice and help of Dylan Yost and Siu Au Lee at Colorado State. Without these collaborators, the progress we saw during my PhD would not have been possible.

Finally, I want to thank my wife Oksana for her support throughout this process. She has moved to a different continent and an unfamiliar country with me. My many late nights at the AD have left care for our son entirely in her hands. She has been there to celebrate my triumphs and listen to my complaints. I am ever grateful for her love, support, and encouragement during what has been a difficult PhD.

# Chapter 1

## Introduction

Antihydrogen, the bound state of an antiproton and a positron, represents a simple, but exotic system. As such, its study has the potential to shed new light on old assumptions. For example, the Standard Model's CPT symmetry invariance gives a definite relationship between matter and antimatter physics. The simplicity of antihydrogen and the precision to which the structure of its matter counterpart has been measured make this anti-atom an attractive means for testing CPT invariance.

Antihydrogen's potential as a probe of such fundamental physics has bolstered interest in its production. Decades of progress toward this goal has finally led to regular trapping of the antimatter system [1,2]. With preliminary spectroscopy already being done [3], we are on the verge of making truly precise comparisons between hydrogen and its antimatter counterpart.

This introductory chapter starts with a history of antihydrogen production and trapping. A more detailed discussion of CPT symmetry and antihydrogen's promise as a means to test for such invariance follows. Finally, I outline the contents of this

thesis. The work described within is my effort, along with the rest of the ATRAP collaboration, to move one step closer to precision antihydrogen spectroscopy.

## **1.1 An Abridged History of Antihydrogen Trapping**

In 1986, it was reported that, for the first time, antiprotons had been confined in a trap [4]. This accomplishment paved the way for antihydrogen trapping and led to the formation of the TRAP (later ATRAP) collaboration. Over the next several years, TRAP took significant steps. Using foils for slowing incoming particles and electrons to cool them once caught, antiproton energies were reduced from 5.9 MeV to under 100 meV [5,6]. Antiprotons and positrons were then caught in the same trap and made to interact, cooling the antiprotons [7,8].

Cold antihydrogen was produced in 2002 by both the ATRAP and the ATHENA (which later split into ALPHA and AEGIS) collaborations [9,10]. ATRAP later demonstrated antihydrogen production through an alternative double-charge exchange method [11]. This was followed by anti-atom generation in a Penning-Ioffe trap [12]. Finally, antihydrogen was produced and confined in magnetic traps. ATRAP produced 5 atoms per mixing trial, confining them between 15 and 1000 s [1]. ALPHA originally reported 0.1 atoms per trial, trapped for a total of 172 ms [13]. They later reached 1 atom per trial, holding it for 1000 s [14]. More recently, ALPHA has reported producing about 14 atoms per trial [2].

## 1.2 CPT Tests with Antihydrogen

With the antihydrogen produced so far, work toward tests challenging physical assumptions has begun. One of these is that nature is symmetric under simultaneous interchange of particles and antiparticles, inversion of spatial coordinates, and reversal of time's direction. These correspond, respectively, to charge conjugation (C), parity transformation (P), and time reversal (T). The invariance of a physical theory under the simultaneous application of all three is called CPT symmetry.

The belief that CPT is a true symmetry of nature is grounded in both theory and observation. It can be shown, under sensible assumptions, that any "local, Lorentz-covariant field theories of point particles" possesses this symmetry [15,16]. On the other hand, measurements performed on a wide variety of elementary particles and bound states establish stringent limits on the symmetry violations that may exist in different sectors (Fig. 1.1). It certainly appears that CPT invariance is universal.

Still, there are reasons to be skeptical. For one, evident symmetries of decades past have been proven incorrect. Assumptions that weak interactions conserve parity lead to the  $\theta$ - $\tau$  puzzle of the 1950s [19]. After the possibility of parity violation in the weak sector was raised, experiments confirmed this to be the case [20,21]. Soon after, CP was proposed as a symmetry of nature [22]. Experiments later forced the rejection of that hypothesis [23,24].

Another reason to doubt CPT symmetry is the observed dominance of matter over antimatter in the universe. A classic treatment of this problem gives three "Sakharov conditions" for the observed matter-antimatter asymmetry to arise: baryon number violation, C and CP violation beyond what is currently known, and thermal non-

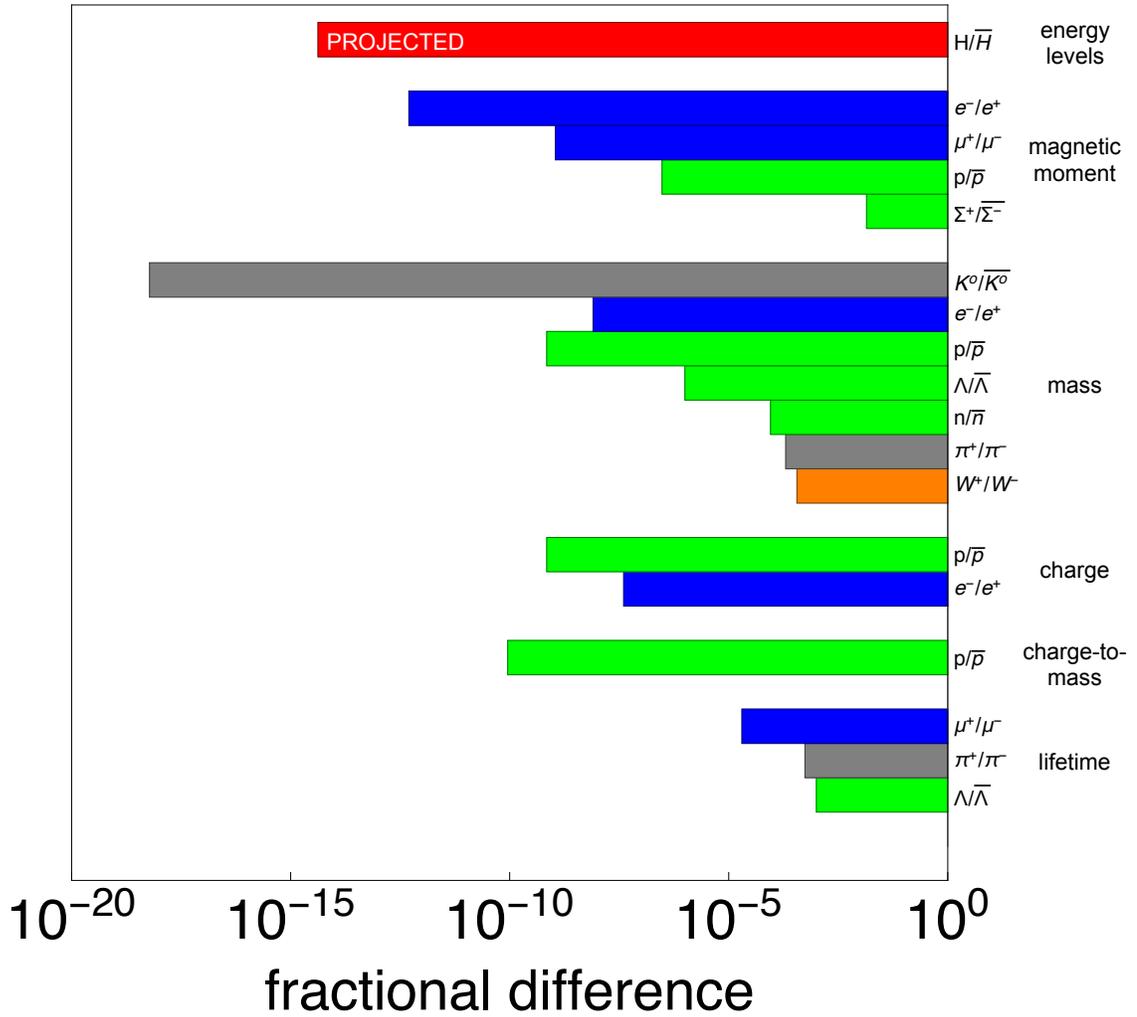


Figure 1.1: Measurements are taken from [17, 18]. The color of each bar indicates whether a test is performed on leptons (blue), baryons (green), mesons (gray), bosons (orange), or a lepton-baryon bound state (red).

equilibrium [25]. However, if CPT is broken, only baryon number non-conservation is needed [26, 27].

Antihydrogen's 1S-2S spectroscopic transition is suited for probing CPT symmetry with particular precision. Hydrogen's 1S-2S frequency has already been measured to four parts in  $10^{15}$  [18]. Although this was done for a hydrogen beam, and less precision may be possible in a magnetic trap, a part-in- $10^{12}$  measurement was carried out on hydrogen's 1S-2S transition in the latter environment [28]. Laser stability limited the precision of the result, so, in principle, even more precise spectroscopy is possible.

### **1.3 Thesis Contents**

This thesis describes both the experimental and computational work I did as a member of ATRAP. While the subjects covered are diverse, each contribution described is a means toward robustly producing trapped antihydrogen, laser cooling the atoms, and performing 1S-2S spectroscopy on the cooled sample.

The opening chapter (Ch. 2) gives an overview of the latest version of the apparatus used by ATRAP to synthesize and trap antihydrogen. While much of it was built when I joined the experiment, the trap still required substantial improvements, occasional part replacement, and an intimate familiarity with its many components. One system that did not exist before my PhD work was the set of laser paths. The trap now boasts three orthogonal routes by which laser light can access trapped atoms.

The central components of the apparatus are discussed in more detail in the following chapters. The Ioffe magnet, which was substantially upgraded in the most recent version of the trap, is responsible for catching antihydrogen atoms produced

within its bore. Ch. 3 explains the advantages of the new Ioffe magnet. It also recounts the difficulties we had, as well as the eventual triumph, when constructing an adequate enclosure for holding its cryogen supply. The Penning trap, tasked with catching and manipulating antihydrogen's precursors, is described in Chs. 4 and 5. Included in that account is my fabrication of new electrodes for compatibility with the new Ioffe enclosure, as well as the more robust vacuum enclosure around the electrode stack that I designed and constructed. Ultimately, these improvements saved a great deal of experiment time by eliminating chronic leaks.

Ch. 6 discusses how we made and trapped antihydrogen. Adjustments to the ATRAP procedure and new components give us smaller plasma radii and a better understanding of our plasma geometries. Experiments performed in 2018 provide a clear direction toward greater progress.

The focus of the final chapters is on computational methods and simulations that show how to make the trapped atoms interact with cooling and spectroscopy lasers. This was a major part of my work. Ch. 7 sets up the atomic theory needed in the simulations that follow while Ch. 8 explains the code used to perform the simulations. Ch. 9 presents simulation results for Lyman- $\alpha$  laser cooling, showing that significant energy reduction is possible for anti-atoms in our Ioffe trap. Ch. 10 follows up on these results, demonstrating both our ability to perform 1S-2S spectroscopy over reasonable times without cooling the atoms and the degree to which laser cooling reduces the widths of the lines we observe. These studies allow us to predict optimal cooling and spectroscopy procedures ahead of the actual experiments, reducing the amount of time-consuming trial-and-error needed. They also give us confidence that

our apparatus' design is compatible with the high-precision measurements that are the goal of this work.

In summary, while we have yet to realize the ultimate goal of precise 1S-2S spectroscopy on antihydrogen, the work described in the following manuscript makes substantial progress in that direction. Our team was able to improve ATRAP's newest antihydrogen apparatus, making it far more robust and adding the components needed for laser access. Antihydrogen was captured for the first time in this apparatus, and an understanding of factors limiting our yield was gained. Finally, ground work was laid for laser studies, promising to save time on future experiments and giving evidence of the feasibility of the envisioned precision measurement.

# Chapter 2

## Apparatus

The ATRAP collaboration's newest experimental device for working with antihydrogen is known as CTRAP. The CTRAP apparatus is an amalgamation of a variety of systems and structures that capture antiprotons, generate and trap positrons, form antihydrogen, and trap it. In this chapter, each of the major components of CTRAP, along with crucial complementary devices, will be described. Later chapters (3-5) cover the heart of the apparatus - the Penning and Ioffe traps - in much more detail. Fig. 2.1 gives a high-level overview of some of the most important parts of CTRAP.

### 2.1 Superconducting Solenoid

Surrounding the trap apparatus is a persistent superconducting magnet, which provides the uniform 1 T field for the Penning trap. The field is uniform to within a factor of  $1.5 \times 10^{-4}$  over a cylindrical volume situated at the center of the winding of length 50 cm and diameter 5 cm. In addition to its field homogeneity, the magnet

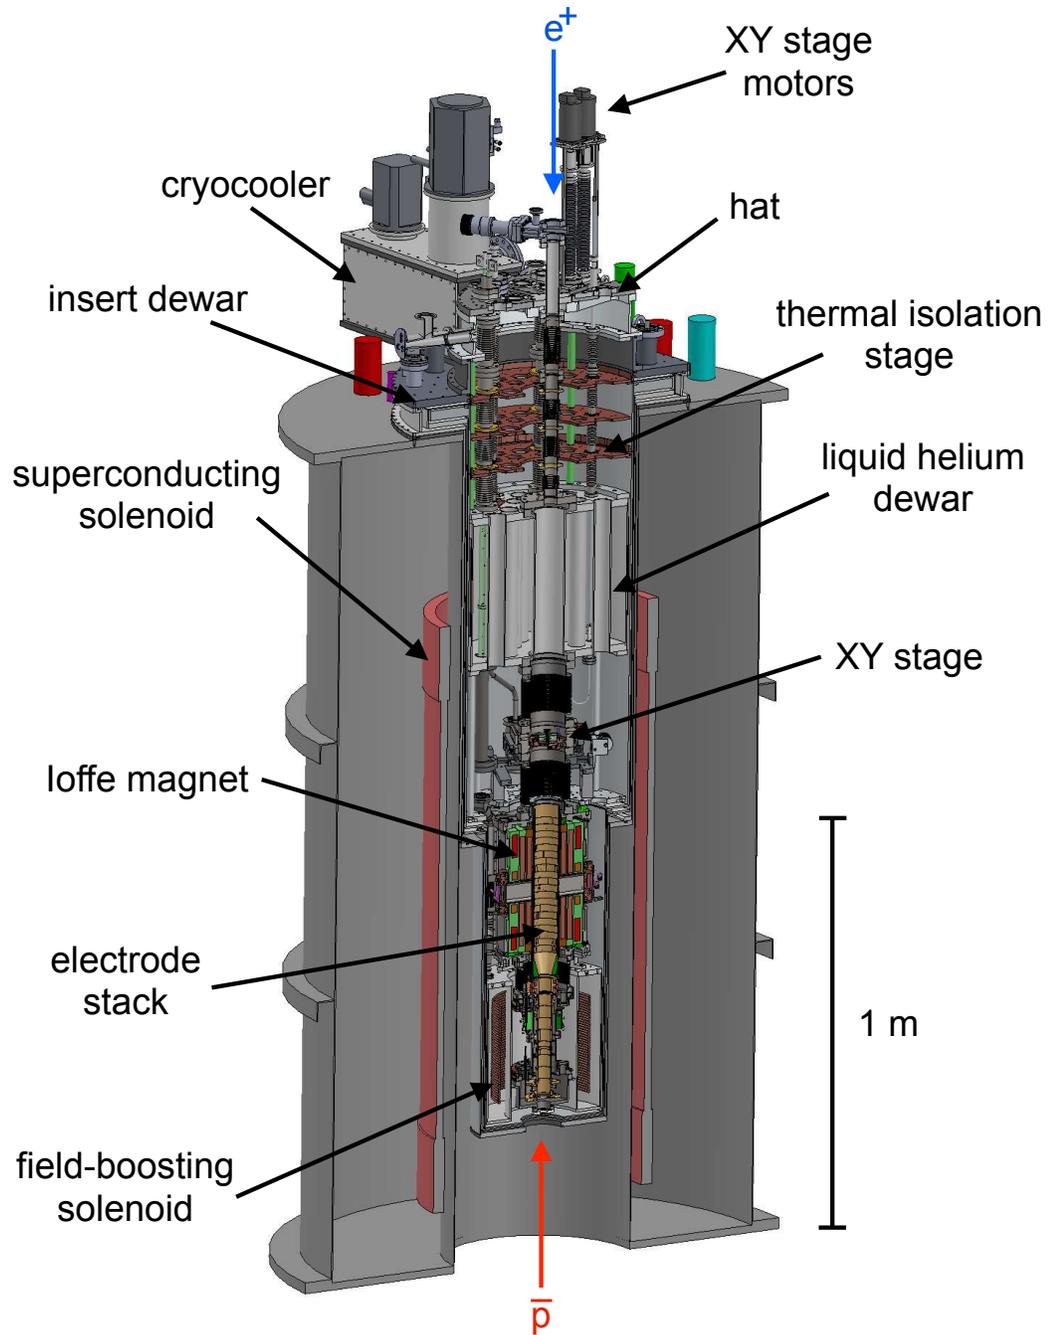


Figure 2.1: The CTRAP apparatus is shown mounted within the insert dewar, which, in turn, sits inside the large Penning trap magnet.

system has the feature of requiring very little maintenance. The 550 L helium space must be topped up once every 7 days to keep the coils immersed, provided that the surrounding liquid nitrogen shield is filled once every 10 days. The magnet operated without quenching during the time that antiprotons were available each year.

## 2.2 Cryogenic and Vacuum System

A well-controlled environment is crucial for creating trappable antihydrogen. An excellent vacuum is required to prevent annihilations between trapped antimatter particles and residual gas. Cryogenic temperatures are needed to maximize the number of antihydrogen atoms synthesized with energies low enough to be caught by the 400 mK deep Ioffe trap. CTRAP is designed to meet both needs.

Careful design has made it possible to keep CTRAP's particle traps at or below 4.2 K with minimal liquid helium consumption. A 48 L liquid helium dewar lies near the top of the apparatus. It conductively cools components below while also ensuring that the Ioffe and field-boosting solenoid coils remain immersed. A continuous fill system keeps this internal dewar's helium level to within a 5 L band by drawing from an external, 500 L storage dewar. This external dewar must be replaced once every two days.

A pumped helium system is in place to further cool the trap electrodes to 1.2 K [29]. Helium flows from the 48 L dewar downward, through a line of low thermal conductivity, and into a reservoir. This flow is controlled by a needle valve, which can be opened and closed from the top of the apparatus using a rotatable feedthrough. The reservoir is pumped through a line running to an external scroll pump. The

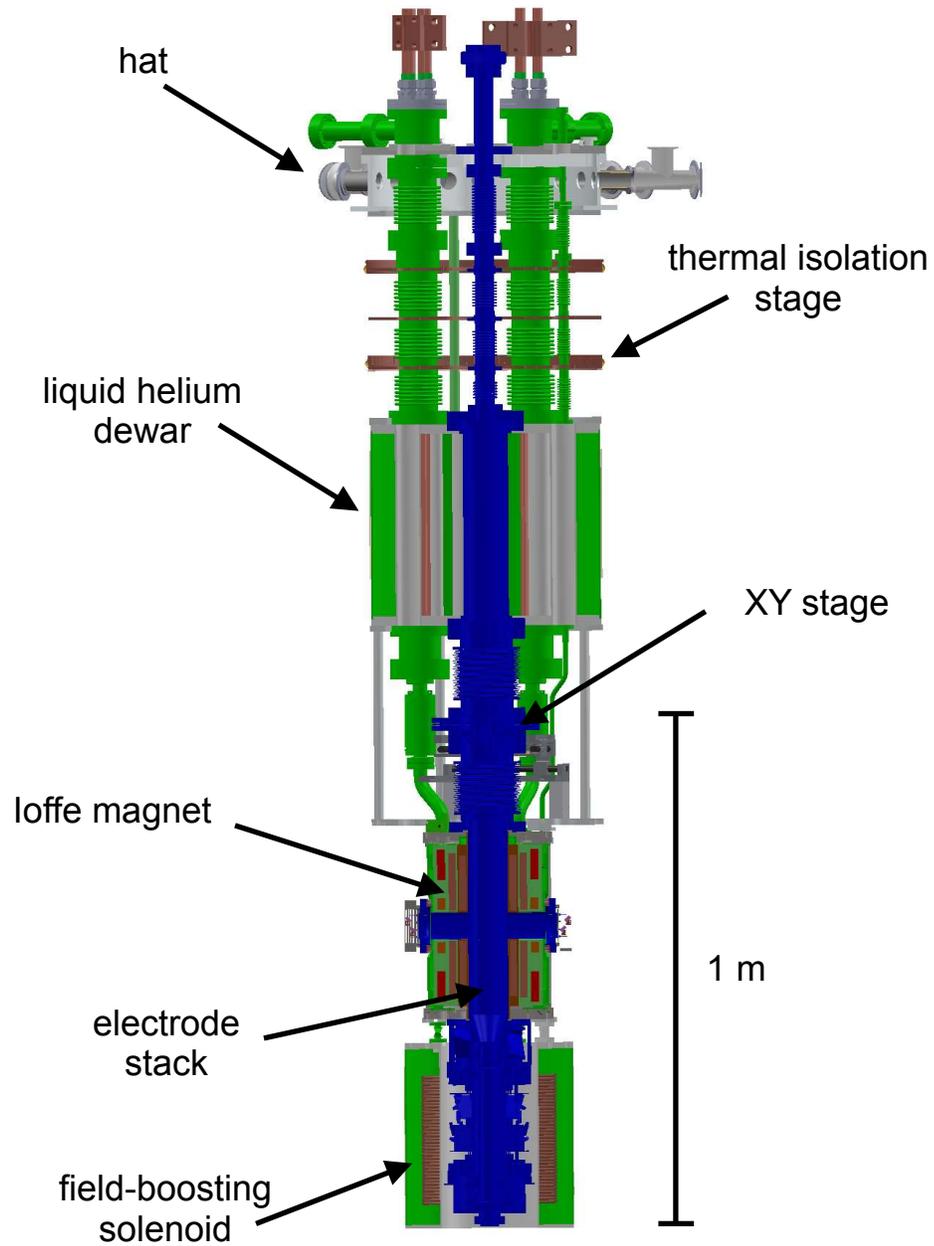


Figure 2.2: The experiment vacuum space (blue) and the helium space (green) are highlighted. The insulating vacuum space surrounds CTRAP, and is bounded by the insert dewar, which is not pictured.

temperature of the helium is thus reduced to around 1.1 K. The helium, which at this temperature is a superfluid, runs through a circuit consisting of a 290 cm long, 3.2 mm diameter tube, whose ends are connected to the top and the bottom of the reservoir. This cooling line's path is such that it passes by both the upper and lower sections of the electrode stack. Clamps ensure good thermal contact between the line and the titanium vacuum components to which the electrodes are fixed. Elsewhere, contact between the tube and 4.2 K components is avoided.

Heat loads on various parts of the apparatus determine whether they can be held at 4.2 or 1.2 K as well as the rate at which helium is consumed. The space around CTRAP is pumped to vacuum in order to insulate it from the walls of the insert dewar and the room-temperature hat. This also prevents ice from damaging sensitive components and obstructing laser paths. To reduce radiative heating, the walls of the insert dewar are cooled by pulse-tube refrigerators. The walls are divided into three thermally isolated layers, which are maintained at room temperature, 60 K, and 20 K, respectively.

The apparatus is suspended from its top, room-temperature portion, which will be referred to as the "hat". This gives a conductive path between room and cryogenic temperatures. G10 rods connect the liquid helium dewar to the hat, providing as much insulation as possible. Three copper thermal isolation stages, each covered in aluminized mylar, also lie in the space between these components to reduce the radiative heat load. The thermal isolation stages are fitted with flexible copper pieces that make contact with the insert dewar walls, keeping each of them at 20 K as well. Finally, the vacuum components and electrical wiring are clamped to each

stage, reducing their temperatures as much as possible before they make contact with colder parts of the apparatus.

It's also important to maintain separation between components at 4.2 K and the electrode stack assembly, which is cooled to 1.2 K. The upper section of the electrode stack fits tightly within the inner bore of the Ioffe trap. While it's impossible to avoid contact, electrodes are surrounded by a G10 or linen phenolic sleeve, which effectively insulates them both thermally and electrically. The only other mechanical connection between the electrode stack and a 4.2 K part is made by a set of edge-welded bellows with 37 convolutions and a wall thickness of 0.1 mm. As the weight of the electrodes are supported by this connection, G10 rods join the upper and lower flanges of the bellows. The only other conductive path to warmer surfaces goes through the electrical lines. Stainless steel micro-coax and thin constantan wires are used to reduce thermal conduction of heat to the cold apparatus.

While the insulating vacuum separating CTRAP from its surroundings has been discussed, we have not yet mentioned the separate vacuum space occupied by trapped particles, commonly referred to as the "experiment vacuum". This vacuum is shared with the positron system. Within CTRAP, this volume is split by the XY stage, which has a single central 1.5 mm diameter hole. The volume below the stage ranges in temperature from 1.2 K to 4.2 K while the trap walls above the stage range from 4.2 K near the hole to room temperature at the top of the trap. The separation between these regions allows for more effective cryopumping in the lower volume, where particles may be trapped for 10s of minutes. While no direct pressure measurement has been taken under operating conditions, antiproton survival was analyzed in a sim-

ilar apparatus, leading to an upper bound in pressure of  $5 \times 10^{-17}$  Torr [30]. While similar analysis was never carried out in CTRAP, the independence of particle count and confinement time has been verified in the trap for confinement times of up to two hours for electrons, positrons, and antiprotons.

To clarify the above discussion, Fig. 2.2 shows the experiment vacuum space in blue and the space in which helium is allowed to flow in green. Note that the outer walls of each space have been colored, so some features external to a given space may appear to be within it.

## 2.3 Ioffe Magnet

The Ioffe magnet performs the central role in antihydrogen experiments of trapping the atoms. For CTRAP, it also represents the principle upgrade over its predecessor, BTRAP. The motivations for the upgrade, the struggles encountered along the way, and the results are presented in a later chapter. For now, we will quickly summarize the latest version.

The Ioffe magnet consists of four NbTi superconducting coils embedded in a G10 block, which was fabricated around them. Two of the coils each resemble pairs of solenoids, situated symmetrically around the central plane of the magnet and concentric with the magnet's axis. These "mirror" coils are responsible for tailoring the shape of the magnetic field along the axis. The other two coils consist of 4 and 8 "racetrack" sub-coils, known as the quadrupole and octupole coils, respectively. These form magnetic fields with different radial dependences. Together, these four coils can be used to construct a wide variety of magnetic minimum atom traps.

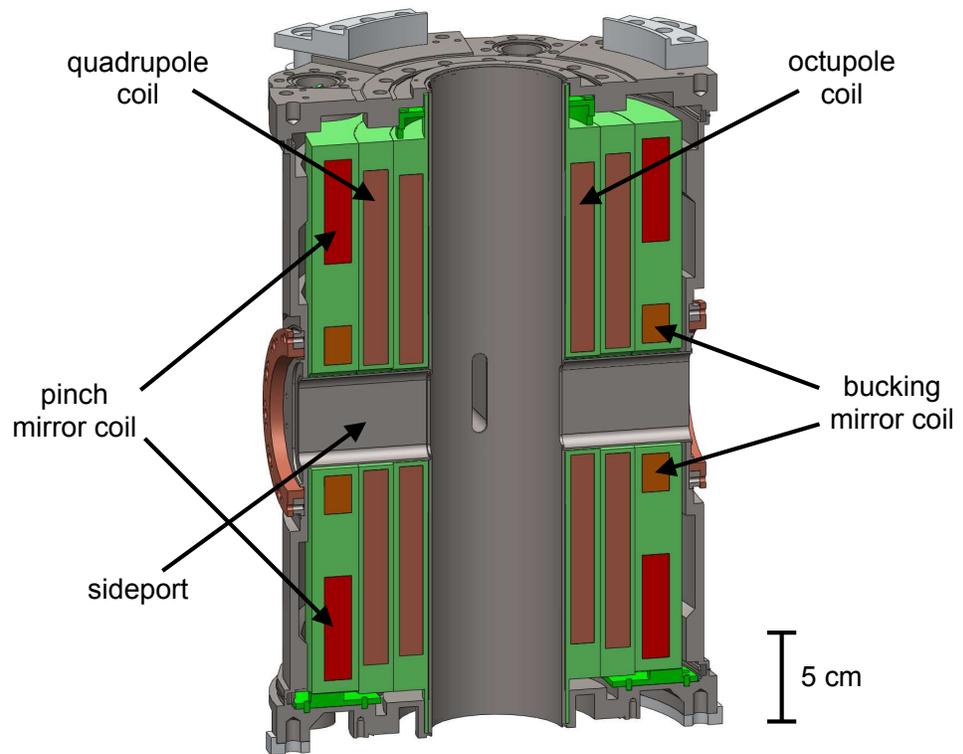


Figure 2.3: The coils of the Ioffe magnet are shown along with its surrounding enclosure and side-access ports.

In addition to housing and stabilizing its coils, the magnet's G10 body is designed to facilitate the types of experiments envisioned in CTRAP. A central bore provides room for a concentric electrode stack, which is used to mix antiprotons and positrons near the magnetic minimum, producing antihydrogen. Sideports, crossing the magnet body radially, give access to lasers for performing cooling and spectroscopy on trapped atoms.

In order to keep the coils superconducting, they must remain immersed in liquid helium. To this end, a titanium alloy enclosure surrounds the Ioffe magnet. This structure was designed to maximize its inner bore's radius, preserve as much volume in the sideports as possible, and interfere with static and changing magnetic fields as little as possible. The challenges of meeting these criteria are described in a later, but in the end, the goals were met and the construction was successful. A cross-section of the magnet and its enclosure is shown in Fig. 2.3.

## 2.4 Electrode Stack

In order to produce antihydrogen, clouds of antiprotons and positrons must be caught, manipulated and merged. CTRAP's electrode stack is designed for flexibly carrying out these tasks, granting plenty of degrees of freedom for modification and tuning. The stack is divided into two parts, separated. Fig. 2.4 serves as a useful reference for the following discussion.

The lower stack is comprised of ten 18 mm radius ring electrodes, a three-electrode high-voltage assembly, and a beryllium degrader. The primary purpose of this segment of the electrode stack is to trap and tailor clouds of antiprotons. To this end, the

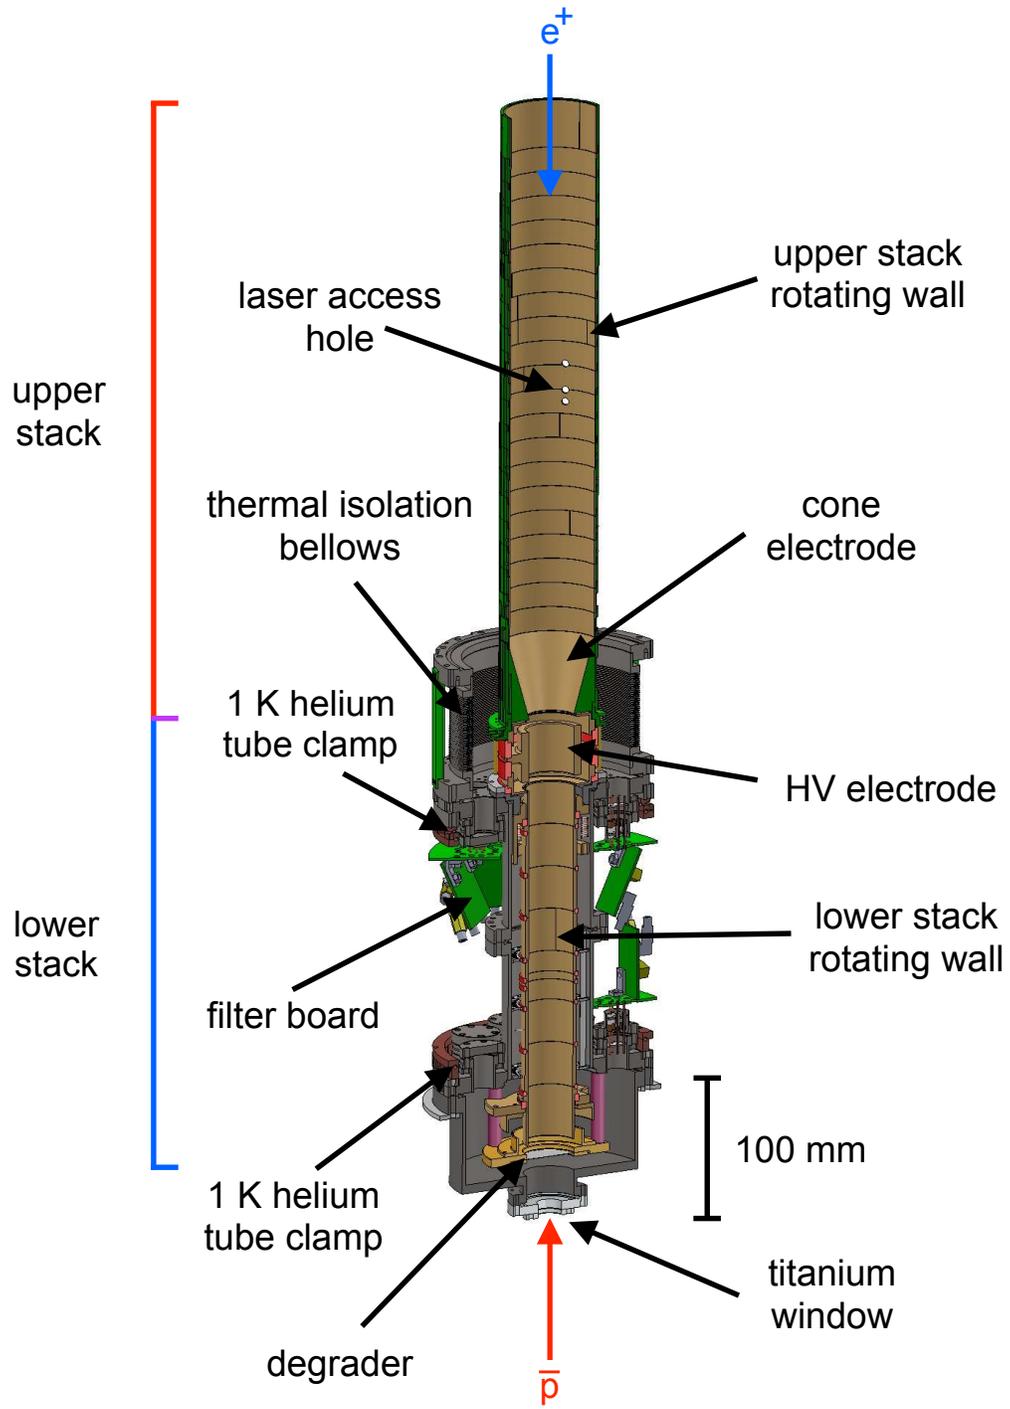


Figure 2.4: Upper and lower electrode stacks

degrader and the central electrode of the high-voltage assembly can each be biased to -5 kV. This allows high-energy antiprotons to remain trapped long enough to cool, through collisions with electrons, into a shallower potential well. The lower stack is also surrounded by a field-boosting solenoid (see Fig. 2.1), which increases trapping rates and decreases cooling times. Finally, a rotating wall electrode in this portion of the trap makes it possible to compress antiproton clouds, which is beneficial for antihydrogen production. These procedures are described in much more detail in a later chapter.

The upper stack is made up of twenty one electrodes of radius 34.1 mm, sitting on a single cone electrode. The cone smooths the transition between the narrow lower stack and the wider upper stack. The principle purposes of the upper stack are to catch and shape positron plasmas and to merge clouds of antiprotons with positrons to create antihydrogen. As is the case in the lower stack, a rotating wall electrode allows us to reduce the radius of trapped positron plasmas. To maximize the antihydrogen trapping rate, the Ioffe magnet and its enclosure surround the majority of the upper stack. This allows antihydrogen's precursors to be combined near the atom trap's minimum. As is shown in a later chapter, the location of this minimum depends on the currents we use. Vertically aligned electrode holes allow laser access to the minima of the most commonly used current configurations. The pattern is aligned with the Ioffe trap's sideports and repeats every 90 degrees.

Several secondary components displayed in Fig. 2.4 are crucial to the electrode stack's operation. The electrodes lie in the experiment vacuum, which, for most of the upper stack, is bounded radially by the inner bore of the Ioffe magnet's enclosure.

The thermal isolation bellows are mounted to the bottom of the enclosure, serving the triple purpose of supporting the weight of the electrode stack, continuing the experiment vacuum, and providing thermal isolation between the 1.2 K electrodes and the 4.2 K magnet. The connected volume continues down to the 10  $\mu\text{m}$  thick titanium window, whose purpose is to admit antiprotons from the AD while separating the experiment and insulating vacuums.

Other important peripheral components include the filter boards. Electrical biases, waveforms, and pulses are sent into CTRAP through feedthroughs at the hat, travel down through the insulating vacuum along either constantan or stainless steel coaxial lines, and arrive at the filter boards. From there, the signals are filtered and sent through a second set of feedthroughs into the experiment vacuum, where they travel directly to the intended electrodes. More details about the boards will be given in the discussion on the electrode stack enclosure's construction.

## 2.5 XY Stage

For the different tasks it performs, CTRAP requires different on-axis devices. For this purpose, a moveable stage that can be translated in two orthogonal directions sits just above the electrode stack. The stage, shown in Fig. 2.5, carries a flange with three components that need to be centered along CTRAP's axis during different phases of the experiment. In the center lies a hole of diameter 1.5 mm and length 20 mm. As noted before, this plays an important role in isolating the fully-cryogenic vacuum below. Activated charcoal is fastened nearby to cryopump background gas in the vicinity. Both positrons and the UV laser used to load electrons into the Penning

trap pass through the aperture, and surrounding Faraday cups assist with steering. To one side of the hole lies a 1"  $\text{MgF}_2$  window, capable of transmitting Lyman- $\alpha$  light for laser cooling. On the other side is a phosphor screen used for imaging trapped plasmas.

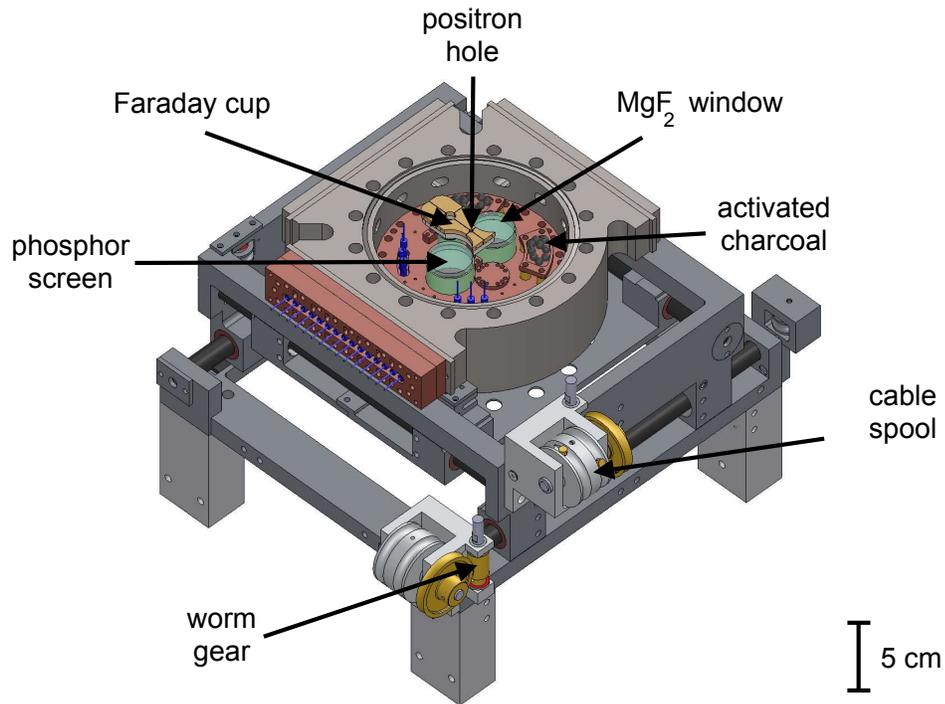


Figure 2.5: CTRAP's translatable XY stage

Moving the stage, which lies deep within the insulating vacuum at a temperature of 4.2 K, is a nontrivial challenge. Motors above the hat turn G10 drive shafts, which extend down to the XY stage's worm gears. These mate with another set of gears, which turn spools of copper wire, pulling the stage along Frelon rods. A 5 k $\Omega$  potentiometer along each axis is used to determine the approximate position of the

stage. For more precise alignment, LEDs mounted to the fixed part of the stage line up with photodiodes on the movable sections when one of the three components described above is on-axis. Pinholes in front of both the LEDs and the photodiodes restrict the length over which the photodiodes respond to the LEDs, allowing alignment to about  $\pm 0.0045$ ". Finally, flexible edge-welded bellows above and below the stage make it possible to maintain a separation between experiment and insulating vacuums.

## 2.6 Detectors

Scintillating detectors surround CTRAP and record annihilations between antimatter and the trap walls. As a diagnostic, these detectors are indispensable. They make it possible to destructively count trapped antihydrogen atoms and antiprotons. Also, positron and antiproton losses that would otherwise go unnoticed during trap manipulations are easily seen and quantified.

As shown in Fig. 2.6, CTRAP is equipped with two sets of detectors. Four layers of 3.8 mm BICRON BCF-12 fibers surround the upper electrode stack, just outside the insert dewar. Two of these layers feature 224 vertical fibers each, with the layers angularly offset by a fiber radius from one another to fill any gaps between the fibers. The other two are helical, with each fiber spanning an arc of  $158^\circ$ . These layers each have 168 fibers and are also displaced from one another by a fiber radius. In principle, the coordinate system arising from the straight and helical layers allows us to determine exactly where particles pass through the fibers.

The second set of detectors lies outside the persistent superconducting solenoid. Rectangular BICRON BC404 paddles are arranged in two layers, separated by 7 mm.

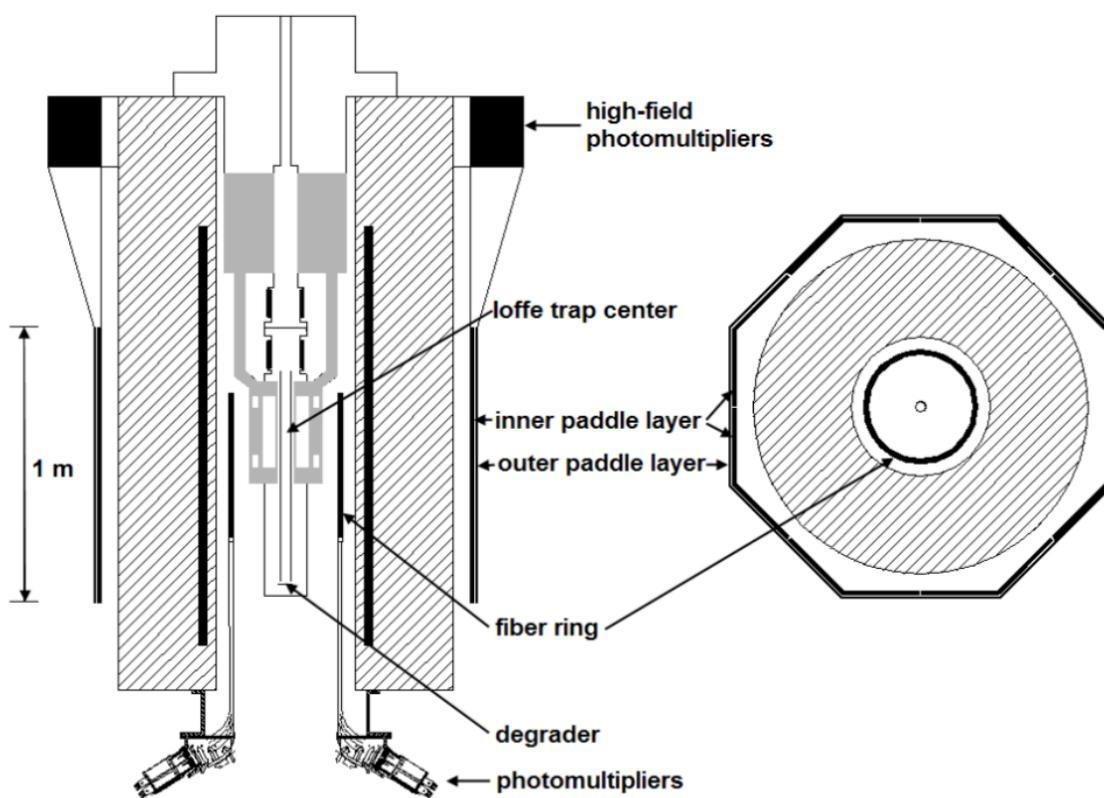


Figure 2.6: Locations of CTRAP's paddle and fiber detectors

The outer layer consists of 8 paddles while the inner layer has 16. Since the solid angle subtended by each of these layers is about the same, electronic noise is reduced by recording signals from the paddles only when two radially adjacent paddles are hit within 40 ns of one another.

To determine the efficiency of components in the detector system, cosmic ray data is collected. Events involving hits on diametrically opposed detectors are flagged, and the proportions of these events seen by intermediate detectors are calculated. The GEANT4 toolkit [31] is used to simulate antiproton annihilations on the gold-plated electrodes of the upper stack. The resultant minimum ionizing particles (MIPs) are tracked to determine the probability that an antiproton annihilation leads to an MIP passing through a given detector. Combining these results, we find detection efficiencies of 68% and 87% for the paddles and fibers, respectively. Filtering events for coincidences between one set of paddles and two fibers results in a detection efficiency of 54%, but the background rate plummets from hundreds of Hz to 41 Hz, allowing for much higher confidence that a given signal comes from an antiproton.

## 2.7 Laser Paths

One of the primary goals of this experiment is to perform 1S-2S spectroscopy on antihydrogen atoms. In order to make this measurement more precise, we also intend to laser cool the atoms. In pursuit of both of these goals, CTRAP is fitted with a number of optics, both at room temperature and at 4.2 K, which make laser access at these frequencies possible.

The hat is fitted with four flange-mounted  $\text{MgF}_2$  windows (Fig. 2.7), offset from

one another by 90 degrees. Transmission through these windows is about 60% for the cooling laser's frequency (121 nm) and upwards of 90% for the spectroscopy laser's frequency (243 nm). A 4 atm pressure differential rating across these components allows them to separate the vacuum of the insert dewar from atmosphere without fear of failure.

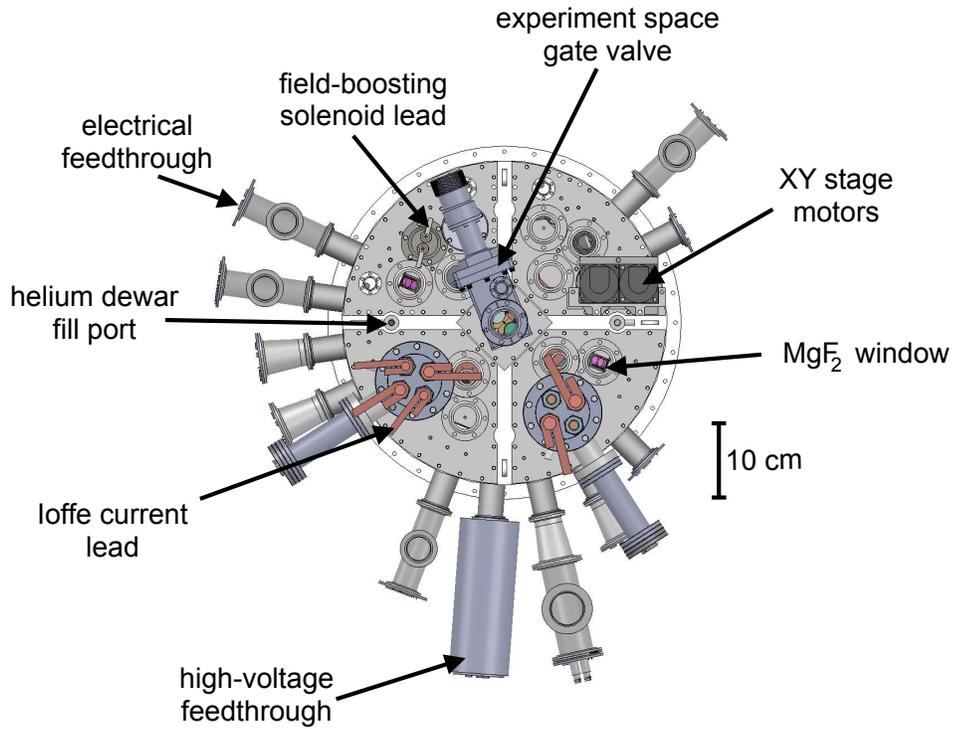


Figure 2.7: CTRAP's hat is shown with several important components highlighted. The sideport-mounted mirrors are visible through two of the  $\text{MgF}_2$  windows. A portion of the XY stage can also be seen through the central gate valve.

Light from these windows travels down through the insulating vacuum. Each flange covering a Ioffe enclosure sideport has a mirror assembly mounted to it, as

shown in Fig. 2.8. This consists of a titanium track and up to two mirror flexure mounts, offset from one another so both are visible from above. The mirrors can be mounted at any vertical position along the track and, once mounted, can be tilted along two axes using M2 brass screws. The surfaces of these mirrors are coated with a  $\text{MgF}_2$  film overlaying aluminum (Acton Optics coating #1200), which allows for up to 83% reflectance at 121.6 nm. Originally, the mirrors were fixed to their mounts by small drops of epoxy. However, the epoxy ran before drying on some pieces, bonding larger surface areas together, which cracked the mirrors. A new design, shown in Fig. 2.9, features a plate over the mirror's edges, fixed to the body of the flexure mount by BeCu springs. This version has proven to be much more robust under thermal cycling.

After a mirror reflects a vertical beam horizontally, toward the Ioffe trap, the light passes through another  $\text{MgF}_2$  window. This window separates the insulating vacuum space from the experiment space. To prevent damage from thermal cycling, the cryogenic seal is made by pressing indium between the window and the copper sideport flange. The sideport flange also uses indium to form a seal with the body of the Ioffe enclosure. Once inside the sideport, the light has an unobstructed path to the center of the trap. Tapped holes on the inside of the sideport flange and on a titanium piece epoxied to the outside of the electrode stack offer the option to mount additional optics inside the sideports.

The path from one window on the hat, through the Ioffe trap, and back out the opposite hat window is several meters long. This makes steering a laser beam, which must hit two 0.5" mirrors and pass through two 0.2" diameter apertures, a challenge.

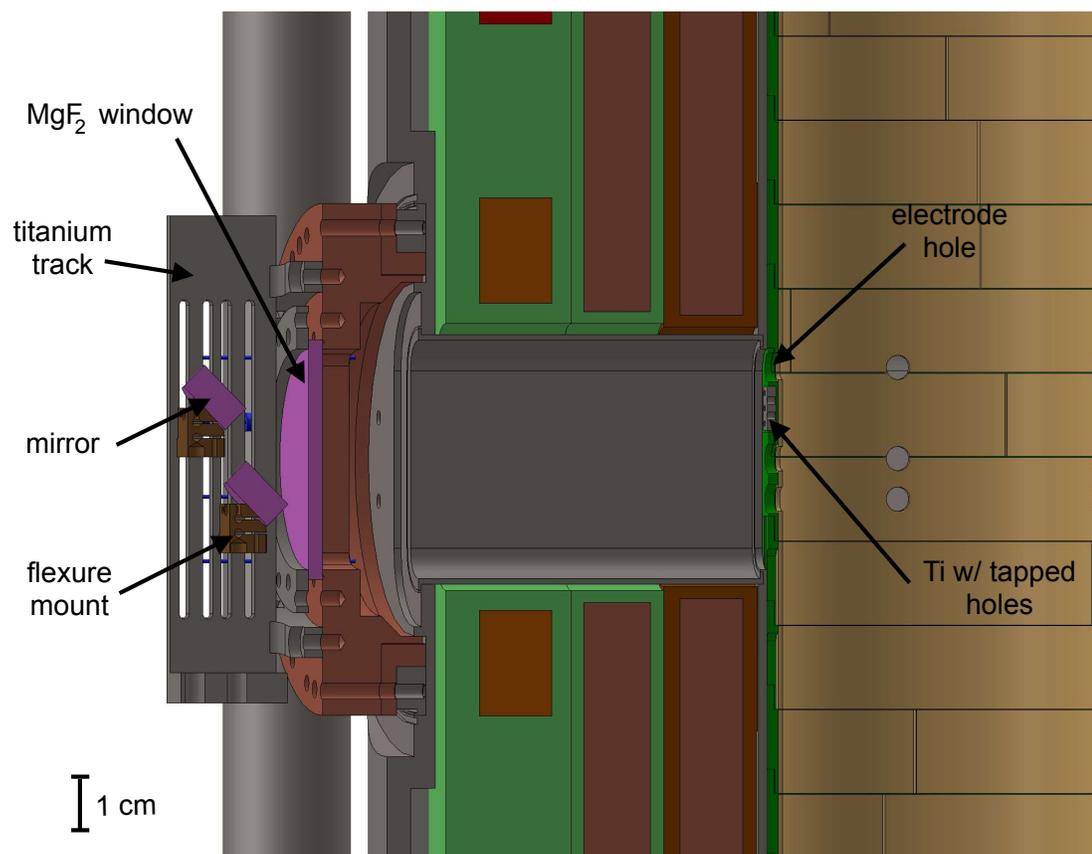


Figure 2.8: Cross-section of the upper stack, Ioffe trap, and a mirror assembly mounted to a sideport flange

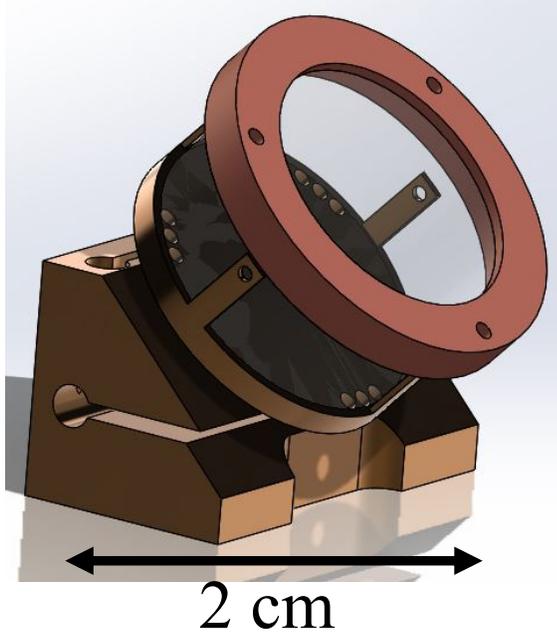


Figure 2.9: Cross-section of the upper stack, Ioffe trap, and a mirror assembly mounted to a sideport flange [32]

Nevertheless, we developed a procedure which has proven successful over multiple cooldowns. A 408 nm tracer laser is mounted on the hat window opposite the entry point of the laser of interest. With CTRAP outside of the insert dewar, the mirror mounts are adjusted until the the tracer arrives at the other hat window. When the trap is moved to the insert dewar, and later, as it cools down, the tracer is adjusted so its exit point stays in about the same place. Once the components of the apparatus have reached their. final temperatures, the laser we were interested in sending through this path is overlapped with the tracer beam, adjustments are made to account for dispersive optics along the path, and the tracer is replaced by a detector specific to the laser being used.

At this time, the only laser path that has been fully constructed belongs to the Lyman- $\alpha$  cooling laser. The major components of this custom-built laser are housed

in a temperature-regulated Faraday cage adjacent to our experimental zone. This system produces pulses at 30 Hz of up to 21 mJ of 365 nm light [33], which is guided to CTRAP's hat. Instead of being mounted directly to the hat, the MgF<sub>2</sub> window at the laser's input accounts for one end of a Kr/Ar gas cell, situated above the hat. The 365 nm light is tripled inside this cell, a SiO<sub>2</sub>-coated optic transmits the remaining 365 nm light and reflects about 17% of the 121 nm light, and the reflected light propagates down toward a sideport mirror, being collimated by a lens along the way. This path, along with a similar axial path, is shown in Fig. 2.10. Currently, 40 nW of the 5.7  $\mu$ W of Lyman- $\alpha$  light produced by the gas cell can be coupled into the trap, but improvements are being explored that should boost this by at least an order of magnitude. Still, the 40 nW we are currently able to provide outpaces the power used for a previous hydrogen laser cooling experiment and a more recent 1S-2P spectroscopic measurement with antihydrogen [34, 35].

## 2.8 Positron Source

Positrons for antihydrogen synthesis are captured and sent to CTRAP from a separate apparatus located across the AD, as shown in Fig. 2.11. A <sup>22</sup>Na source with a solid Ne moderator produces positrons with a few eV of energy. These are captured in room-temperature positron accumulator, where they are collisionally cooled into a Penning trap by N<sub>2</sub> and SF<sub>2</sub> at low pressures [36, 37]. The radius of the positron cloud is kept to 3 mm by a rotating wall drive.

The process described above happens every 30 seconds, and is synchronized with the AD's cycle to make the magnetic environment more predictable. After accumu-

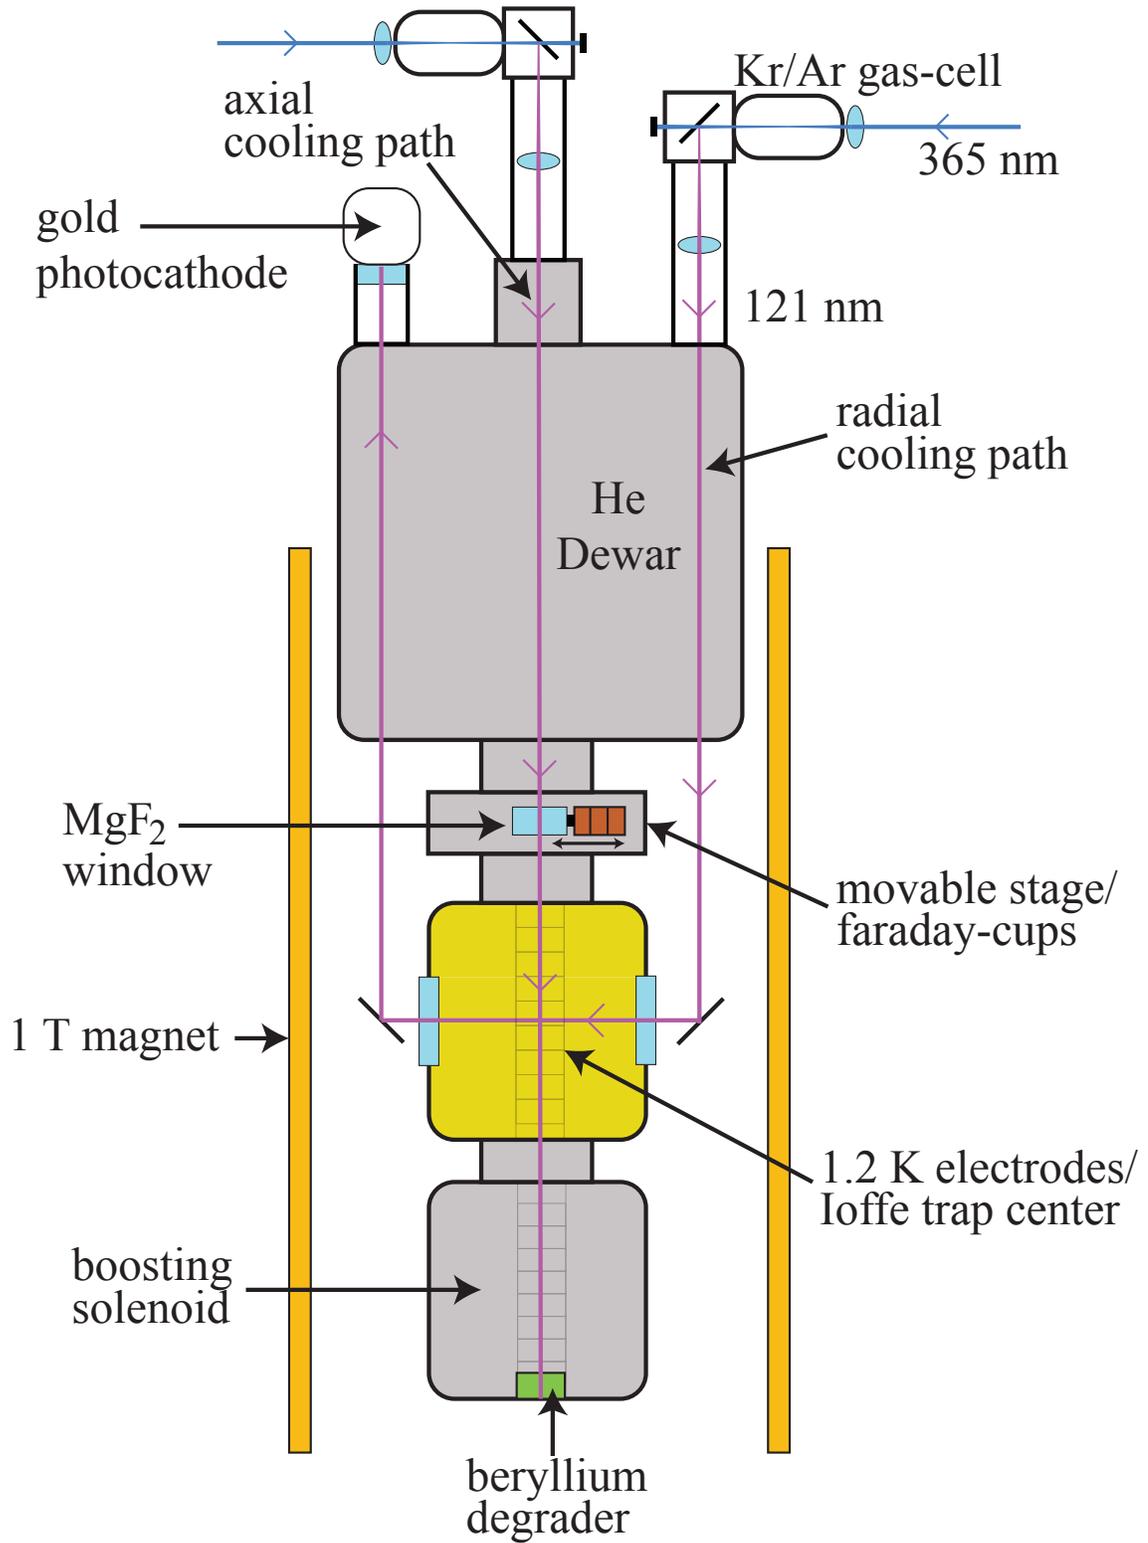


Figure 2.10: Laser paths in CTRAP [32]

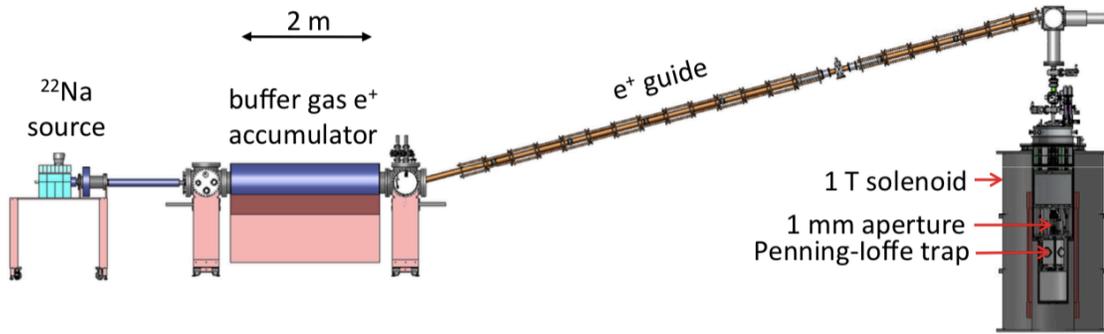


Figure 2.11: The positron source, accumulator, and transfer line are shown along with the CTRAP apparatus, which is mounted inside the persistent solenoid.

lation, positrons are pulsed into a 10 m transfer line, which uses a set of more than 70 magnets to guide the positrons into CTRAP. In the end, 3 million positrons are delivered every 30 seconds, and about half of these are captured.

## 2.9 Summary

It has been shown in this chapter that the systems built into CTRAP make it well-equipped for antihydrogen production, trapping, and, soon, laser manipulation. Vacuum and cryogenic spaces strengthen one another while providing the optimal environment for forming and mixing cold plasmas. The electrode stack and the Ioffe magnet perform the central tasks of capturing and manipulating antihydrogen atoms and their constituents. While some details have been given on these trapping components, the next three chapters explain them more thoroughly, including upgrades made in recent years and the advantages they bring. The detector system gives information on antimatter annihilations with trap walls, making it a useful diagnostic when performing manipulations with antiprotons and positrons while also allowing

us to count trapped antihydrogen atoms. The XY stage gives us the flexibility to use different components on the trap axis during different stages of an experiment. Finally well-designed laser paths make both laser cooling and spectroscopy of trapped anti-atoms a near-term possibility. It is on this collection of systems that ATRAP's ambitions for probing fundamental physical questions with antihydrogen lie.

# Chapter 3

## Ioffe Magnet

The Penning trap effectively confines charged particles, but the ultimate goal of the CTRAP experiment is to trap and analyze uncharged antihydrogen atoms. Since electric potentials are ineffective at confining even low-energy atoms, the magnetic moments present in some antihydrogen states allow for magnetic trapping. In this section, we discuss the magnetic Ioffe trap that holds antihydrogen atoms in the CTRAP apparatus.

### 3.1 Antihydrogen in Magnetic Fields

A full treatment of antihydrogen's behavior in a magnetic field is presented in 7.1. Here we list only the facts relevant to the kinematics of anti-atom confinement.

There are four states in the 1S manifold. Two have energies that increase with magnetic field, while the other two have energies decreasing with field. Energies as a function of field are shown in Fig. 3.1. Since static fields can only produce a

magnetic minimum [38], only states with energies positively correlated with field may be trapped.

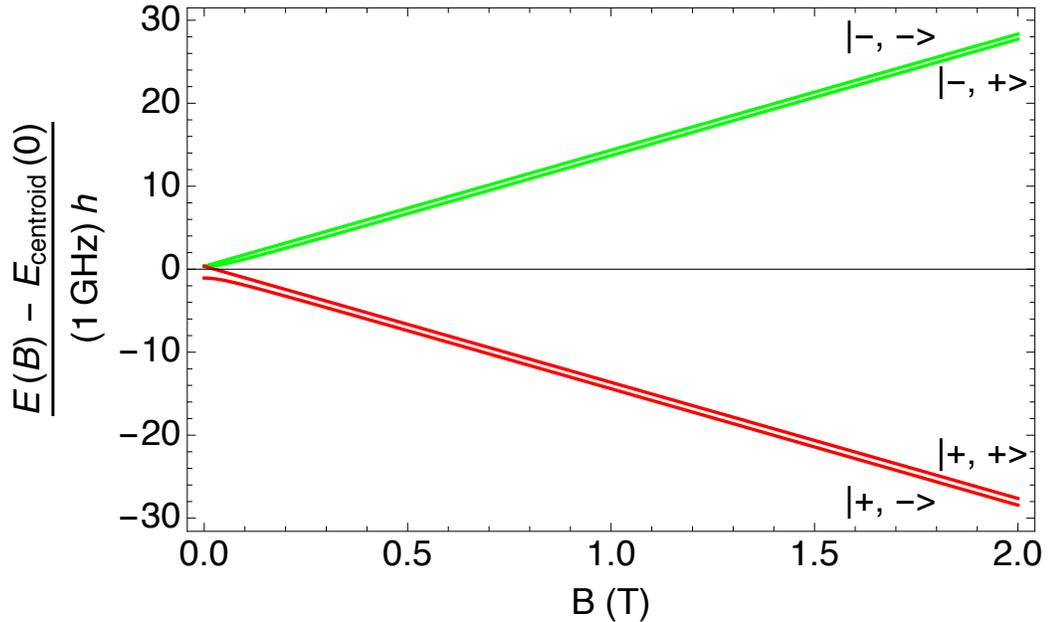


Figure 3.1: Energies of trappable (green) and untrappable (red) 1S states relative to the 1S zero-field centroid

From Fig. 3.1, states appear to have energies linearly dependent on field. While this is not strictly true (see Eqns. 7.5-7.6), the linear approximation fails by less than 2 parts in 1000 for the fields present in CTRAP’s Ioffe traps ( $B > 0.9$  T). We can thus further approximate the potential energies of particles in different states by

$$U(B) = \pm\alpha B, \quad (3.1)$$

where the sign depends on the state (Fig. 3.1) and  $\alpha \approx (14 \text{ GHz/T})h$ . For  $E = \frac{3}{2}k_B T$ , the energy changes at the rate 0.45 K/T. So chosen,  $\alpha$  differs from state-specific linear approximations by less than 0.2% for fields of interest.

## 3.2 Ideal Ioffe Traps

Ioffe traps represent one family of current configurations capable of creating local magnetic minima. Ideally, a Ioffe trap is formed by  $2n$  parallel current bars and a pair of circular mirror coils. We use a Cartesian coordinate system in which the bars run parallel to the  $z$ -axis and intersect the  $xy$ -plane at points  $(\rho_0 \cos k\pi/n, \rho_0 \sin k\pi/n)$ , where  $k$  runs from 0 to  $2n$ . Neighboring bars have equal currents that are oppositely directed. On the other hand, the  $z$ -axis coincides with each mirror coil's symmetry axis. The mirror coils lie on opposite sides of the  $xy$ -plane, equal distances away. The currents in these two coils flow in the same direction and have identical magnitudes. Examples of Ioffe configurations are shown in Fig. 3.2.

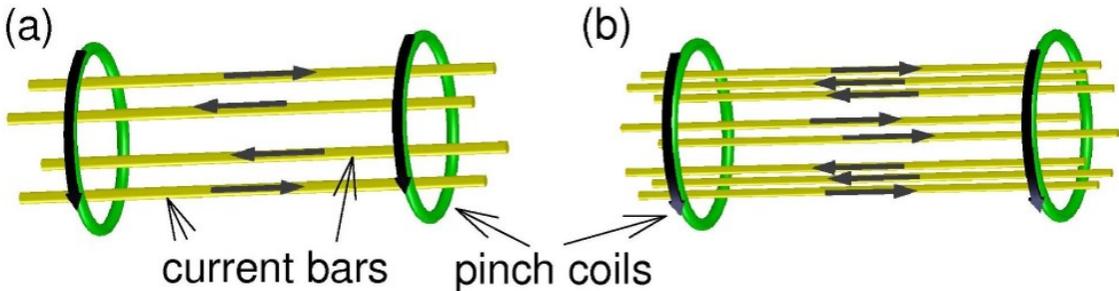


Figure 3.2: The magnetic minimum at the center of quadrupole (a) and octupole (b) Ioffe traps are produced by  $2n = 4$  and  $2n = 8$  current bars, respectively, plus orthogonal mirror coils.

The exact expression for the field arising from even a simplified Ioffe trap with infinite current bars is involved and not particularly enlightening [39]. However, expanding the contribution of the current bars to lowest order in  $\rho/\rho_0$  gives a simple, if only approximate, way to understand the basic differences between traps with

different numbers of bars. Performing this expansion yields

$$\vec{B}_{bars}(\rho, \theta) = B_0 \left( \frac{\rho}{\rho_0} \right)^{n-1} \left[ \cos(n\theta)\hat{\rho} - \sin(n\theta)\hat{\theta} \right]. \quad (3.2)$$

In the absence of other fields,  $|\vec{B}| \propto (\rho/\rho_0)^{n-1}$ . Consequently, adding current bars leads to a flatter minimum near  $\rho = 0$ , which is advantageous for laser operations (e.g. cooling or spectroscopy) which are sensitive to the Zeeman effect.

We have so far neglected the effects of the mirror coils. They contribute to both the radial and axial fields in ways that depend on their distance from the xy-plane, their radius, and the magnitude of the current flowing through them. Let's focus on how the radial dependence of  $|\vec{B}|$  is modified by the additional axial field. This field is taken to be constant ( $B_z\hat{z}$ ) and large compared to the radial field over the range of  $\rho$  of interest. While this implies certain constraints on the geometry and currents in an ideal Ioffe trap, these conditions are approximately met in CTRAP's Ioffe trap due to a background 1 T field. Under these assumptions, the magnitude of the total field is

$$|\vec{B}(\rho)| = \sqrt{B_z^2 + B_0^2 \left( \frac{\rho}{\rho_0} \right)^{2(n-1)}} \approx B_z + B_0 \frac{B_0}{B_z} \rho^{2(n-1)}. \quad (3.3)$$

The addition of an axial field, as in Eqn. 3.3, has three important consequences. First, the radial dependence near the origin is flattened. Second, the change in  $|\vec{B}|$  between the axis and any fixed radius for which our approximations hold is decreased. Since the potential energy of an antihydrogen atom varies linearly with field, this means that the trap depth is lowered by the presence of  $B_z\hat{z}$ . Finally, the added axial field prevents the total field magnitude from vanishing at the origin. This suppresses Majorana spin flips, which allow trapped atoms to escape [40].

While many approximations have been made in this discussion, qualitative statements about the effect of adding more current bars and an axial field generalize to CTRAP's Ioffe traps. Before presenting our current magnetic trap, it's worthwhile to briefly review an earlier version.

### 3.3 BTRAP's Ioffe Trap

In the years leading up to and including CTRAP's construction, our collaboration produced antihydrogen using a different apparatus, known within the group as BTRAP. This trap facilitated one of the earliest observations of trapped antihydrogen while also beating competing production rates by about a factor of 5 [1, 14]. It was also the only antihydrogen trap equipped with side windows, which can be seen in Fig. 3.4, allowing radial access to the atoms.

Despite the success, BTRAP's magnet has a few shortcomings, limiting its usefulness for precision antihydrogen measurements. To maximize trap depth with modest ( $< 100$  A) currents, a large number of windings are used. The resulting high inductance leads to a 15 minute turn-on time, during which large numbers of positrons and antiprotons are lost. Moreover, for our detectors to separate trapped atom signals from background when the magnet is shut off, the ramp-down time should be as short as possible. Methods were found to intentionally quench the coils, allowing for a  $< 1$  second discharge, but this causes other problems. Quenches may shorten the lifespan of the magnet. Also, following a quench, hours of recovery time are required before full current may be reached again. Finally, it seemed prudent to include the option of a flatter octupole ( $2n = 8$  current bars).

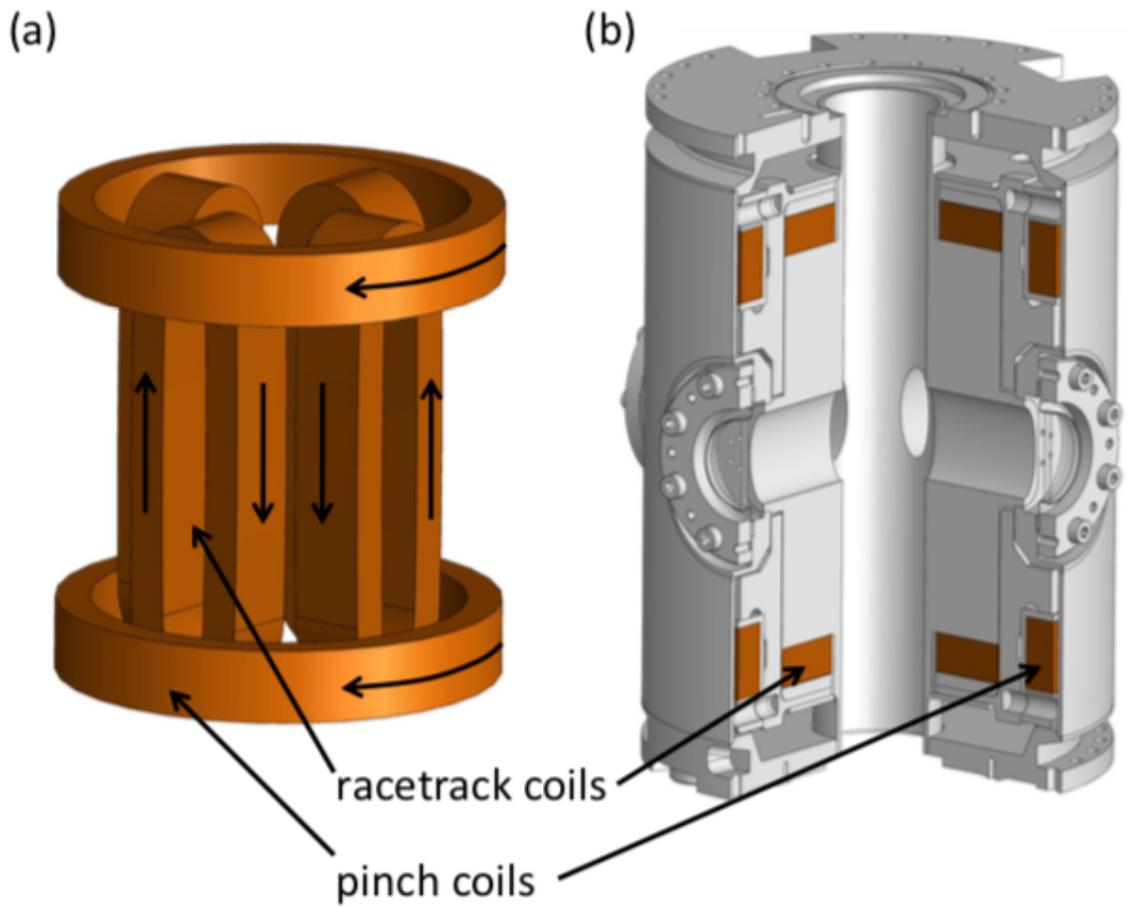


Figure 3.3: The BTRAP Ioffe trap shown as (a) the bare windings and (b) the windings within the surrounding titanium enclosure. [41]

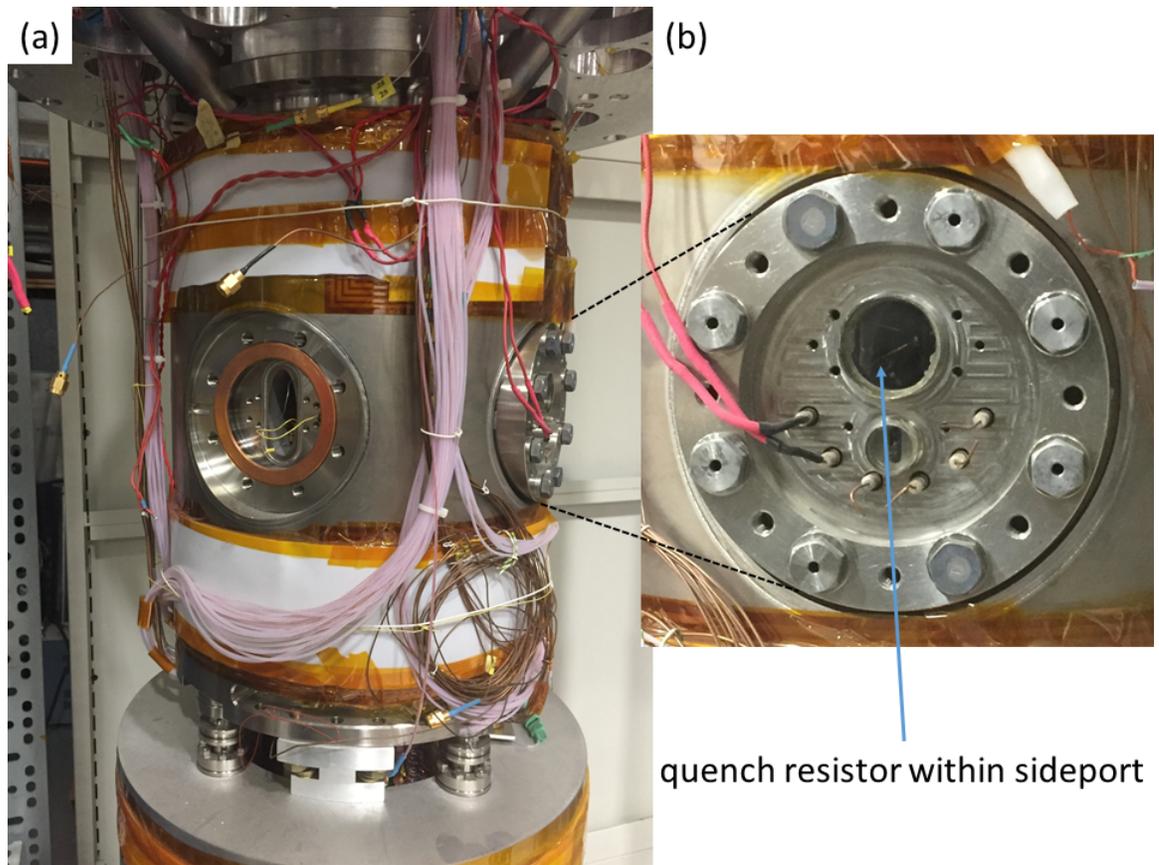


Figure 3.4: BTRAP Ioffe trap enclosure (a), with 2 of 4 side ports showing, and a  $\bar{p}$  loading solenoid below. A quench resistor attached to one side port is visible through the window in (b). Typically 10 W would cause a quench within 1 sec.

## 3.4 CTRAP's Ioffe Trap

With these considerations a new Ioffe magnet was constructed. Several significant design changes were implemented.

First, the new "CTRAP" magnet was given both octupole and quadrupole configurations of current bars. The positioning of the windings is shown in Fig. 3.5, with  $\otimes$  and  $\odot$  denoting relative current directions. One commonality between this magnet and its predecessor is the space made for radial access. In Fig. 3.3, the titanium enclosure has four sideports. Fig. 3.5 shows no enclosure, but the space between the coils and the x and y axes is clear in the cross-sectional view. The rolled out view makes this even more explicit. Near  $z = 0$ , the octupole coils bend away from lines of constant azimuthal angle at  $0^\circ$ ,  $\pm 90^\circ$ , and  $\pm 180^\circ$ . While allowing for three-axis laser cooling and spectroscopy transverse to the trap's axis, the choice to include side-access has consequences for the octupole field, which will be discussed later.

In addition to the octupole coil, a pair of new mirror coils was added to the original design in order to deepen the trap. Consequently, the CTRAP magnet has two pairs of mirror coils: the outer "pinch" coils, responsible for trapping atoms axially, and the inner, axial field cancelling "bucking" coils. The radius of each of these is the same, with the pinch coils lying further from the magnet's symmetry plane than the bucking coils. Fig. 3.6 shows the effect of the bucking coils on trap depth in the axial direction. It's also important to remember that that reducing  $B_z$  increases the radial depth due to the transverse multipole trapping fields (see Eqn. 3.3).

Adding coils changed the potentials we were able to make, but additional modifications were required to facilitate a quick field changes. For the BTRAP Ioffe magnet,

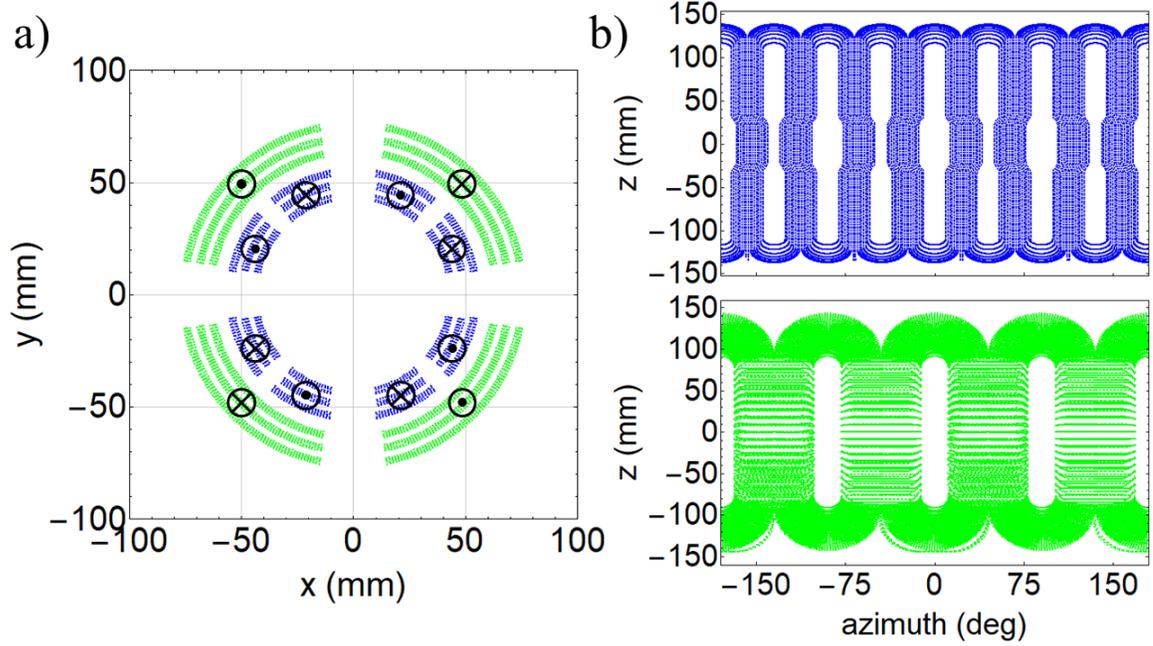


Figure 3.5: The (green) quadrupole and (blue) octupole coils are shown from two perspectives. In (a), a cross-section at the magnet's midpoint captures radial and angular spacing. In (b), the unrolled view shows both how current bars are joined and the existence of bends in the octupole current bars.

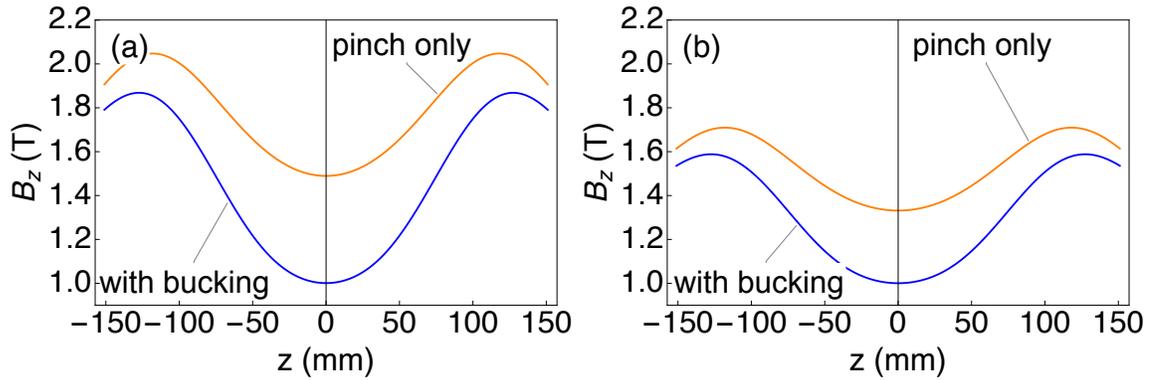


Figure 3.6: The effect of the bucking coils on the field along the trap axis is shown for the (a) quadrupole and (b) octupole current configurations. The 1 T bias field is included, but the multipole fields are not.

generation	trap	$I_{multipole}$ (A)	$I_{pinch}$ (A)	$I_{bucking}$ (A)	trap depth (mK)
BTRAP	quadrupole	68	80	-	250
CTRAP	quadrupole	470	310	264	350
CTRAP	octupole	680	210	179	270

Table 3.1: Currents and trap depths in the BTRAP and CTRAP Ioffe traps

each current bar consists of around  $2933 \pm 8$  windings [42]. On the other hand, the CTRAP magnet uses only 744 and 216 windings for each bar of the quadrupole and octupole coils, respectively. Consequently, the BTRAP quadrupole has an inductance of 3.3 H [43] while the CTRAP quadrupole and octupole coils have inductances of 113 mH and 19 mH, respectively. The details of charging and discharging the magnet will be discussed later, but it should be clear that the lower inductance in the new magnet allows current changes to proceed more quickly at a given voltage. This benefit comes at the cost of higher currents, however, as shown in Table 3.1. It ought to be noted that even with fewer windings per coil, CTRAP’s traps are deeper than what was achievable in BTRAP.

Fig. 3.7 shows the fields resulting from our design choices. A few interesting observations can be made from the plots. Distortion from bends around the sideports is obvious and severe. The field minimum is moved by several millimeters from the magnet’s symmetry plane ( $z = 0$ ), which, unfortunately, makes the optimal laser access point different for the octupole and quadrupole traps. More subtle distortion is seen along the radial direction in the form of off-axis minima. As small as the side wells appear to be, they are still large enough to have important consequences. The dependences of several quantities on field (at 1 T) are shown in Table 3.2. For instance,

the Doppler limit for laser cooling hydrogen (and consequently antihydrogen) is 2.4 mK [44]. This is equal to the potential energy change due to a 5.3 mT difference in field. The cooling schemes simulated in a later chapter don't quite reach this limit, but enough atoms are cooled into the off-axis regions to affect the results of subsequent laser experiments. As will be pointed out as they are presented, spectral lineshapes also bear features arising from the side wells.

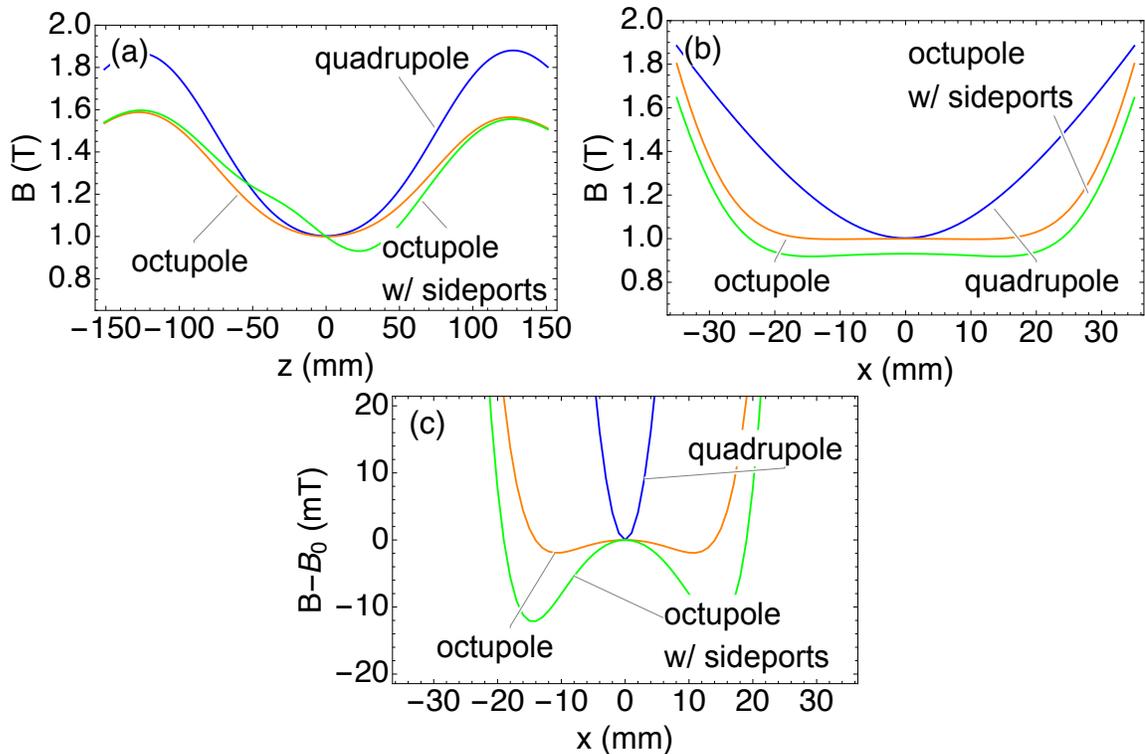


Figure 3.7: The field magnitudes for each trap are shown for (a)  $x, y = 0$  and (b)  $y = 0, z = z_{min}$ . The plot in (c) zooms in on (b), with each curve displaced so the on-axis  $B - B_0$  is zero. Note that curves for "octupole w/ sideports" represent the octupole trap discussed in this section while "octupole" curves describe a hypothetical version of CTRAP's octupole magnet for which the current bars are not bent around the sideports. The latter will be considered in later simulation chapters.

Despite its imperfections, the CTRAP magnet is a significant upgrade over its

quantity	change per mT
energy	0.45 mK
1S-2S frequency	0.96 kHz
laser cooling frequency	14 MHz

Table 3.2: Dependence of various antihydrogen quantities on local field

BTRAP counterpart. Four independent coils grant significant flexibility in choosing potential shapes. Furthermore, we are able to create deeper wells, allowing for higher antihydrogen trapping rates. Finally, by reducing the number of windings in each coil, we make it easier to charge and discharge the magnet quickly.

### 3.5 CTRAP Magnet Enclosure

Once the magnet was built, a significant amount of effort went toward constructing an enclosure compatible with the experiment’s goals. An enclosure is needed for two purposes. First, in order to maintain the magnet’s coils in a superconducting state, they must be immersed in liquid helium. It is seen from Fig. 2.2 that both the insulating and experiment vacuum spaces border this volume. The enclosure thus allows for submersion of the Ioffe coils in liquid helium without compromising either of the required vacuum spaces. Second, the magnet itself has no taps or other structures for mounting onto CTRAP or for the mounting of smaller parts (e.g. mirrors for spectroscopy) onto the magnet. Tapping the G10 magnet body directly is unacceptable because of the risk of hitting a coil, the possibility of weakening the structure surrounding the coils, and the ease with which threads might be stripped. Conse-

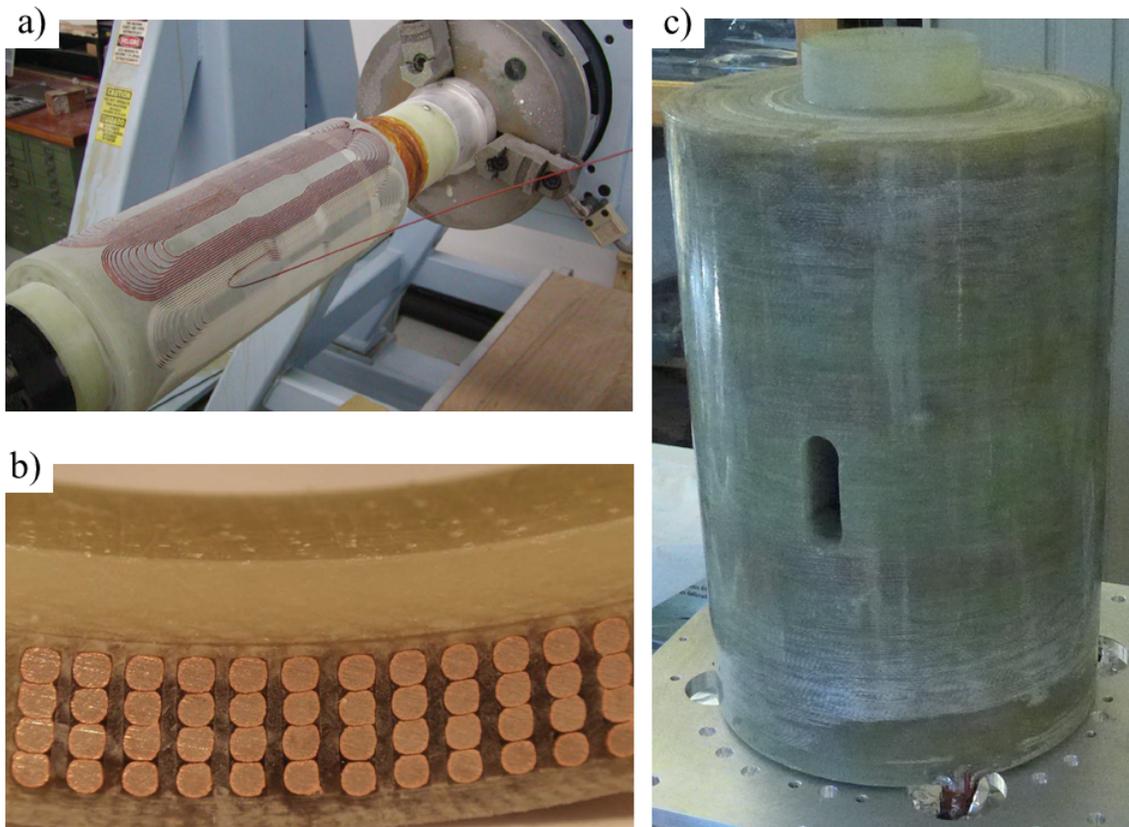


Figure 3.8: The Ioffe windings were set into grooves machined into a G10 epoxy fiberglass tube (a) and then epoxied into those grooves (b). The end result is a solid epoxy-fiberglass structure with embedded Ioffe coils, featuring an open tube in its central axis for Penning trap electrodes, 4 radial slots for sideports, and small channels between coil layers to improve LHe cooling (c).

quently, an enclosure separate from the magnet with at least some metallic parts is necessary.

While designing an enclosure, a number of possible problems were identified. Since we intended to ramp up and shut off the new magnet more quickly than was done in BTRAP, it was necessary to choose materials with high resistivities so eddy currents didn't slow either process. Another issue was that adding an extra layer of material to the inner magnet bore decreases the the volume in which atoms may be trapped. The atoms, which come from a  $\sim 17$  K distribution [14, 45], are confined in a sub-Kelvin trap. Narrowing the radius reduces the trap depth, especially for the octupole field, so it's important to keep the enclosure walls as thin as possible. Finally, it was desirable for the enclosure material to perturb the magnetic field as little as possible.

With these barriers in mind, we first designed a mixed-material vacuum enclosure, shown in Fig. 3.9. Plates at the top and bottom were made of aluminum with welded-in bi-metal CF flanges (Atlas Technologies). These plates were epoxied to a G10 shell surrounding the magnet, to which was epoxied G10 sideport tubes and copper rings around the external entrances to the sideports. These rings were to facilitate indium seals between copper flanges and the enclosure, which would both separate the external insulating vacuum from the internal experiment vacuum and allow for the mounting of mirrors and other small parts. While early tests were encouraging, we were unable to develop a robust method for cooling CTRAP to 4.2 K without developing leaks in this enclosure.

After the failure of the mixed-material enclosure, we set about making the best all-metal enclosure we could. Among the metals we considered, titanium alloys had

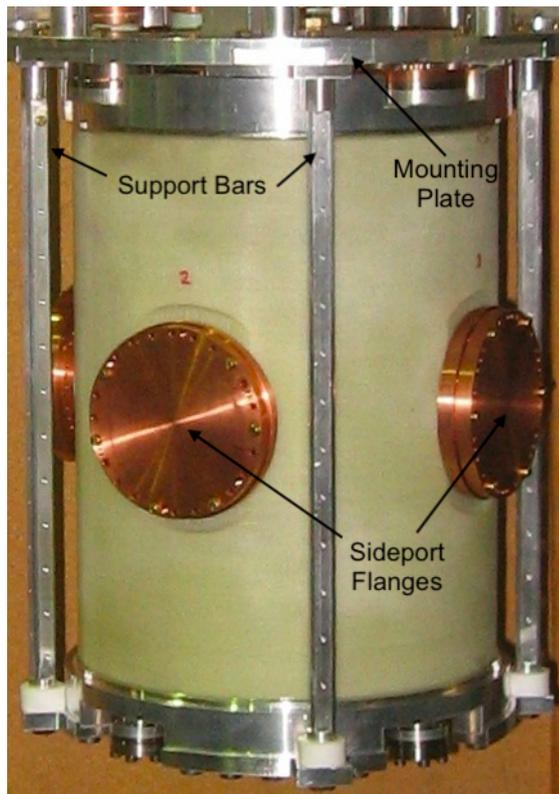


Figure 3.9: Mixed-material vacuum enclosure for the CTRAP magnet. The support bars hold the weight of the bottom plate and the antiproton catching magnet (not pictured) mounted below, taking stress off of epoxy joints.

Ti grade	Al %	V %	Sn %
5	5.5-6.75	3.5-4.5	< 0.1
6	4.0 - 6.0	< 0.1	2.0 - 3.0
9	2.5-3.5	2.0 - 3.0	< 0.1

Table 3.3: Impurity contents, by weight, of different grades of titanium [47]. Other elements are limited to 0.5% by mass or less.

the most attractive combination of high strength, high resistivity, and low magnetic susceptibility. However, there was significant concern that certain variants may become superconductors under experimental conditions [46].

To ensure our enclosure would not become a superconductor during normal use, I used a SQUID magnetometer (Quantum Design MPMS XL) at Harvard’s Laukien-Purcell Instrumentation Center to test the magnetic behaviors of several titanium alloys. We originally believed that three alloys would be available in the quantities needed to produce an enclosure for CTRAP’s magnet - grades 5, 6, and 9 (Table 3.3). For this test, I was able to procure a cylinder of 3 mm in both height and diameter of grade 6 titanium and two such cylinders from different samples of grade 5 titanium. Because of the time-sensitivity of this project, the only grade 9 titanium I had access to was a large, thin-walled cylinder. This cylinder was cut into nearly flat rectangular pieces, which were then packed securely into the magnetometer’s sample holder so the volume filled was of dimensions similar to the cylinders used for other alloys.

The results of my tests on different alloys are given in Fig. 3.10 and Fig. 3.11. Each was performed by starting at zero field, cooling a sample to the intended temperature, then cycling  $H_z$  through the following path: 0 Oe  $\rightarrow$  15000 Oe  $\rightarrow$  -15000 Oe  $\rightarrow$

0 Oe. Measurements were taken both at 4.2 K, the expected temperature of the enclosure, and 3.5 K, so a superconducting transition just below 4.2 K would not be missed. The grade 5 hysteresis curves show clear signs of superconductivity at both temperatures [48]. Similar results have been reported elsewhere [46]. The data taken for alloys of grades 6 and 9 appears to show simple paramagnetic behavior, though a strange, reproducible feature occurs near zero field for the grade 9 sample at 3.5 K. This feature is shown in more detail in Fig. 3.12.

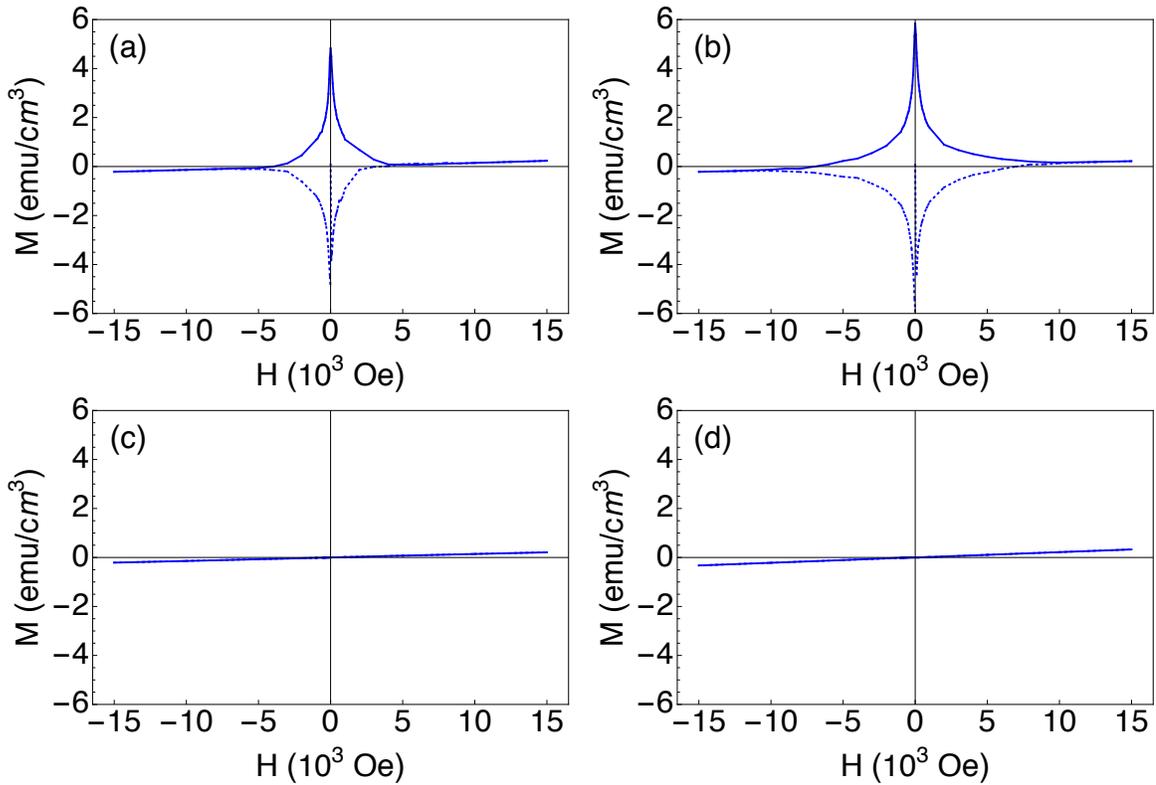


Figure 3.10:  $M$  vs.  $H$  hysteresis curves at 3.5 K for titanium alloys of grades (a) 5 (sample 1), (b) 5 (sample 2), (c) 6, and (d) 9. The curve is dotted when the field is increasing and solid when the field is decreasing.

The results of my tests on different alloys suggested that grade 5 titanium should be avoided, grade 9 was likely okay, and grade 6 posed the least risk. Unfortunately,

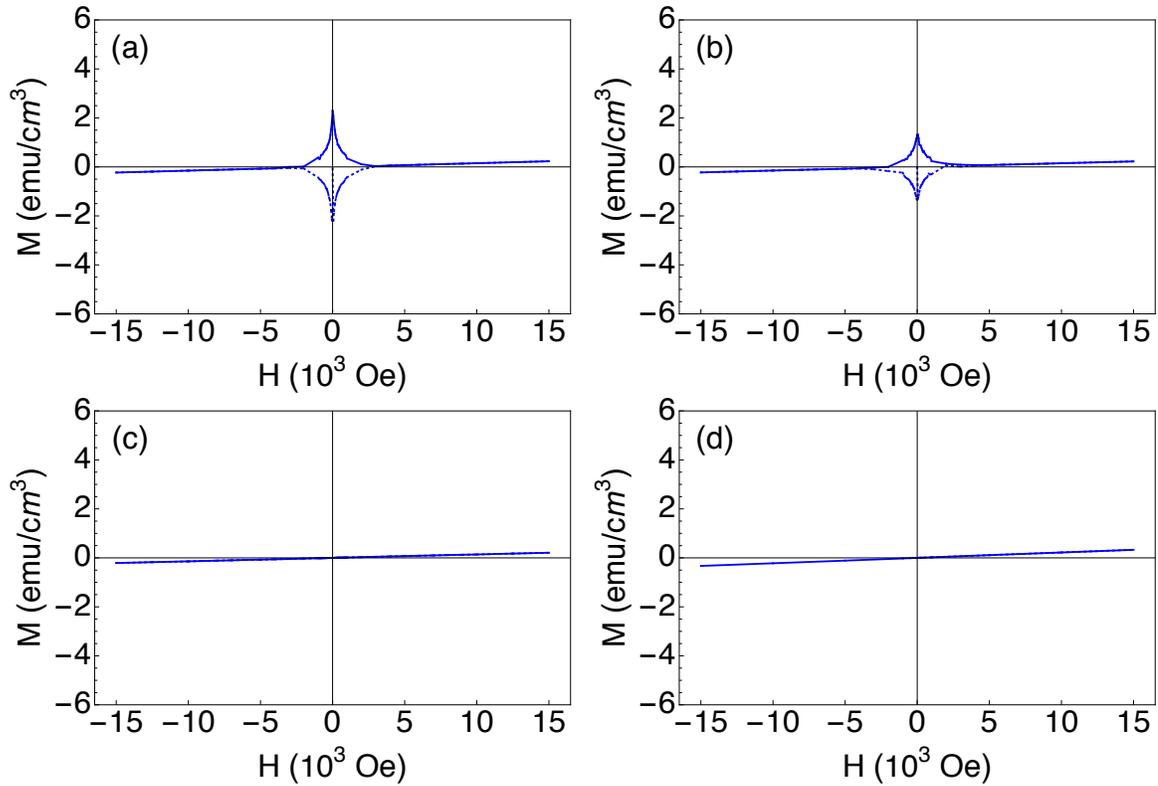


Figure 3.11:  $M$  vs.  $H$  hysteresis curves at 4.2 K for titanium alloys of grades (a) 5 (sample 1), (b) 5 (sample 2), (c) 6, and (d) 9. The curve is dotted when the field is increasing and solid when the field is decreasing.

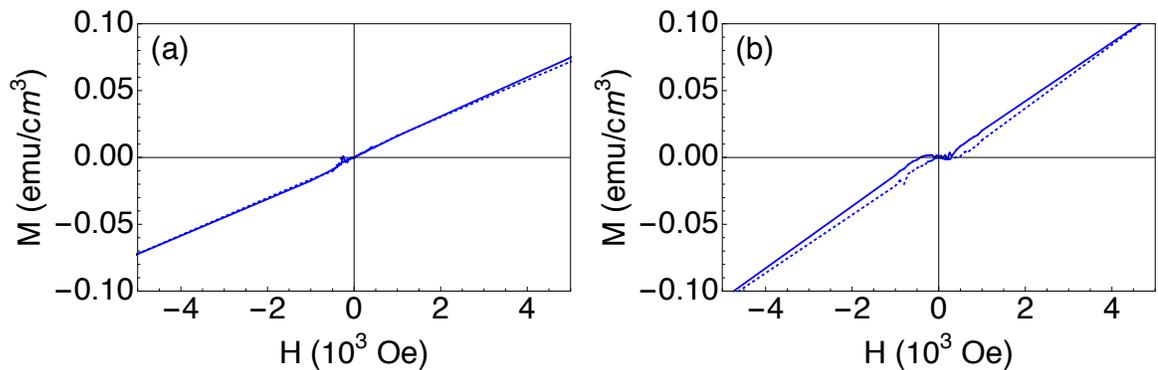


Figure 3.12: Subfigures (a) and (b) are zoomed-in views of parts (c) and (d), respectively, in Fig. 3.10. The structure in the grade 9 sample, part (b) of this figure, persisted through multiple data collection cycles.

after performing these tests, the only supplier we found with enough grade 6 titanium to construct an enclosure revealed that they were unsure whether the stock they had was actually grade 6. They also refused to give us a small sample for testing purposes. We were thus forced to use grade 9.

Altemp Alloys provided us with five pieces of grade 9 titanium, sized as needed for the magnet enclosure. From each of these pieces, I removed two 3 mm diameter, 3 mm tall cylinders. I made sure to take the samples from opposite sides of each piece to strengthen the argument that the results of local measurements generalize over the entire piece of titanium. The measurements performed were similar to those done on the four samples of different titanium alloys. Three modifications were made to the data collection process due to a lack of time, both for magnetometer use and for making the decision to use the Altemp stock for the enclosure. Only half of the  $M$  vs.  $H$  curve was probed, since no more is needed to determine whether a sample behaves as a superconductor.  $H$  was only increased to  $10^4$  Oe, since the two branches of the hysteresis curves I'd seen to that point merged below this field. Finally, time was only available for measurements at one temperature, so the expected operating temperature, 4.2 K, was chosen.

The results, shown in Fig. 3.13 and Fig. 3.14 give strong evidence of superconducting behavior in the material intended to form the inner bore of the enclosure. The piece to be used for the sideport rings (used to form an indium seal with the sideport flanges) also shows hysteresis akin to that shown in Fig. 3.12, but the curves of increasing and decreasing field merge around 2000 Oe. For both pieces, the minimum field experienced during CTRAP magnet manipulations lies comfortably above those

for which the metal behaves as a superconductor. The rest of the titanium stock appears to behave paramagnetically. It should also be pointed out, as is particularly visible in Fig. 3.14, that samples from different locations in a single piece of titanium give similar curves, lending credibility to the hypothesis that a hysteresis curve from a small sample reflects the general properties of the piece from which it was taken.

With confidence that the material we bought would not behave as a superconductor under operating conditions, we constructed a magnet enclosure. Harvard's shop machined the titanium stock, and Joining Technologies (electron-beam) welded the pieces around the magnet. A great deal of planning went into the design of the enclosure and the execution of the welds as we aimed for several conflicting goals:

- Maximize the radius of the inner bore to preserve as much trap depth as possible.
- Fix a precise alignment between the magnet and the enclosure to maximize the available sideport volume and increase the accuracy of our field calculations.
- Prevent stresses resulting from differential contraction between the magnet's G10 body and the titanium enclosure (Table 3.4).
- Avoid heat damage to the magnet during welding.

Welding was also a challenge due to difficult-to-access joints deep within the sideports. Fig. 3.15 shows the weld joints along with some of the clamps used for alignment and heat sinking.

Despite the difficulties along the way, we were ultimately successful in constructing a robust titanium magnet enclosure compatible with sideports. Since its assembly, the magnet-enclosure ensemble has been successfully cycled between room temperature

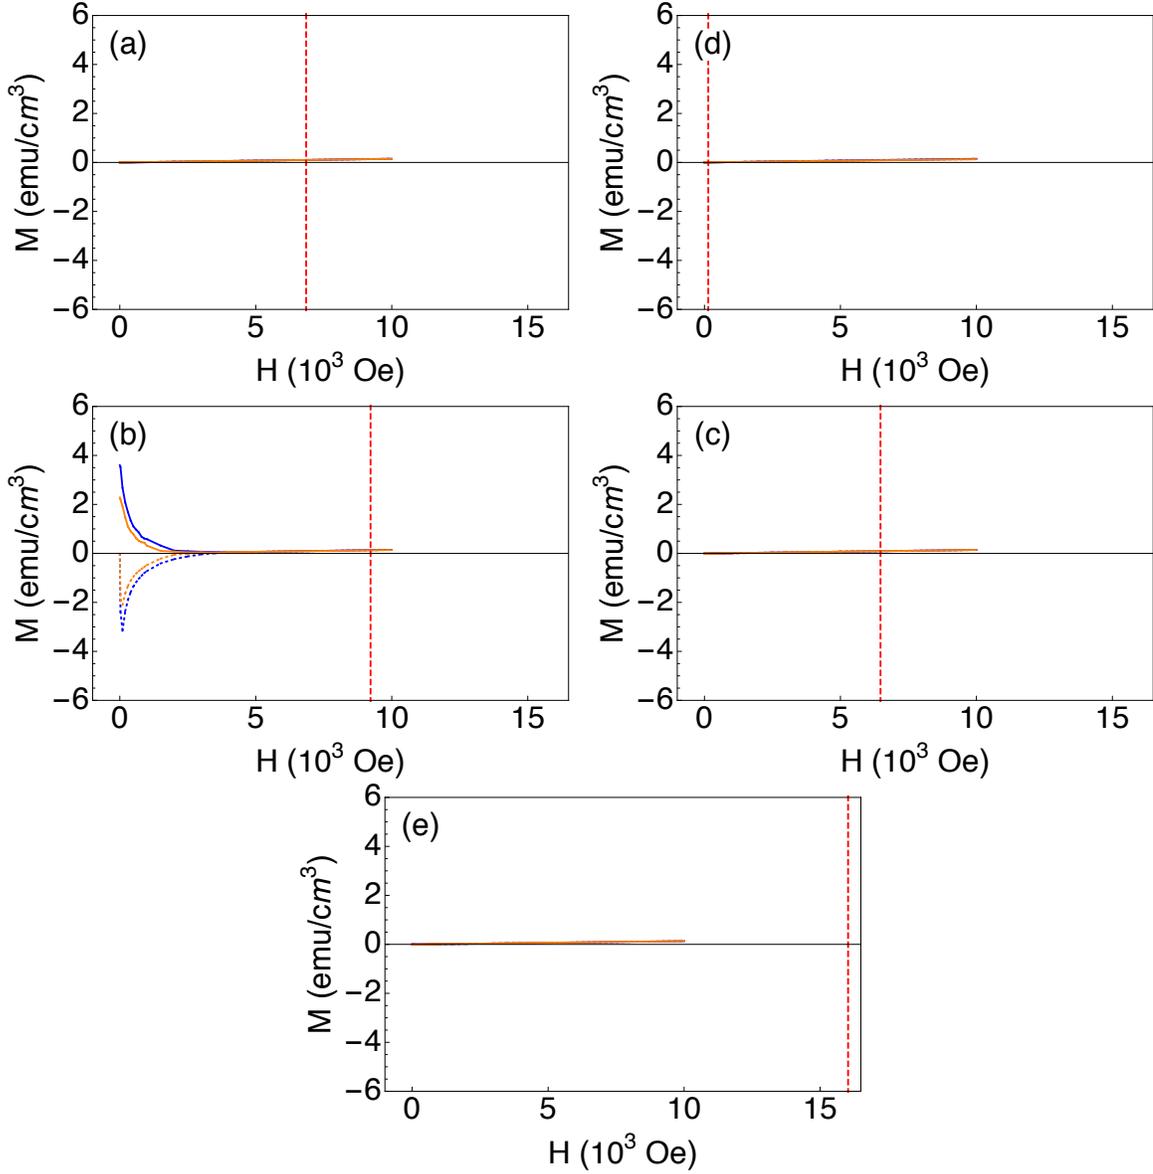


Figure 3.13:  $M$  vs.  $H$  hysteresis curves at 4.2 K for the grade 9 titanium stock used to produce the (a) top and bottom plates, (b) outer shell, (c) inner bore, (d) sideports, and (e) sideport rings. Two samples are taken from each piece of material with one represented in blue and another in orange. The convention of using a dotted curve for increasing field and a solid curve for decreasing field from Figs. 3.10-3.12 continues here. A red dashed line gives the minimum field experienced by each piece during normal trap manipulations.

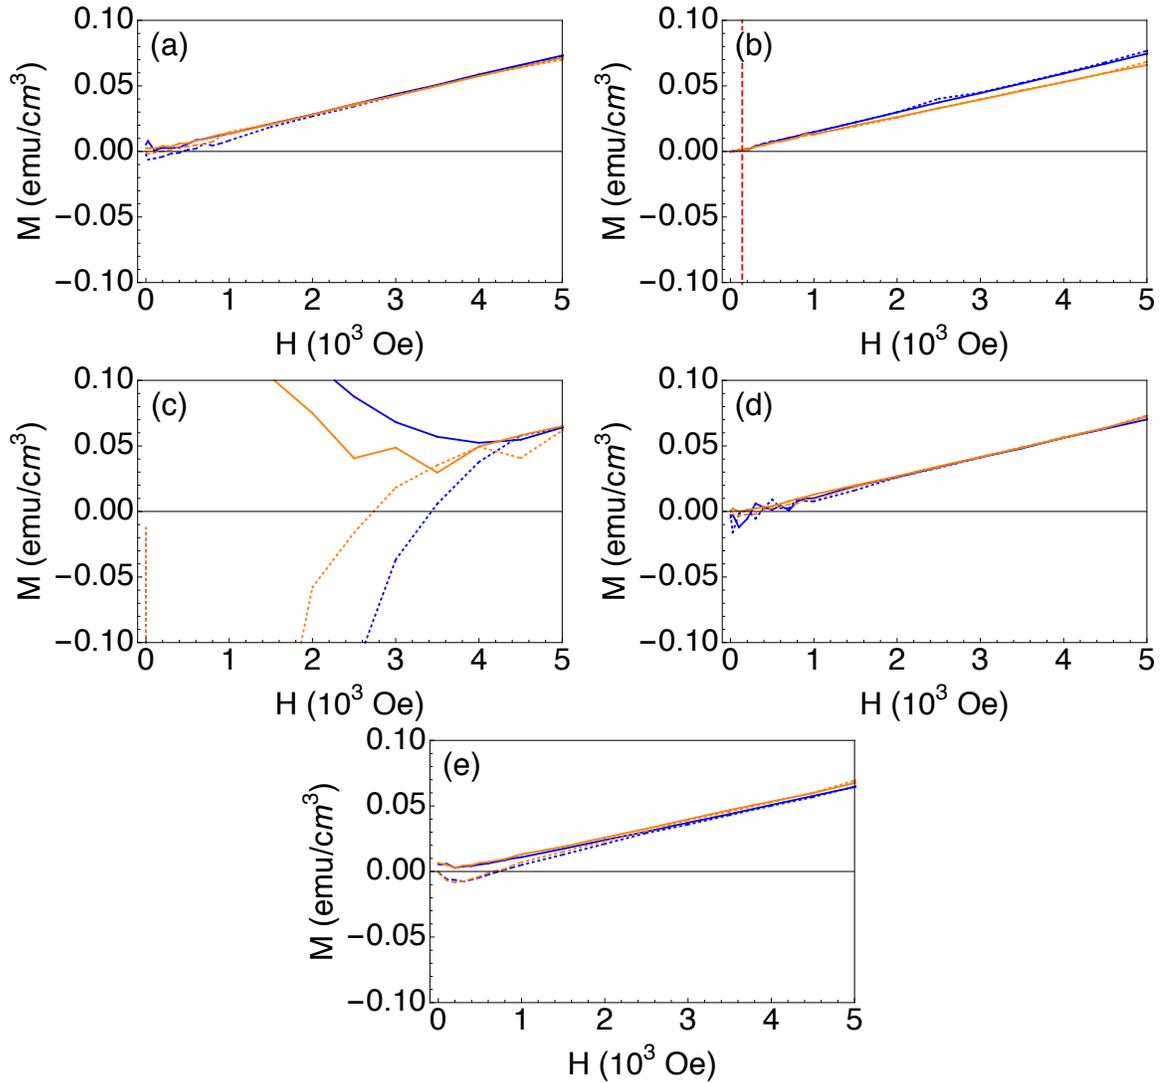


Figure 3.14: A zoomed-in version of Fig. 3.13 shows  $M$  vs.  $H$  curves for material used to form the (a) top and bottom plates, (b) outer shell, (c) inner bore, (d) sideports, and (e) sideport rings. Note that the red dashed line is cut off in some plots due to the reduction in the considered range of  $H$ .

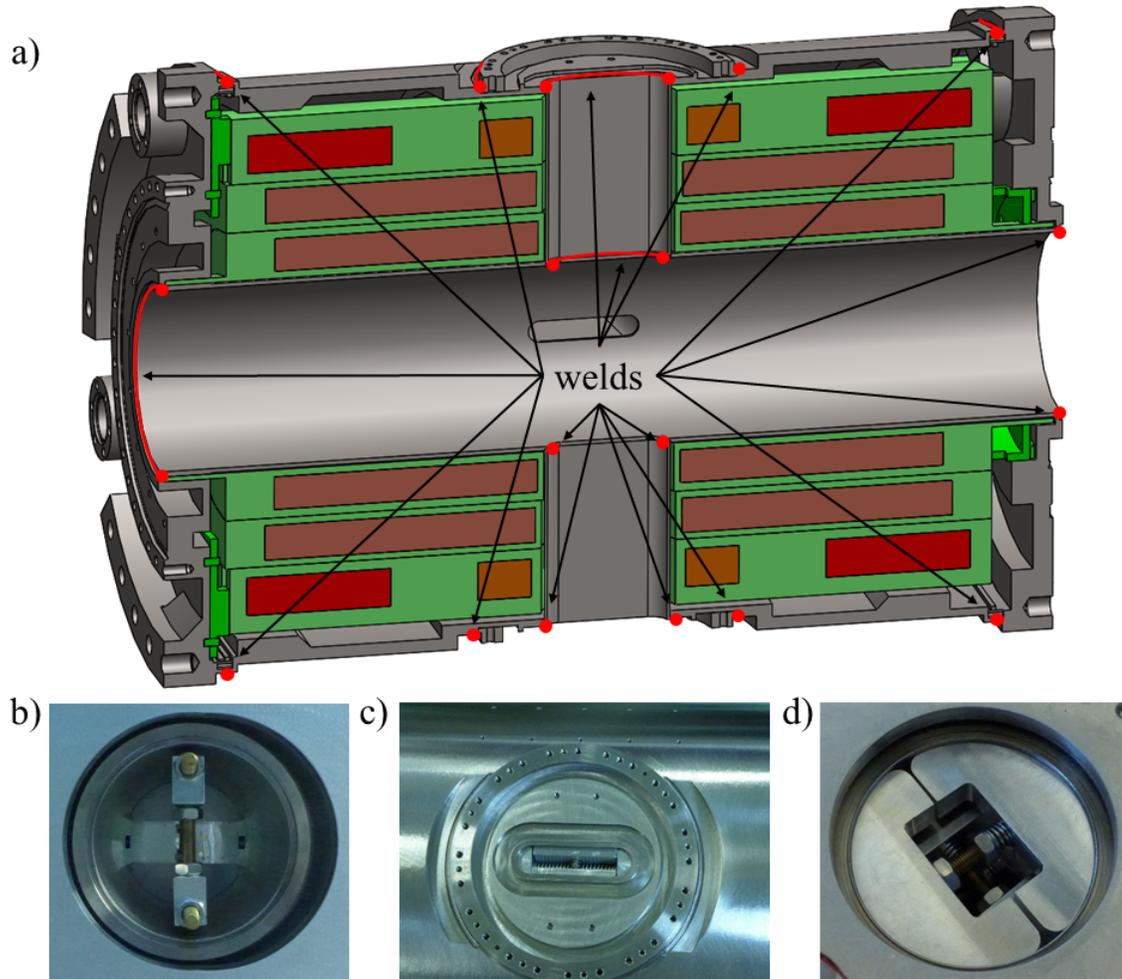


Figure 3.15: (a) A diagram of the magnet and enclosure is shown with weld paths colored red (intersections with the cross-sectional surface shown as dots), highlighting the difficulty of interior welds and the close

Material	$(L(293K) - L(4K))/L(293K) (10^{-3})$
Ti	1.51
Gr 5 Ti	1.73
G10    to fibers	2.41
G10 $\perp$ to fibers	7.06

Table 3.4: Thermal contraction coefficients for titanium and G10 [49]. Limited information is available for grade 9 titanium, so its contraction coefficient is assumed to be somewhere between that of pure titanium and the grade 5 alloy.

and 4 K eight times, spending a total of 67 weeks cold, without leaking. Unfortunately, the extra material needed for the inner bore decreases the depths of the octupole and quadrupole traps by %20 and %10, respectively, compared with the mixed-material design. However, as will be shown in the next section, our effort to make frequent, well-timed, fast current ramps possible was a resounding success.

## 3.6 Magnet Operation

The limiting factor for the BTRAP Ioffe magnet's ramp rate is the set of diodes set up in parallel with each coil to protect against quenches. These redirect current in case a quench raises the voltage drop across a coil above 5 V. Due to the high inductances of the coils, changing currents at even modest rates can cause such voltage drops. For instance, the BTRAP quadrupole coil is limited to current ramps of 1.5 A/s by its diode.

To avoid these limitations, we implemented a better quench protection system

for CTRAP. Quenches are detected by measuring voltage drops over each half of a given coil using voltage taps built into the magnet. The inductive voltage over each half of a coil should be nearly the same during magnet charging and discharging. However, a quench is unlikely to propagate symmetrically around a coil's center. Consequently, quenches are expected to cause the voltage drops over symmetric coils halves to diverge.

Once a quench is detected, the current must be diverted from the coil as quickly as possible. In place of the passive BTRAP system, we designed an active protection circuit, shown in Fig. 3.16. A separate system uses the method described above to monitor for quenches. When it detects an asymmetric voltage drop across the halves of a given coil, exceeding some user-defined threshold, a signal is sent through an optical fiber to the control circuit of the quench protection system. The IGBT is then opened, the coil's power supply is shut off, and the event is recorded in a high-frequency data stream recording fast current dumps. We also allow current to be dumped on command, as is desirable when trying to detect trapped antihydrogen atoms, by sending a TTL to the control circuit. This initiates the same sequence of events as a detected quench.

The results of these changes are shown in Fig. 3.17. The more nuanced quench-detection system allows for a full trap ramp-up in about two minutes. Current dumps are completed in tens of milliseconds, with radial confinement disappearing after 35 ms. In practice, we are able to ramp up currents and perform fast dumps as frequently as desired. No recovery time is needed in between.



## 3.7 Summary

In this chapter, we discussed the neutral atom trap that the ATRAP collaboration uses to capture antihydrogen. First, the general theory behind Ioffe traps and antihydrogen trapping was treated. Then, we looked at the predecessor to CTRAP's Ioffe magnet. Several problems with this magnet which interfered with ATRAP's experimental goals were identified. This motivated the construction of a replacement, which formed the subject of the remainder of the chapter.

After describing the new magnet and its coils, we outlined the problems we had surrounding it with an enclosure for holding liquid helium. The first attempt, a mixed-material enclosure, was unable to survive thermal cycling and had to be discarded. An all-metal enclosure was built in its place and shown to be robust under cooling to 4.2 K. A new quench protection system was added alongside the new magnet, allowing for faster current changes. Overall, the process of upgrading the BTRAP Ioffe magnet was long and rife with difficulties. However, the result of this project is a far better trap for capturing and detecting antihydrogen atoms.

# Chapter 4

## Penning Traps in Theory

In any antihydrogen experiment, antiprotons and positrons must be brought together to form the atoms. A typical choice for the tool to accomplish this task is known as a Penning trap. This chapter starts with a description of an ideal Penning trap, along with the motions of a single particle in the trap. While this gives some intuition for the types of motion relevant to the clouds of particles used to form antihydrogen, the description is incomplete. The rest of the chapter will complement earlier sections with the effects of interactions between trapped particles. In particular, equilibrium plasma configurations as well as collisional effects are discussed.

### 4.1 Fields and Single-Particle Motion

One of the simplest schemes for confining charged particles is known as a Penning trap. The development will closely follow [50]. The setup consists of a uniform, static

magnetic field,  $B_0\hat{z}$ , superposed on a quadratic electric potential,

$$\varphi(\rho, z) = \frac{V_0}{2d^2}(z^2 - \frac{1}{2}\rho^2) \quad (4.1)$$

$$d^2 = \frac{1}{2}(z_0^2 + \frac{1}{2}\rho_0^2), \quad (4.2)$$

where  $\rho_0$  and  $z_0$  are defined in Fig. 4.1 and  $V_0$  is the difference between the endcap and ring potentials. One way to establish such a potential, as suggested by the figure, is to position hyperbolic electrodes along the equipotentials specified by Eqn. 4.1.

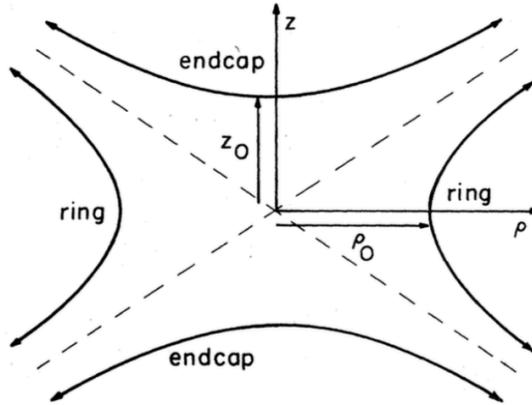


Figure 4.1: Cross-section of hyperbolic electrodes used for a Penning trap's quadratic potential (from [50])

Although the CTRAP experiment generally manipulates clouds of  $10^4$  particles or more, it's instructive to understand single-particle motion in Penning fields. We make the definitions

$$\omega_z^2 = \frac{qV_0}{md^2} \quad (4.3)$$

$$\omega_c = \frac{qB_0}{m}, \quad (4.4)$$

where  $q$  and  $m$  represent the charge and mass of the confined particle, respectively.  $V_0$  is taken to have the same sign as  $q$ . The Lorentz force law then yields equations

of motion

$$\ddot{x} = \frac{\omega_z^2}{2}x + \omega_c \dot{y} \quad (4.5)$$

$$\ddot{y} = \frac{\omega_z^2}{2}y - \omega_c \dot{x} \quad (4.6)$$

$$\ddot{z} = -\omega_z^2 z. \quad (4.7)$$

The axial coordinate decouples and undergoes simple harmonic motion with frequency  $\omega_z$ . The equations of motion for the transverse degrees of freedom can be combined as

$$\ddot{u} + i\omega_c \dot{u} - \frac{\omega_z^2}{2}u = 0, \quad (4.8)$$

where  $u = x + iy$ . The solution is a linear combination of oscillations at frequencies

$$\omega_{\pm} = -\frac{\omega_c}{2} \mp \frac{1}{2}\sqrt{\omega_c^2 - 2\omega_z^2}. \quad (4.9)$$

We choose positive frequencies for counter-clockwise rotation in the x-y plane and negative frequencies for clockwise rotation, where  $\hat{z}$  is parallel to  $\vec{B}$ . Returning to Cartesian coordinates, we can decompose a particle's position as follows.

$$\vec{r}(t) = \vec{r}_+(t) + \vec{r}_-(t) + z(t)\hat{z} \quad (4.10)$$

As was previously explained, the last term in Eqn. 4.10 represents oscillation parallel to the z-axis with frequency  $\omega_z$ . The terms  $\vec{r}_{\pm}(t)$  rotate in the x-y plane with frequencies  $\omega_{\pm}$ . The sum of these motions is illustrated (not to scale) in Fig. 4.2.

The relative magnitudes of the Penning trap frequencies are worth emphasizing. First, let's make the more conventional definitions

$$\omega_m \equiv \omega_- \quad (4.11)$$

$$\omega_c' \equiv \omega_+ = -\omega_c - \omega_m. \quad (4.12)$$

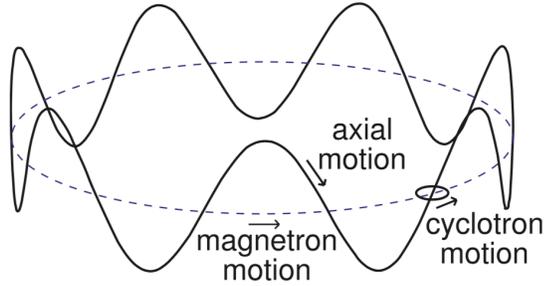


Figure 4.2: Cyclotron, axial, and magnetron motion in a Penning trap (from [41])

References are often made to  $\omega_m$  as the magnetron frequency and  $\omega_c'$  as the modified cyclotron frequency. In many situations, including the one found in our trap, the expressions for these quantities may be simplified and a hierarchy emerges. Consider  $|\omega_z/\omega_c|$ . From the definitions in Eqns. 4.3-4.4,

$$\omega_z^2/\omega_c^2 = \frac{mV_0}{qd^2B_0^2}. \quad (4.13)$$

In CTRAP, potentials change on the 100 V scale over the length of about a centimeter. Taking  $B_0$  to be 1 T,  $\omega_z/\omega_c$  is around  $10^{-3}$  for electrons (and positrons) and  $10^{-1}$  for antiprotons. Consequently  $\omega_m$  may be approximated to lowest order in  $\omega_z^2/\omega_c^2$  as

$$\omega_m \approx -\frac{1}{2} \frac{\omega_z^2}{\omega_c^2} \omega_c. \quad (4.14)$$

Since  $|\omega_z/\omega_c| < 1$ , we arrive at the hierarchy

$$|\omega_c'| > |\omega_z| > |\omega_m|, \quad (4.15)$$

where  $|\omega_z/\omega_c|$  is (approximately) the factor by which the frequencies change going left to right in Eqn. 4.15. Note that this factor is proportional to the square root of the particle's mass, so the hierarchy is much stronger for electrons and positrons than it is for protons and antiprotons.

## 4.2 Plasmas in Equilibrium

The discussion in the previous section focuses on single particles in an ideal quadrupole Penning trap. In CTRAP, typical manipulations involve clouds of  $10^4$ - $10^8$  particles. Both the Coulomb interaction and collisions between trapped charges modify the single-particle picture. This section describes equilibrium configurations for clouds of trapped, charged particles.

We assume that our charges are confined in a Penning trap with a uniform magnetic field  $B_0 \hat{z}$  and an electric potential  $\varphi$  which is cylindrically symmetric, but not necessarily pure quadrupole. The Boltzmann distribution for the trapped particles in cylindrical coordinates is

$$f(\vec{r}, \vec{v}) \propto \exp \left( -\frac{1}{k_B T} \left[ \frac{1}{2} m (\vec{v} - \omega_r \rho \hat{\phi})^2 - \frac{1}{2} m \omega_r \rho^2 (\omega_c + \omega_r) + q\varphi(\rho, z) \right] \right), \quad (4.16)$$

where  $\omega_c = qB/m$  and  $\omega_r$  is a parameter to be determined [51, 52]. Inter-particle interactions are ignored in this treatment ( $\Gamma = \left( \frac{q^2}{4\pi\epsilon_0 a} \right) / (k_B T) < 0.2$  for plasmas of interest). In coordinates rotating with angular frequency  $\omega_r$  about the  $z$  axis, the velocity-dependent part of Eqn. 4.16 is proportional to a Maxwell-Boltzmann distribution. We can thus interpret Eqn. 4.16 as describing some spatial distribution of charges with temperature  $T$ , rotating around the  $z$  axis with angular frequency  $\omega_r$ .

To determine this spatial distribution, we integrate over the velocities and multiply by the total number of trapped particles, arriving at an expression for particle density.

$$n(\rho, z) = n_0 \exp \left( -\frac{1}{k_B T} \left[ -\frac{1}{2} m \omega_r \rho^2 (\omega_c + \omega_r) + q\varphi(\rho, z) \right] \right) \quad (4.17)$$

The potential  $\varphi(\rho, z)$  has been defined to vanish at the origin. Eqn. 4.17 still has two undetermined quantities:  $\varphi(\rho, z)$ , the potential resulting from both trapped charges

and electrode biases, and  $\omega_r$ . The potential may be found by solving Poisson's equation with charge density  $n(\rho, z)$ . Integrating  $n(\rho, z)$  over all space then gives a relation of the form

$$N = n_0 g(\omega_r), \quad (4.18)$$

where  $N$  is the total number of particles in the cloud. If  $N$  is known, Eqn. 4.18 can, at least locally  $n_0$  and  $\omega_r$ , be used to find the rotation frequency from the central density and vice versa.

The Poisson equation referenced above is often too difficult to solve analytically. Numerical solutions suggest that as long as the Debye length ( $\lambda_D^2 = \frac{k_B T}{q^2 n_0 / \epsilon_0}$ ) is small compared to the dimensions of the cloud,  $\varphi(\rho, z)$  is such that  $n(\rho, z)$  is approximately constant over some volume before falling to zero along the boundary over a distance of about a Debye length [51]. Typical CTRAP plasmas are more than 80 times  $\lambda_D$  wide along their narrowest dimension, so this picture should accurately describe the clouds we manipulate.

If the density is nearly constant in the plasma's interior, the potential is approximately described by

$$\varphi(\rho, z) = \frac{m}{2q} \omega_r (\omega_c + \omega_r) \rho^2. \quad (4.19)$$

Poisson's equation then yields

$$n_0 = -\frac{2m\epsilon_0}{q^2} \omega_r (\omega_c + \omega_r) \quad (4.20)$$

For plasmas in CTRAP  $|\omega_r/\omega_c| \ll 1$ , so this may be inverted to yield

$$\omega_r = -\frac{n_0 q^2}{2m\epsilon_0 \omega_c} = -\frac{n_0 q}{2\epsilon_0 B}. \quad (4.21)$$

As the density depends on the plasma's rotational frequency, torques may be used to change its radius. More will be said about this later.

The radius may also change as a result of changes in the magnetic field. The Hamiltonian for a collection of particles in a Penning trap may be written [53]

$$H = \sum_i \frac{1}{2m} \left[ p_{\rho,i}^2 + \frac{1}{\rho_i^2} \left( p_{\phi,i} - \frac{1}{2}qB\rho_i^2 \right)^2 + p_{z,i}^2 \right] + \sum_i q\varphi(\vec{r}). \quad (4.22)$$

Cylindrical symmetry ensures that this expression is invariant under translations in  $\phi$ , so

$$p_\phi = \sum_i \left[ m\rho_i^2 \dot{\phi}_i + \frac{1}{2}qB\rho_i^2 \right] \quad (4.23)$$

is a constant [54]. The first term in Eqn. 4.24 is negligible for our plasmas. This can be seen by noting that, for a fixed radius, the Boltzmann distribution gives  $\langle \dot{\phi} \rangle = \omega_r$ . Thus, the ratio of the first term to the second is  $2\omega_r/\omega_c \ll 1$  for clouds in CTRAP. A more rigorous justification, as well as a bound on the size of the first term, is given in [53]. Dropping the kinetic term, we find that

$$p_\phi \approx \frac{1}{2}qB \sum_i \rho_i^2 = \frac{1}{2}qBN \langle \rho^2 \rangle \quad (4.24)$$

does not change as long as cylindrical symmetry is maintained. In particular,  $\langle \rho^2 \rangle$  changes as the inverse of the magnetic field.

Many characteristics of a confined plasma's equilibrium configuration have been explored, but we haven't yet given information on such a cloud's shape. For simplicity, let's assume that the electric potential is an ideal quadrupole. In the zero-temperature limit, the plasma may be shown to be a uniform-density spheroid [55,56]. Taking the

spheroid to have a radius of  $r_p$  and a half-length of  $z_p$ , we make the definitions

$$\alpha_a \equiv z_p/r_p \tag{4.25}$$

$$\omega_p^2 \equiv \frac{n_0 q^2}{\epsilon_0 m}, \tag{4.26}$$

which are referred to as the aspect ratio and plasma frequency, respectively. With the definition of  $\omega_z$  in terms of the potential from Eqn. 4.3, it is possible to show that

$$\frac{\omega_z^2}{\omega_p^2} = Q_1^0 \left( \frac{\alpha_a}{(\alpha_a^2 - 1)^{1/2}} \right) / (\alpha_a^2 - 1), \tag{4.27}$$

where  $Q_l^m$  is the associated Legendre function of the second kind [57]. The right-hand side of this expression monotonically decreases with  $\alpha_a$ . Consequently, the equation may be inverted to solve for the aspect ratio. Also, for a fixed  $q$ ,  $m$ , and  $d$ , the left-hand side is proportional to  $V_0/n_0$ .

Taken together, Eqn. 4.27 and the conservation of 4.24 paint a clear picture of how changes in trap fields alter the shape of a confined plasma. A simple calculation reveals that  $\langle \rho^2 \rangle = 2\rho_0^2/5$  for a spheroid. Thus, when the magnetic field goes from  $B_0$  to  $B_1$ ,  $\rho_0$  is multiplied by a factor of  $(B_0/B_1)^{1/2}$ . On the other hand, changes in the magnitude of the quadrupole potential, parameterized by  $V_0$ , don't change  $\rho_0$ . Consequently,  $z_p \propto \alpha_a$ , where the relationship between  $\alpha_a$  and  $V_0$  is described by Eqn. 4.27.

In this section, equations have been given relating a number of different parameters used to describe plasmas and the Penning traps that hold them. All formulas are true, with stated assumptions, if the potential takes the form of a pure quadrupole. Some hold with the weaker assumption of cylindrically symmetric fields. Though an extensive study has not been done, we expect the intuitions gained from the ideal

case to hold when the potential is sufficiently close to Eqn. 4.1. In cases requiring more exact solutions, the numerical solver EQUILSOR may be employed [58]. This program has previously been shown to give results in agreement with experiment [59].

### 4.3 Collisions and Cooling

The previous section describes trapped plasmas in thermal equilibrium. However, clouds of charged particles held in CTRAP's Penning trap are often disturbed by changing electric and magnetic fields. It is therefore desirable to understand the time scale on which equilibrium is restored. At the same time, our experimental goals require more of our trapped plasmas. Because the antihydrogen trap is so shallow ( $< 1$  K), positrons and antiprotons involved in anti-atom synthesis must be moving as slowly as possible. An efficient cooling mechanism must therefore be realized. In this section, we examine the time scales for both cooling and thermalizing trapped plasmas.

Consider the motion for a particle in a trapped plasma between collisions. Take the magnetic field to be the background  $B\hat{z}$  and the electric field to be derived from Eqn. 4.19. The resulting equations of motion for transverse degrees of freedom are identical to Eqns. 4.5-4.6, with  $\omega_z$  replaced by  $\omega_p$ . Since the  $z$  degree of freedom, which undergoes force-free motion, is decoupled in this situation as well, the solution for the transverse components is again the sum of rotating vectors with frequencies given by Eqn. 4.9 after substituting  $\omega_p$  for  $\omega_z$ .

Accelerating charges radiate at a rate given by Larmor's formula [60]

$$P_{rad} = \frac{q^2 a^2}{6\pi\epsilon_0 c^3}, \quad (4.28)$$

where  $a$  is the magnitude of the acceleration. Let's consider only the high-frequency component of a plasma particle's motion. This component,  $\vec{r}_+$ , obeys

$$\ddot{\vec{r}}_+ = (\omega_+ \hat{z}) \times \dot{\vec{r}}_+. \quad (4.29)$$

Since  $\vec{r}$  is perpendicular to  $\hat{z}$ ,

$$a_+^2 = (\ddot{\vec{r}}_+)^2 = \omega_+^2 v_+^2 = 2E_+ \omega_+^2 / m. \quad (4.30)$$

Thus, the energy in this motion decays as

$$\dot{E}_+ = -\frac{q^2 \omega_+^2}{3\pi\epsilon_0 m c^3} E_+ \approx -\frac{q^4 B^2}{3\pi\epsilon_0 m^3 c^3} E_+, \quad (4.31)$$

where the last (approximate) equality comes from  $\omega_+ \approx \omega_c$ . Eqn. 4.31 predicts, for a 1 T field, an energy decay time of 2.6 s for electrons/positrons and 1.6e10 s (506 years!) for antiprotons. The stark difference comes from the  $m^3$  dependence of this time.

We haven't yet treated the component rotating at  $\omega_-$ . However  $|\omega_-/\omega_+|$  never exceeds 1.3e-4 for electrons/positrons or 1.0e-2 for antiprotons in CTRAP. From the quadratic dependence on frequency in Eqn. 4.31, we see that no species is significantly cooled over the duration of our experiments by rotation at  $\omega_-$ .

So far, it appears that radiation only effectively cools leptons' cyclotron-like motion. It's important to ensure that other degrees of freedom are cooled for electrons and positrons. Antiprotons also require a cooling mechanism. To solve these problems, we rely on lepton-lepton scattering and electron-antiproton scattering, respectively.

A summary of results on scattering and equilibration rates found in the literature is presented in [61]. Before writing the relevant formulas, it's important to explain the oft-used parameter  $\Lambda_{st}$ . This is known as the Coulomb logarithm for particles  $s$  and  $t$ , and it is defined by

$$\log \Lambda_{st} = \log (r_{max}^{(s,t)} / r_{min}^{(s,t)}). \quad (4.32)$$

Here,  $r_{max}^{(s,t)}$  is taken to be the Debye length of the plasma of interest, and  $r_{min}^{(s,t)}$  is an approximation of the distance of closest approach between  $s$  and  $t$  particles. The value is often between 10 and 20 and only changes logarithmically with parameters like density and temperature [62]. Consequently, we will take  $\log \Lambda_{i,j} = 10$  in what follows. Perfect precision is unimportant, since the order of magnitude of scattering and relaxation times is what we aim to estimate.

The first quantity of interest is the collision rate between particles of type  $s$  moving through a field of  $t$  particles. This rate,

$$\nu_{st} = \frac{2^{1/2} n_t e^4 \log \Lambda_{st}}{12\pi^{3/2} \epsilon_0^2 \mu_{st} m_s [(k_B T_s / m_s) + (k_B T_t / m_t)]^{3/2}}, \quad (4.33)$$

is most easily interpreted by noting that it satisfies

$$\dot{\vec{V}}_s = - \sum_t \nu_{st} (\vec{V}_s - \vec{V}_t), \quad (4.34)$$

where  $\vec{V}_k$  is the average velocity of the  $k$  particles and  $\mu_{st}$  is the reduced mass of  $m_s$  and  $m_t$ . It's important to note that, in general,  $\nu_{st} \neq \nu_{ts}$ . This is easily seen if  $s$  is far more abundant than  $t$ , but  $m_s = m_t$ . In this case, the average  $t$  will have far more encounters with  $s$  than the average  $s$  particle will have with  $t$ , so  $\nu_{st} \ll \nu_{ts}$ . It should also be emphasized that  $s$  and  $t$  may represent the same type of particle.

Cloud	$(\nu_{e^\pm e^\pm}^E)^{-1}$ ( $\mu\text{s}$ )	$(\nu_{\bar{p}\bar{p}}^E)^{-1}$ ( $\mu\text{s}$ )	$(\nu_{\bar{p}e^-}^E)^{-1}$ ( $\mu\text{s}$ )	$(\nu_{e^-\bar{p}}^E)^{-1}$ ( $\mu\text{s}$ )
pure $e^-$	0.05	-	-	-
pure $e^+$	0.02	-	-	-
$e^-$ s with trapped $\bar{p}$ s	0.004	20	3	300
$\bar{p}$ s with residual $e^-$ s	400	20	$3 \cdot 10^5$	300

Table 4.1: Thermalization time scales are given for plasmas commonly used in CTRAP with  $T = 20$  K.

The rate of energy exchange can also be described in similar terms. The thermalization rate between species  $s$  and  $t$  is given by

$$\nu_{st}^E = 2 \left( \frac{m_s}{m_s + m_t} \right) \nu_{st}. \quad (4.35)$$

This quantity satisfies

$$\dot{T}_s = \sum_t \nu_{st}^E (T_t - T_s), \quad (4.36)$$

making  $\nu_{st}$  the rate of equilibration between species  $s$  and  $t$ . The interpretation is slightly different when  $s = t$ . In that case,  $\nu_{ss}$  represents the rate at which the energy distribution of  $s$  approaches Maxwell-Boltzmann.

Table 4.1 uses Eqns. 4.33 and 4.35 to describe equilibration time scales in CTRAP at  $T = 20$  K [45]. Recall that the lepton cyclotron degree of freedom is cooled on a 2.6 s scale in the 1 T field permeating CTRAP. For catching and preparing protons, a 3.7 T field is used, giving a 0.2 s cooling time. In pure lepton plasmas, thermalization thus happens much more quickly than cooling in either field, suggesting that the clouds stay in thermal equilibrium while their temperatures slowly drop. It's worth noting that simulations suggest an equilibration rate between orthogonal degrees of

freedom three orders of magnitude lower than what we have predicted because of the magnetic field [63]. This rate is still much faster than that of cyclotron cooling, so our qualitative understanding of lepton plasmas is unchanged.

The situation for mixed clouds is slightly more complicated. To describe their evolution, we use the equations [64]

$$\dot{T}_{\bar{p}} = \nu_{\bar{p}e^-}^E (T_{e^-} - T_{\bar{p}}) \quad (4.37)$$

$$\dot{T}_{e^-} = \nu_{e^-\bar{p}}^E (T_{\bar{p}} - T_{e^-}) + \nu_c (T_{eq} - T_{e^-}), \quad (4.38)$$

where  $\nu_c$  is the cyclotron energy decay rate in Eqn. 4.31 and  $T_{eq}$  is the equilibrium temperature. These equations must be numerically integrated due to the complicated temperature dependence of the  $\nu_{st}^E$ .

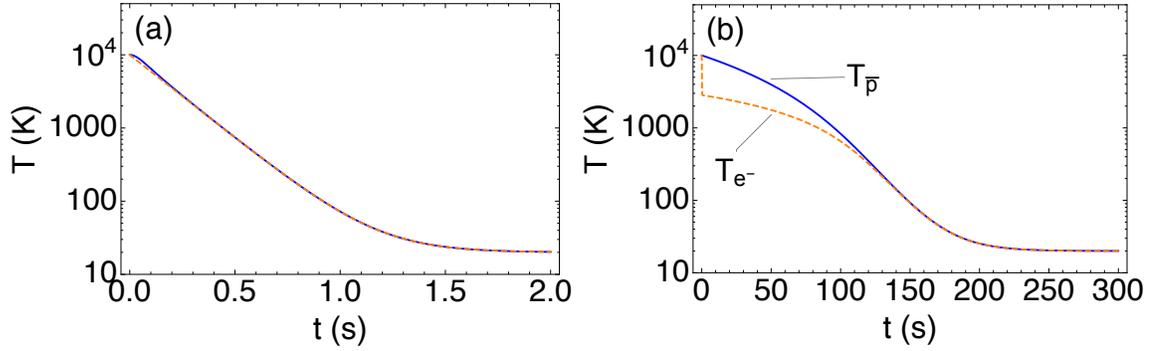


Figure 4.3: Temperatures of the components of a plasma containing electrons (orange, dashed) and antiprotons (blue, solid) are shown as the electrons cool by cyclotron radiation. In (a), the ratio of electrons to antiprotons is  $10^2$  with the geometry taken after antiprotons have been caught and the cloud has been compressed. Part (b) shows cooling after the electrons have been pulsed out, where the electron to antiproton ratio is  $10^{-3}$ .

Fig. 4.3 shows two example plasmas in CTRAP, each heated to  $k_B T \approx 1$  eV. When electrons heavily outnumber antiprotons, as is the case in (a), the electrons cool at a rate similar to what would be observed in the absence of antiprotons. Because

of the tight collisional coupling between antiprotons and electrons (Table 4.1), the antiproton temperature stays close to the electron temperature. On the other hand, (b) describes an instance where  $10^3$  electrons are left to cool  $10^6$  antiprotons. The poor coupling between the two species, exacerbated by high temperatures, allows the electrons to radiatively cool as if no antiprotons were present for a short time. Eventually, the radiative cooling is balanced by heating from scattering antiprotons, and the electron temperature slows its descent. As both species cool, their coupling improves and their temperatures converge.

The model presented in this section for plasma cooling and thermalization is simple, but useful. With it, we are able to determine the times needed for various plasmas to relax to low-temperature equilibrium states. Still, additional insights may be gained from further analysis. Simulations, for instance, should be able to account for effects related to the magnetic field and centrifugal separation in multi-species plasmas [65, 66].

## 4.4 Summary

In this chapter, we discussed certain aspects of charges confined in a Penning trap. We started by summarizing the fields and single-particle motion expected in an ideal quadrupole trap. Equilibrium configurations of low-temperature plasmas suspended in such a trap were then described, and equations connecting properties related to geometry, rotation, and density were given. Finally, radiation caused by cyclotron motion was offered as an effective mechanism for cooling electrons and positrons. Thermalization rates were estimated and used to understand how the cooling of light

lepton cyclotron motion might propagate to other degrees of freedom and heavier particles.

With the theoretical groundwork laid by this chapter, we're ready to move on to Ch. 5. There, the implementation of a Penning trap as part of the CTRAP apparatus is outlined. Although its fields deviate from the ideal quadrupole configuration used for much of the analysis in this chapter, efforts are made to keep the differences small. The ideas discussed here still hold as good approximations.

# Chapter 5

## CTRAP's Penning Trap

The practical details of this chapter, which describe CTRAP's Penning trap, supplement the more theoretical discussion of the previous chapter. The trap consists of the large persistent solenoid surrounding the experiment and the stack of ring electrodes lying at the center of the CTRAP apparatus. The electrode stack had been assembled by the time I joined this experiment [41]. However, changes to the Ioffe magnet's enclosure and design flaws leading to repeated leaks and long repair times made it necessary to rebuild both the stack and the vacuum enclosure holding it. The fabrication procedure for the electrode stack and the design details for its enclosure are laid out in what follows.

### 5.1 CTRAP Electrodes and Potentials

While a trap with hyperbolic electrodes has an attractive simplicity, there are a number of drawbacks. First, as is clear from Fig. 4.1, holes must be introduced into

the ideal electrode configuration in order to load particles. These lead to anharmonic terms, which depend upon hole sizes and may be compensated for by additional electrodes [50]. A more serious issue is that the electric potential can only effectively confine charges of a single sign. To efficiently make trappable antihydrogen, positrons and antiprotons need to be confined in close proximity. A related problem is the lack of flexibility in the possible potentials. As will be discussed later, a diverse set of potential structures are used to tailor positron and antiproton clouds for optimal antihydrogen production.

To address these issues, we use a stack of independently biased cylindrical electrodes. Such a scheme maintains the cylindrical symmetry of the hyperbolic electrode trap while allowing ample space for particle loading. Dozens of degrees of freedom make it possible to construct complex potentials, simultaneously confine particles of opposite charge, and tune away anharmonicities. This design also has the practical advantage of making electrodes easier to machine and interchangeable.

For cylindrical electrodes, the potentials are most easily found by solving Laplace's equation in cylindrical coordinates. We imagine a single electrode of half-length  $z_0$  and radius  $\rho_0$  at potential  $V_0$ , centered at the origin of a cylindrical coordinate system. On either end, for  $|z| > z_0$  and  $\rho = \rho_0$ , lies a grounded cylinder representing the rest of the electrodes. For calculational simplicity, we introduce grounded planes at  $\pm L$ , where  $L \gg z_0$ . Using rotational invariance, we can write a separable solution to Laplace's equation with  $z$ 's boundary conditions as

$$\psi(\rho, z) = R(\rho)Z(z). \quad (5.1)$$

Inserting it into the partial differential equation and performing the usual manipula-

tions, we get

$$\rho^2 R''(\rho) + \rho R'(\rho) - \rho^2 k^2 R(\rho) = 0 \quad (5.2)$$

$$Z''(z) + k^2 Z(z) = 0. \quad (5.3)$$

Eqn. 5.3 and the boundary conditions yield

$$Z(z) = A_n \cos(k_n z), \quad (5.4)$$

where  $k_n = \frac{\pi}{2L}(1 + 2n)$ ,  $n \in \mathbb{Z}$ , and  $A$  is a constant. On the other hand, Eqn. 5.2 can be reformulated by setting  $\bar{\rho} \equiv k_n \rho$  and  $\bar{R}(\bar{\rho}) \equiv R(\rho) = R(\bar{\rho}/k_n)$ :

$$\bar{\rho}^2 \bar{R}''(\bar{\rho}) + \bar{\rho} \bar{R}'(\bar{\rho}) - \bar{\rho}^2 \bar{R}(\bar{\rho}) = 0. \quad (5.5)$$

This is just the modified Bessel equation with  $\alpha = 0$ . The solution can be written in terms of the modified Bessel function that stays finite at the origin,

$$\bar{R}(\bar{\rho}) = B_n I_0(\bar{\rho}) \quad (5.6)$$

or, in terms of the original  $\rho$  coordinate,

$$R(\rho) = B_n I_0(k_n \rho). \quad (5.7)$$

Putting these components together, we arrive at an expression for the potential.

$$\varphi = \sum_{n=0}^{\infty} C_n I_0(k_n \rho) \cos(k_n z) \quad (5.8)$$

The boundary conditions at  $\rho = \rho_0$  still need to be fixed. Since, for  $n, m \geq 0$

$$\int_{-L}^L dz \cos(k_n z) \cos(k_m z) = L \delta_{n,m}, \quad (5.9)$$

we can integrate Eqn. 5.8 multiplied different factors of  $\cos(k_n z)$  to determine the unknown coefficients. After carrying out this process, we arrive at the following expression.

$$C_n = \frac{1}{LI_0(k_n \rho_0)} \int_{-L}^L dz V(\rho_0, z) \cos(k_n z) \quad (5.10)$$

$$= \frac{2V_0}{LI_0(k_n \rho_0)} \int_0^{z_0} dz \cos(k_n z) \quad (5.11)$$

$$= \frac{2V_0 \sin(k_n z_0)}{Lk_n I_0(k_n \rho_0)} \quad (5.12)$$

With  $C_n$  determined for all  $n \geq 0$ , Eqn. 5.8 becomes

$$\varphi = \sum_{n=0}^{\infty} \frac{2V_0}{Lk_n} \frac{I_0(k_n \rho)}{I_0(k_n \rho_0)} \sin(k_n z_0) \cos(k_n z). \quad (5.13)$$

The potential generated by several electrodes at different potentials may be calculated by translating and superposing expressions of the form Eqn. 5.13.

While cylindrical coordinates are natural for computing the potentials inside cylindrical electrodes, our goal is to approximate a quadrupole field. Note that

$$\frac{V_0}{2} \left(\frac{r}{d}\right)^2 P_2(\cos \theta) = \frac{V_0}{2d^2} \left(z^2 - \frac{1}{2}\rho^2\right), \quad (5.14)$$

which is the expression for the quadrupole potential in Eqn. 4.1. The general expression for solutions Laplace's equation in spherical coordinates is

$$\varphi(r, \theta, \phi) = \frac{V_0}{2} \sum_{l=0}^{\infty} D_{2l} \left(\frac{r}{d}\right)^{2l} P_{2l}(\cos \theta) \quad (5.15)$$

when the boundary conditions exhibit rotational symmetry about the z-axis, there is reflection symmetry across the xy-plane, and the origin is in the region of interest. Note that  $d$ , defined by Eqn. 4.2 with  $\rho$  and  $z_0$  taking on the meanings assigned in the current context, and  $V_0$  make each  $D_{2l}$  dimensionless. With a solution of the form Eqn. 5.15, setting up a quadrupole potential amounts to minimizing  $D_{l \neq 1}$ .

Determining the coefficients  $D_{2l}$  from the nontrivial spherical versions of the boundary conditions discussed above is difficult. Instead, let's set Eqn. 5.15 equal to Eqn. 5.13, set  $\rho = 0$  in cylindrical coordinates, and set  $\theta = 0, r = z$  in spherical coordinates. After expanding cosines in powers of  $z$  and setting terms of the same order in  $z$  equal, we arrive at

$$D_{2l} = \frac{2}{V_0} \frac{(-1)^l}{(2l)!} \sum_{n=0}^{\infty} C_n (k_n d)^{2l} \quad (5.16)$$

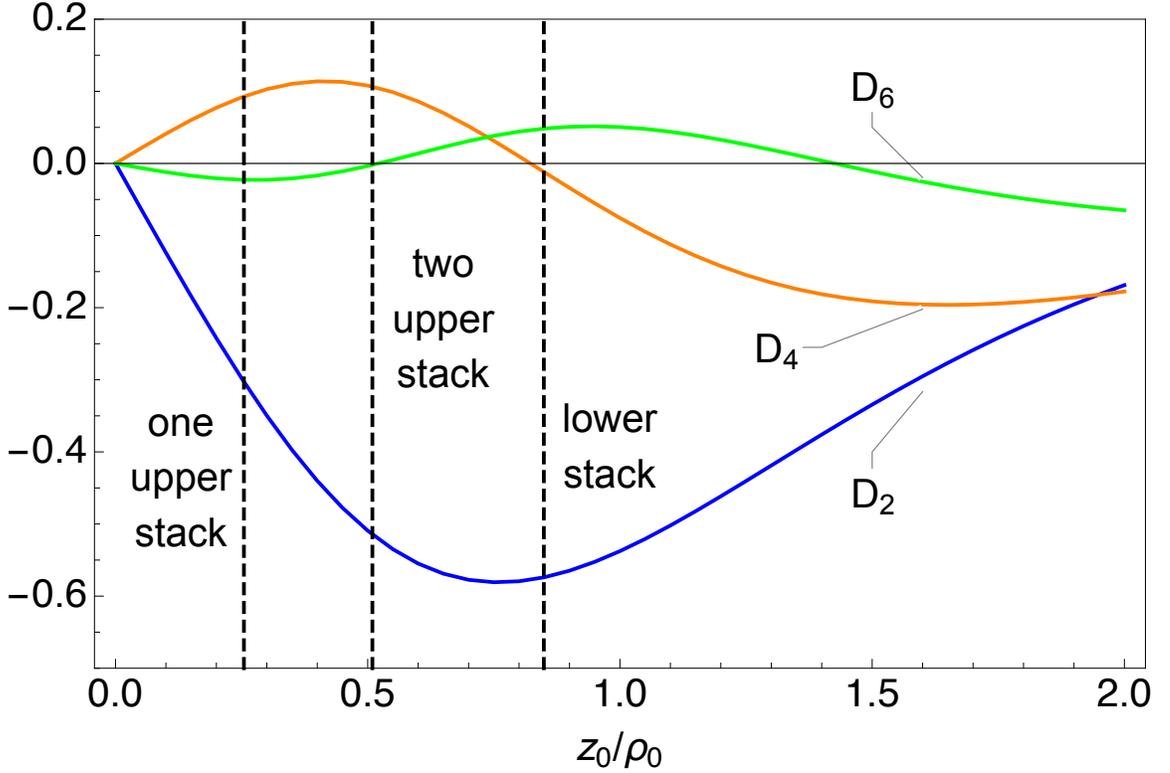


Figure 5.1:  $D_{2l}$  coefficients from Eqn. 5.16, plotted as a function of the ratio  $z_0/\rho_0$ . The dimensions of one and two upper stack electrodes and one lower stack electrode are indicated by the vertical lines. An ideal quadrupole has  $D_{l \neq 1} = 0$ . Here,  $L = 100\rho_0$  and 1000 terms are used in Eqn. 5.16. Increasing both of these by an order of magnitude does not perceptibly change the appearance of the curves.

The three lowest-order coefficients are plotted as a function of electrode dimensions

in Fig. 5.1. Lower stack electrode lengths were chosen to minimize the 4th-order term while keeping the 6th-order term small. Electrodes in the upper stack are optimized to have a large radius in order to make the magnetic Ioffe trap as deep as possible. With radii of 34 mm, they would need to be about 58 mm long to match the length-to-radius ratio in the lower stack. Since the Ioffe magnet is slightly longer than 300 mm, 5 electrodes of this length would fit in the bore. In order to buy more flexibility for creating complex wells during antihydrogen synthesis and detection, we initially chose a value of  $z_0/\rho_0$  that minimizes the 6th order term while keeping the 4th-order term relatively small. There were still concerns about the number of degrees of freedom for potential shaping, so the length was cut in half. This allows for reasonably harmonic two-electrode wells without compromising the control needed for perfecting antihydrogen production.

Estimating potentials in terms of power series is useful for understanding their properties near the origin. However, the volume over which these series represent good approximations may be limited by a finite radius of convergence or by the number of terms included in the expansion. This problem may be overcome by numerically solving Laplace's equation. This is the method used for calculating and visualizing potentials in real time in CTRAP. A comparison is made between the numerical and analytic estimates in Fig. 5.2.

## 5.2 Electrode Construction

When I joined the experiment, both the upper and lower electrode stacks of CTRAP had already been constructed. However, modifications to the Ioffe mag-

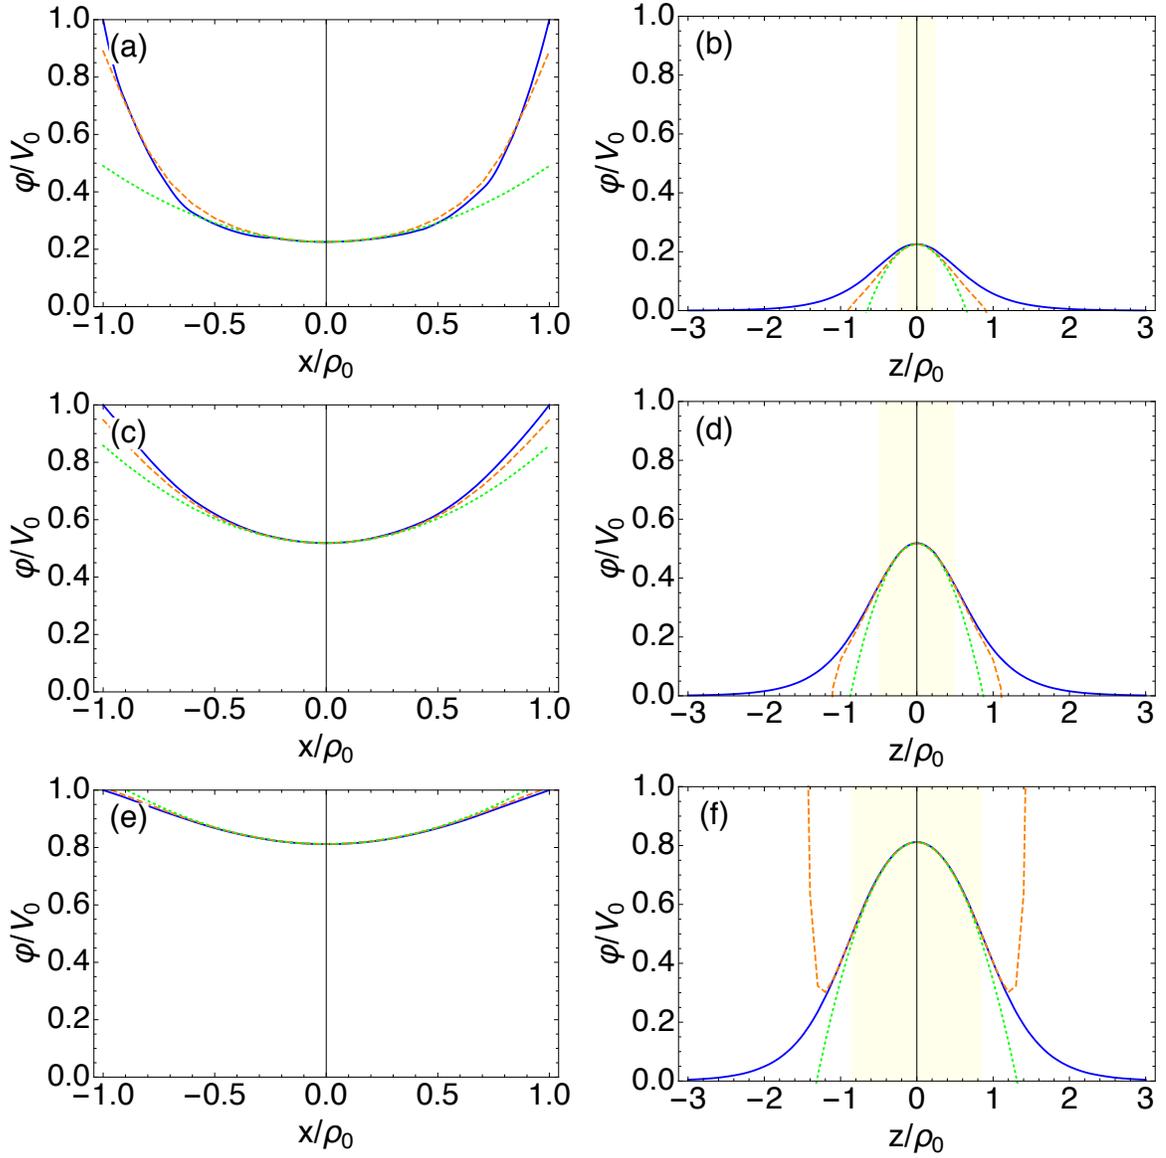


Figure 5.2: Potentials are determined by numerical calculation (solid blue) and series expansion (dashed orange), along with the harmonic approximation (dotted green). Cross sections are shown at  $y = z = 0$  (a, c, e) and  $x = y = 0$  (b, d, f), with the length subtended by the electrode colored yellow. The first pair of plots are for the one-electrode upper stack geometry from Fig. 5.1, the second pair is for the two-electrode upper stack geometry, and the final pair is for the lower stack single-electrode geometry.

net enclosure reduced the radius of the inner bore in which the upper stack sits, forcing us to rebuild that part of the trap.

The original upper stack was designed and built by Phil Richerme. Sample electrodes are shown in Fig. 5.3. Features include a slit to prevent eddy currents during fast magnet manipulations, a tab for minimizing the amount of metal heated when soldering in wires, and walls with large internal radii to maximize the Ioffe trap depth. The G10 sleeves both give rigidity to the thin, slitted copper electrodes and maintain separation between neighbors. They also have channels to accommodate the microcoax cables attached to each electrode.

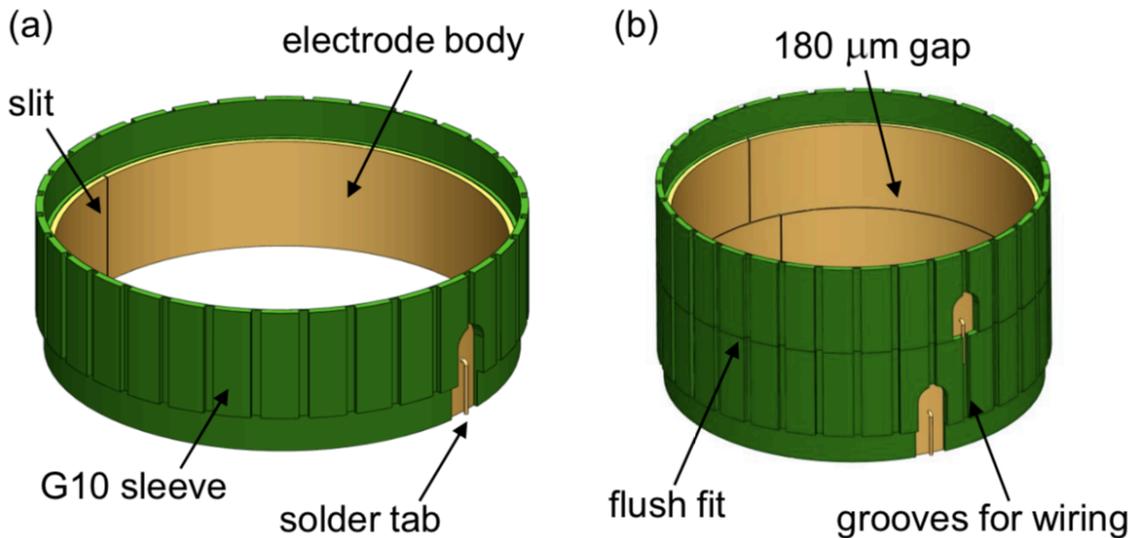


Figure 5.3: Original CTRAP electrode design [41] with (a) showing a single electrode and (b) demonstrating the fit between a pair of electrodes.

The procedure I used to produce electrodes for the new stack roughly follows the one outlined in Phil's thesis [41]. Slight modifications are made for special electrodes, like the four-segment rotating wall electrode and the conical electrode at the base of the upper stack. To start, an oxygen-free copper cylinder of the appropriate dimen-

sions is machined in Harvard's shop. The cylinder is then brought to T.O.C. Finishing to polish the interior surface. Profilometer measurements show resulting average surface roughnesses of between 2 and 4 micro-inches. Next, a small strip of copper is pressed flush to the electrode's outer wall while heat guns uniformly increased the temperature of the ensemble. When hot enough, solder is used to bind the copper tab to the cylinder. With the tab in place, the electrode is sent to Absolute Metal Finishing, where an acid pre-wash strips away oxidation from the previous step before a gold plating is applied. Importantly, the nickel layer usually deposited between copper surfaces and gold plating is not used because of its magnetic properties. The results are shown in (a) through (c) of Fig. 5.4.

With the electrode finished, the next step is to cut a slit and add the outer sleeve. Before cutting an electrode, it must be secured in a clamp, as in (d) of Fig. 5.4. This holds the electrode in place during the cut, and it prevents the freed edges from springing outward, which would deform the cylindrical body. The electrodes are wire EDM cut at Boston University's machine shop with a 0.007" wire. Finally, with the clamp still in place, a sleeve, machined in Harvard's shop, is slipped around the electrode, clamped in place using an overhead plate and threaded rods passing through the holes at the four corners of the jig shown in Fig. 5.4 (d) and (e), and glued to the electrode using Stycast 2850 epoxy. After a 24 hour curing period, the finished ensemble is removed from its clamps and is ready for use.

A few lessons were learned in the course of preparing these electrodes. Its useful to enumerate them as words of caution and advice for future electrode builders.

- Most electrodes had three glue joints holding them to their sleeves: one near



Figure 5.4: The state of an electrode after different stages of production are shown. These stages include (a) polishing, (b) soldering, and (c) plating. Electrodes are also shown before (d) the slit is cut and (e) the sleeve is glued in place.

the solder tab and two straddling the slit. Many of these failed after several thermal cycles, causing shorts. Replacements, designed with more holes in their sleeves for glue joints, have so far fared better.

- Because its magnetic susceptibility is two orders of magnitude smaller than G10's at 4 K [49], we decided to make the sleeves from linen phenolic. This material has several unfortunate properties. First, edges along which it's been cut have lots of fibers that are difficult to remove. Second, when the electrodes are assembled, it's often necessary to twist one to line up their channels. This is surprisingly hard to do with the linen phenolic sleeves, forcing stack builders to exert large torques on a delicate system. Finally, the linen phenolic sleeves are much more easily bent than their G10 counterparts. As they are responsible for keeping the electrodes from springing out of round once cut, this weakness may lead to warped electric potentials.
- It's known that copper diffuses into gold plating. However, the fact that our 10-20 micro-inch plating showed evidence of this after a few weeks was surprising given known rates for this process [67]. The affected electrodes were re-plated, but those produced after this batch were given 80-100 micro-inch layers of gold.

The process outlined in this section successfully produced electrodes compatible with the new Ioffe trap enclosure in a quick and efficient manner. Problems arose with the new upper stack, but swift responses provided solutions that allowed the experiment to continue to make progress. Still, a new design is clearly needed. In particular, a means more robust than cryogenic epoxy joints must be found to keep electrodes separated without sacrificing the large inner radius we've achieved.

## 5.3 Enclosure Construction

In addition to rebuilding the upper electrode stack, I also redesigned and rebuilt the vacuum enclosures (also referred to as the pinbases) for both the upper and lower stacks. This was not strictly necessary, but leaks repeatedly slowed experimental progress. The new enclosures have proven to be more robust than their predecessors while also allowing for faster repairs in case a leak does arise.

The old enclosure is shown in Fig. 5.5. The portion shown is just below the upper stack. Not shown is the thermal isolation bellows and the Ioffe enclosure, which usually wrap around the upper stack. The weaknesses of this design are obvious.

- Thirty feedthrough pins are brazed to a custom crescent flange. If one pin breaks, the entire flange must be replaced, which is expensive and slow.
- The way the pins are laid out is problematic. It's easy for a tilted pin or a small fleck of metal to cause unintended shorts. Operating on any pin is similarly difficult as collisions with nearby pins and dropped pieces of solder can lead to breaks or shorts.
- Routes from pins to their electrodes can overlap, causing messy, error and damage-prone wiring.
- The spaces between the filter boards and the flanges are extremely tight. In order to tighten or remove the flanges, custom angled screwdrivers are required. A stripped screw requires the entire filterboard to be removed.
- If more indium is needed to make the seal between a copper crescent flange

and the titanium it mates to, all electrical connections need to be removed. Attempts have been made to insert indium without doing this, but so far, they have all been unsuccessful.

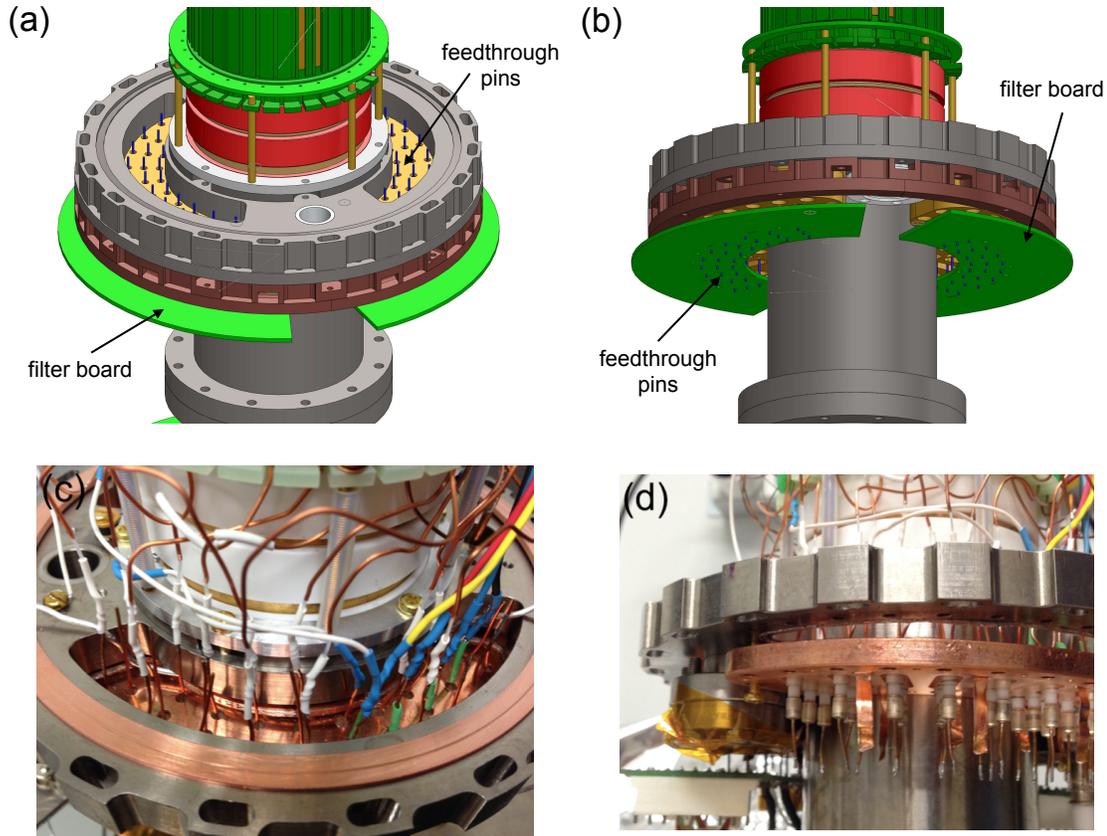


Figure 5.5: (a) and (c) Old upper stack pinbase from the top (experiment vacuum side), (b) and (d) Old upper stack pinbase from the bottom (insulating vacuum side)

The problem, noted above, about changing a flange's indium is both understated and a frequent issue. Approximately once every second trap cool down, the crescent flanges, which separate experiment and insulating vacuum spaces, leaked. Tightening the flanges more rarely fixed the leak. Consequently, the electrode stack needed to be removed from the Ioffe enclosure, all electrical connections had to be taken off of the

problematic flange, the thin micro-coaxes were often damaged during their removal so much of the electrode stack needed to be rewired, the flange would be put back in place, the connections with the microcoax would be reestablished, and the electrode stack would be reinserted into the Ioffe enclosure. This whole process took two to two and half weeks, which quickly became unacceptable for one leaky flange.

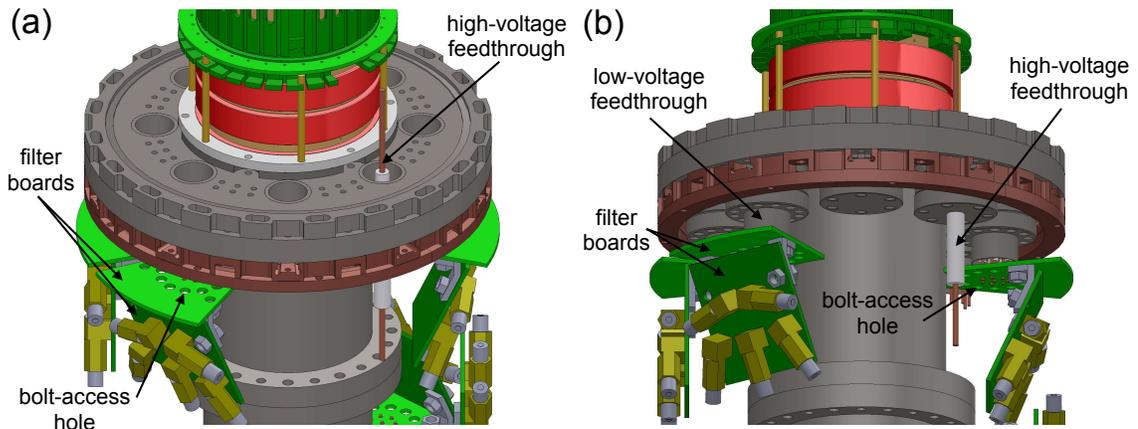


Figure 5.6: (a) New upper stack pinbase from the top (experiment vacuum side), (b) New upper stack pinbase from the bottom (insulating vacuum side)

As a fix to these problems, I designed a new, all-titanium, all-conflat (CF) vacuum enclosure for the electrode stack. While the feedthrough flanges and blanks are stock parts, many of the other components are custom-made (MPF Products). The filter boards are also custom jobs, fitting the same circuits and connectors as the old filter boards. The finished products are shown in Figs. 5.7 and 5.8.

There are several important features that make this version more robust and easy to work with. A subset are listed below.

- All vacuum seals are CF, so no tightening is necessary between thermal cycles.
- The main bodies of the upper and lower stack vacuum pieces are interchange-

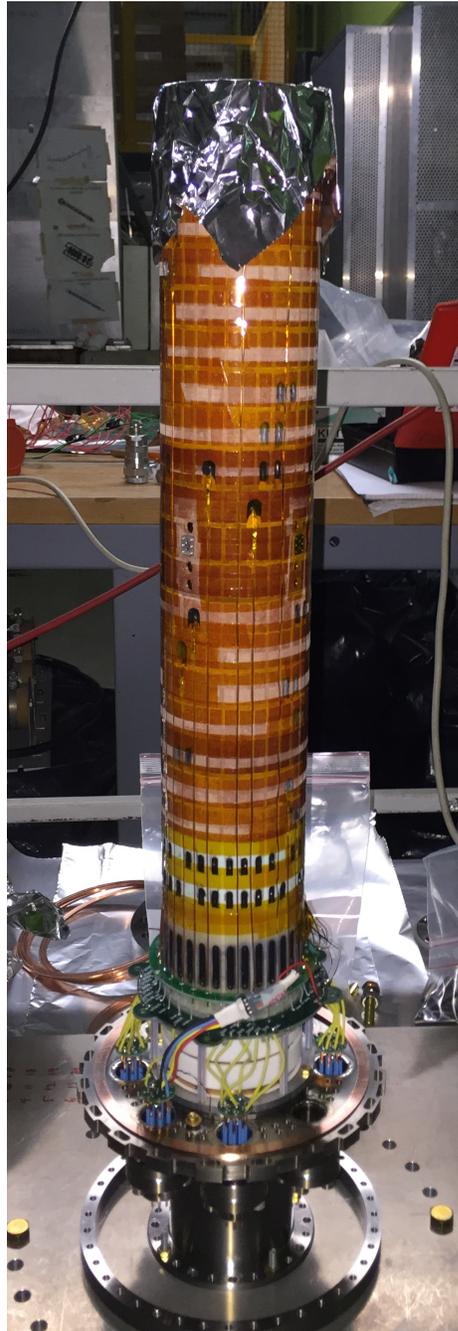


Figure 5.7: New all-CF vacuum pinbase with the upper stack built on top. The foil above the electrode stack prevents concrete dust from the AD from contaminating electrode surfaces. The kapton tape wrapped around the electrodes holds the microcoax in place and prevents it from being pulled when the stack is inserted into the Ioffe enclosure. Notice that the bottom three (replacement) electrodes have G10 sleeves and lots of glue joints.

able.

- Two types of filter boards were produced. One holds eight copies of the general circuit used by most electrodes. The other holds two of these general circuits as well as specialized circuits for two specific types of electrodes (rotating wall and split, though the split circuit wasn't used in CTRAP). Thus, the filter board layout minimally constrains where each electrode is connected.
- Holes are built into the filter boards to ensure easy access to CF bolts.
- Redundancies are built in so stripping a tapped hole does not require us to replace expensive custom vacuum parts. For most CF seals, only half of the available holes are used. If a tap is stripped, all bolts may be moved by one hole and make an equivalent seal. Where this is not possible, plate nuts are used.
- Paths from electrodes to their feedthrough pin follow a two-step path. Microcoax goes straight down the upper stack, through the sleeve channels, and connects to a distribution board. Traces in the distribution board connect the microcoax to copper wires (yellow insulation in Fig. 5.7), which join the distribution board and small feedthrough boards via short paths. This way, the messiness of routing different electrode wires is handled by circuitboard traces. Finally, the small boards on top of the feedthroughs make connections between copper wires and feedthrough pins less awkward.
- Although this has already been stated, it's worth reiterating that all feedthroughs are stock parts, and therefore quickly, cheaply, and easily replaceable.



Figure 5.8: Pinbases for the upper and lower stacks are joined together, the upper stack has been inserted into the Ioffe enclosure, and lines coming down from the hat are plugged into the filter boards. Stainless steel coaxial cables mate to the filter board SMA connectors while the white rectangular blocks join the low-frequency bias lines.

The new pinbases have been a part of CTRAP since 2016. While the old design made these components a constant maintenance chore, the new version has operated fault-free. Even if a leak arises, the process for fixing it will be much quicker and easier. In both cases, this upgrade improves CTRAP's uptime for performing experiments.

## 5.4 Summary

In this chapter, many of the details of CTRAP's system allowing for the trapping and manipulation of charged particles were explained. It was shown that although an ideal quadrupole potential is not generated by the electrodes, their dimensions may be set so it is easy to make nearly-quadratic potentials. A description of the electrode fabrication procedure I used to build the upper stack is then given. Finally, a summary of the electrode enclosure design explains the flaws of the old system and the ways in which the new components fix these problems. As a result of these construction projects, the electrode stack ensemble is compatible with the new magnet enclosure, holds vacuum robustly under thermal cycling, and is far quicker to maintain when problems arise.

# Chapter 6

## Experimental Procedures and Results

The (A)TRAP collaboration has accumulated decades of experience manipulating and trapping electrons, positrons, and antiprotons. Over that time, certain techniques have become standard. More recent antihydrogen trapping uses methods that are still being perfected. This chapter collects the procedures, both new and old, that we used to create and trap antihydrogen in 2018. The results of our trapping experiments are also reported.

### 6.1 Particle Counting

One of the most important things to understand about the clouds held in our Penning traps is the number of particles they contain. There are two principle methods used to get this information. For antimatter species, annihilations with trap electrodes cause detectors to be triggered. Known detection efficiencies can be used to determine the number of annihilated particles. On the other hand, by deposit-

ing a plasma on an electrode connected to a circuit designed to measure charge, the population of that plasma may be deduced.

The antiproton plasmas used include  $10^5$  to  $10^6$  particles, which makes charge counting susceptible to noise. To count annihilations, the antiprotons are confined in a potential well at a single electrode, and a ramp is constructed such that escaping particles are sent to the degrader. A trigger is then sent over fiber to a multi-channel scaler (MCS), a high-speed analogue-to-digital converter (HSADC), and the source of the bias for the electrode holding the antiprotons. Once triggered, the bias source decreases the electrode's voltage until the well is inverted. The MCS and HSADC record the number of detected particles and the changing voltage, giving the number of antiprotons lost as a function of the potential of the confining electrode.

Electron and positron plasmas are usually made up of  $10^7$  to  $2 \cdot 10^8$  members. Clouds in this population range are easily measured by our charge counting system, which has difficulty distinguishing fewer than  $10^5$  particles at once from noise. To count these species, we use a charge-sensitive preamplifier. The potential setup is the same as for the annihilation measurement except that the degrader is held at a higher voltage than nearby electrodes to suppress the emission of secondary electrons, which may skew the count.

To begin, the barrier separating the potential well holding the particles from the ramp to the degrader is lowered by 5 V for 75 ns. Particles spill onto the degrader, charging up the 1 pF capacitor in Fig. 6.1 and causing the oscilloscope to see a spike in voltage. The voltage read on the oscilloscope can be used to determine the number

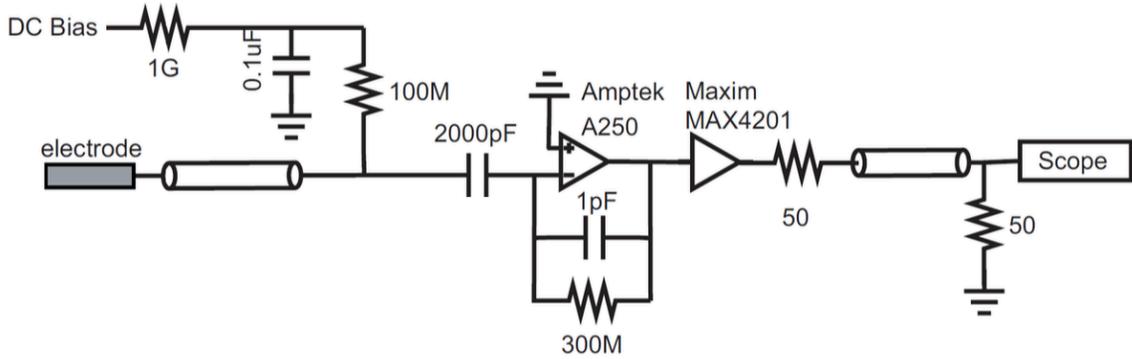


Figure 6.1: Diagram of CTRAP's charge-sensitive preamplifier [41]

of particles of charge  $q$  hitting the degrader by

$$N = -2VC/q, \quad (6.1)$$

where  $C$  is modified from 1 pF by other capacitances in the system, but can be determined beforehand using calibrated charges. The potential on the confining electrode is then increased or decreased by 3 V for positive and negative species, respectively. This process repeats until the potential well confining the particles has been inverted.

## 6.2 Plasma Geometry

Recall that in thermal equilibrium, a plasma occupying a Penning trap with a purely quadratic potential can be approximately described as a steadily-rotating, constant-density spheroid. The quantities used to describe such plasmas include  $N$ ,  $n_0$ ,  $\omega_r$ ,  $r_p$ ,  $z_p$ ,  $\alpha_a$ , and  $\omega_p$ . (Respectively, these are the number of particles, density, rotational frequency, radius, half-length, aspect ratio, and plasma frequency of the plasma.) Eqs. 4.21, 4.25, 4.26, and 4.27, along with

$$N = n_0 \frac{4}{3} \pi r_p^2 z_p \quad (6.2)$$

connect these quantities. Consequently, if two are measured and the electric and magnetic fields are known, the rest can be deduced. We have already discussed how  $N$  is determined for different species. Only one more measurement is needed to fully specify a plasma to a good approximation. Trappable antihydrogen production rates depend critically on the density of the positron plasma and on the radii of both the positron and antiproton plasmas. Consequently, such knowledge is critical for refining CTRAP's antihydrogen synthesis techniques.

In this section, we discuss two measurements that supplement  $N$  to round out the understanding of plasmas in CTRAP's Penning trap. The first is non-destructive and uses resonances to determine  $\alpha_a$ . The second is a (destructive) direct image of the plasma on the XY stage's phosphor screen. This not only serves to measure  $r_p$ , but also has the potential to give information on whether the radial distribution of the plasma matches the uniform-density spheroid assumption.

Historically, we have only used modes to determine plasma geometries. Consequently, the procedure for making this measurement is, at present, better developed and more repeatable. The motivation for this technique comes from [68]. There, modes are described for zero-temperature plasmas in an ideal Penning trap under small perturbations. The frequencies of the azimuthally symmetric modes are given by

$$1 - \frac{\omega_p^2}{\omega_l^2} = \frac{k_2 P_l(k_1) Q_l'(k_2)}{k_1 P_l'(k_1) Q_l(k_2)}, \quad (6.3)$$

where

$$k_1 = \frac{\alpha_a}{(\alpha_a^2 - 1 + (\omega_p/\omega_l)^2)^{1/2}} \quad (6.4)$$

$$k_2 = \frac{\alpha_a}{(\alpha_a^2 - 1)^{1/2}}. \quad (6.5)$$

The equation for  $l = 1$  is identical to Eqn. 4.27, with  $\omega_1$  replacing  $\omega_z$ . This means that  $\omega_z$  can be measured directly. The second mode we measure, which can be interpreted as an oscillation in  $\alpha_a$ , has the frequency  $\omega_2$ . Fig. 6.2 shows the low-order mode frequencies, relative to  $\omega_z$ , as a function of aspect ratio.

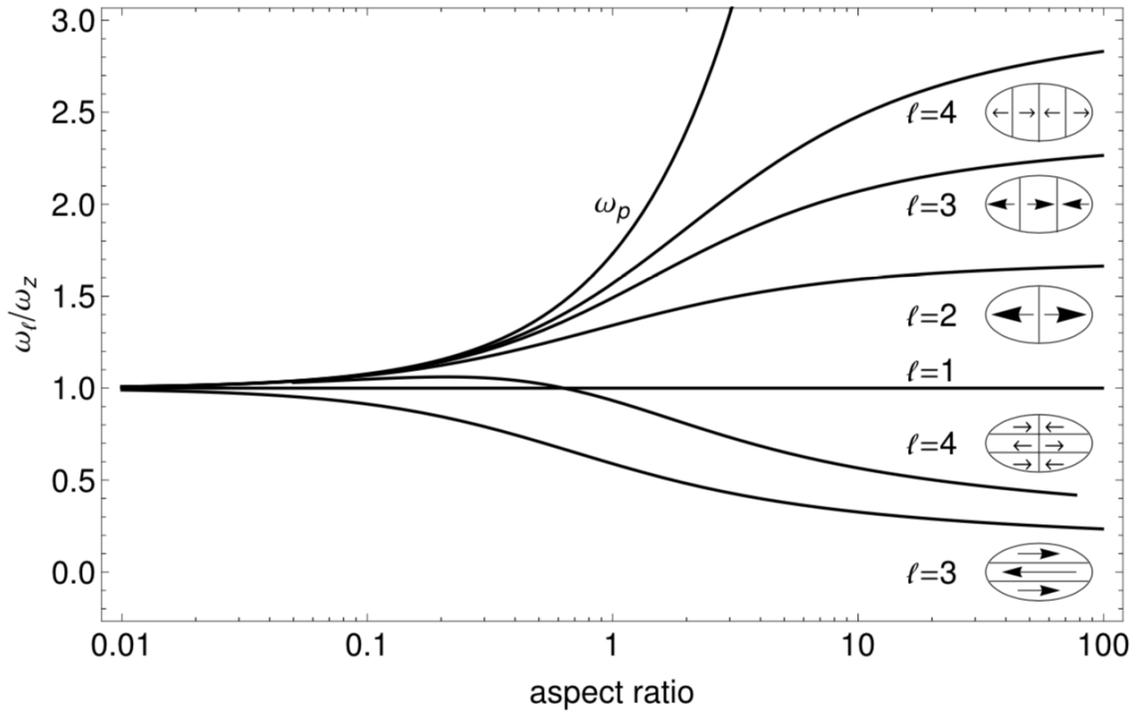


Figure 6.2: Low-order mode frequencies as a function of aspect ratio, adapted in [41] from [69]

To make a mode frequency measurement, the plasma to be queried is first put into a nearly harmonic well. A frequency synthesizer (PTS 250) is then set to a specified frequency. RF switches allow the drive to reach an electrode adjacent to the center

of the well for  $1 \mu\text{s}$ . The excited modes continue to ring after the drive has ended, inducing oscillating image charges on nearby electrodes. Voltage drops caused by these currents are taken at the hat from one of the harmonic well's electrodes, sent through a series of filters and amplifiers, and mixed with the drive frequency. The resulting signal is low-pass filtered ( $f_c = 2.5 \text{ MHz}$ ), measured  $3 \mu\text{s}$  after the drive ends by a 24-bit oscilloscope (PXI 5922), and fast Fourier transformed, revealing any nearby resonances.

While this measurement technique is robust, giving nearly identical results for identically prepared plasmas, there are some significant drawbacks. In the lower electrode stack, we are able to detect modes for as few as 2 million particles while in the upper stack, where the electrodes have almost twice the radius, we're limited to plasmas containing more than 10 million particles. Both of these are well beyond the sizes of our largest antiproton plasmas. Another problem is that the measurement requires the particles to be confined in a near-harmonic well. This can be circumvented by using a program like RATTLE [70] to compute modes of different plasma shapes for different potentials. However, without somehow speeding up the calculation, this would be too slow for real-time diagnostics when experimenting with new potential structures. Neither of these are necessarily insurmountable, but they do represent areas where significant work is needed to expand the usefulness of this test.

As a solution to these problems, as well as a source for information on non-spheroidal plasma geometries, we developed an imaging system. Fig. 6.3 shows the underside of the flange carried by the XY stage. The imaging assembly and the microwave system were both added in 2018. While at moment, more development is

needed to make it useful, the latter is intended for measuring magnetic fields in the trap, selecting trapped antihydrogen states, and performing hyperfine spectroscopy. Plasma imaging, on the other hand, has already been successfully implemented in CTRAP.

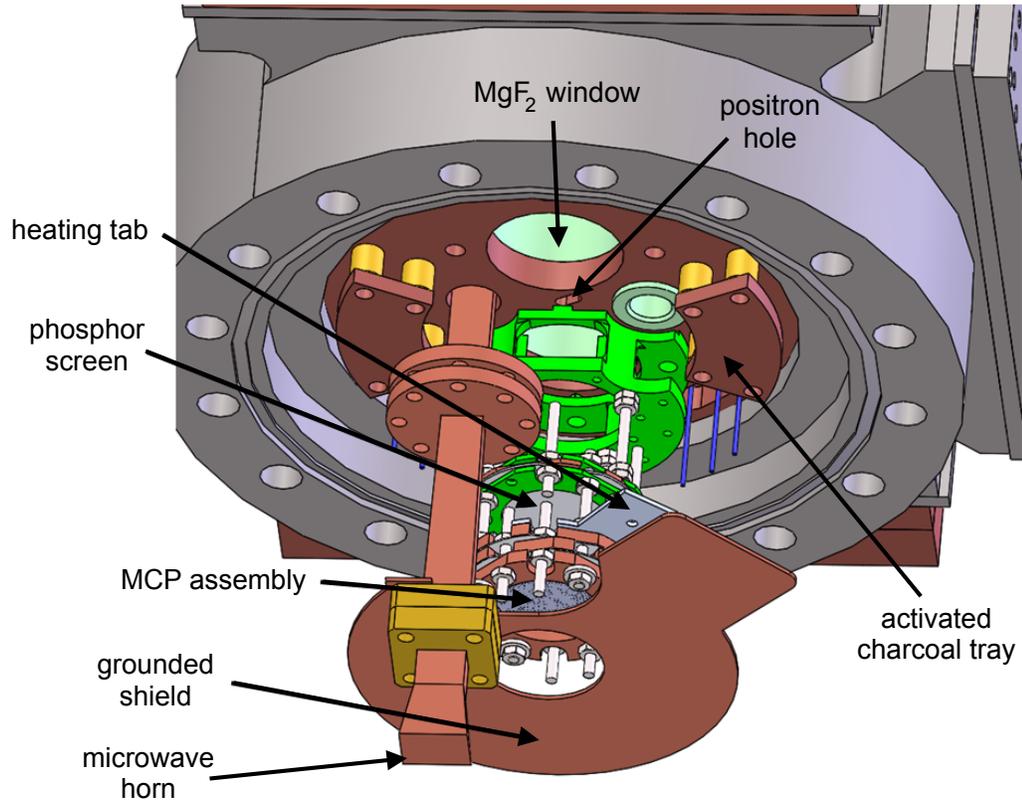


Figure 6.3: Plasma imaging system

At the bottom of the imaging assembly is a grounded piece of copper, which both prevents the insulators above it from charging up and shields the potentials near the top of the electrode stack from those on the multi-channel plates (MCPs) above. The MCP assembly is a stack of two MCPs (Hamamatsu F1094-01) sandwiched between three rings which can be independently biased. During normal operation the plates,

from bottom to top, are at 100 V, 1 kV, and 2 kV. As multiple images deplete the number of electrons in an MCP's channel at its observed temperature of 12 K, a heating tab is used, together with a diode laser, to raise the MCPs to 20 K between images. This has shown to be both necessary and sufficient to have the system function as expected. Finally, above the MCP assembly lies the phosphor screen (Kimball Physics PHOS-UP22GL), biased to 3 kV in order to pull in electrons from the MCP. The screen emits at around 550 nm and has a decay time of 4 ms.

To use this system, the plasma to be imaged is first confined to a single electrode. A voltage ramp is constructed such that released particles are forced to the bottom plate of the MCP assembly, with a "back door" preventing them from traveling the opposite direction. Once the proper potential structure has been constructed, a trigger is sent both to a pulse generator and a camera located outside of the vacuum system, about 1.5 m above the XY stage. The camera takes a series of pictures, and one that appears to capture a strong signal from the phosphor screen is manually selected. The pixel intensities are then fit to a Gaussian density profile, and a radius is reported. The visible section of the plasma screen, which is 20 mm in diameter, sets the scale for the measurement.

Fig. 6.4 shows a typical set of images from antiprotons, positrons, and electrons. As suggested by the pictures, many of the electron and positron plasmas appeared, at least near their centers, to have Gaussian profiles. It's also worth noting that, when modes measurements were performed prior to imaging, the radii generally agreed to within 20%. On the other hand, the antiproton clouds were often large, lighting up a large portion of the phosphor screen, but without definite, reproducible structure.

Consequently, fits were usually poor.

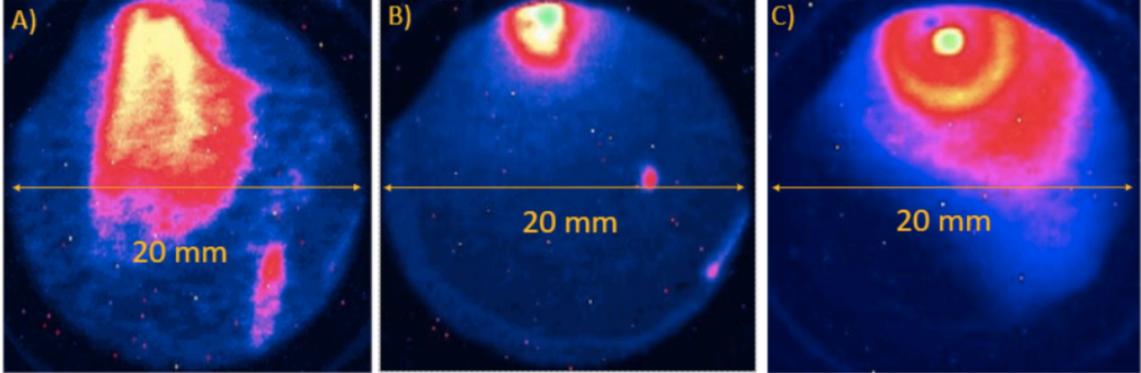


Figure 6.4: Plasmas are shown with (a) 360 thousand antiprotons, (b) 20 million positrons, and 30 million electrons [71].

While not yet fully refined, the progress made so far with plasma imaging represents a substantial step forward in our ability to understand the clouds of particles confined within CTRAP. In the absence of the time-crunch imposed by daily beam shifts, the following improvements can be made.

- With CTRAP cold, the motors could not exert enough torque to get the phosphor screen centered. Consequently, many plasmas are cut off by the edge of the screen. Improvements to the XY stage's translation mechanism are needed to remedy this problem.
- More work needs to be done studying the trajectory of particles pulsed to the screen. Clouds imaged from different parts of the electrode stack appear to have slightly different centers. One possible explanation is that the axes of the electrodes and the magnetic field are not perfectly aligned. A misaligned or slight fringing of the magnetic field may also account for distortions in the antiproton plasmas [72].

- Despite appearances from Fig. 6.4, many of the raw images are quite faint. This may be fixed by changing the voltage drops across the imaging assembly, moving the camera closer to the phosphor screen, or adding more intermediate MCPs. A crisper picture should give more detailed information on our plasmas.

With these modifications, imaging in CTRAP should serve as a solid guide for optimizing plasma manipulations.

### 6.3 Rotating Wall

It is important to be able to manipulate both the density and the radius of clouds held in CTRAP. Getting these parameters right is critical when trapping antiprotons and positrons, moving particles through a changing field, and making antihydrogen. From Eqn. 4.20, we see that the density can be increased by forcing the plasma to spin faster. Eqns. 4.25-4.27 show that this forces the radius to decrease as well.

The technique of applying a rotating wall to trapped plasmas is one way to achieve faster rotations, higher densities, and smaller radii. This approach uses a ring electrode, cut vertically into  $n$  segments, each subtending the same angle about the symmetry axis. Sinusoidal drives offset by  $2\pi/n$  are applied to each segment, giving rise to a rotating electrical field. In CTRAP, we use an  $n$  of 4.

The idea of a rotating wall is simple in principle. A rotating electric field exerts a torque on the trapped plasma, changing  $\omega_r$ . However, the reality is slightly more complicated. Early experiments in plasmas used weak drives to excite Trivelpiece-Gould (TG) modes in electron [73] as well as ion [74] plasmas, leading to compression or expansion. Density manipulations were only effective around discrete frequencies.

Increasing the amplitude of the drive was shown to increase the rate of density change, heat the plasmas, and broaden the bands at which the drive worked.

More recently, a strong-drive regime has been explored [75, 76]. With even higher amplitude drives than were used in previous studies, compression was observed over a wide (8 MHz), unbroken range of frequencies. While TG mode excitation generally leads to plasmas rotating much more slowly than the applied drive, these demonstrations prepared clouds with  $\omega_r \approx \omega_{drive}$  over much of the frequency interval probed. Consequently, this technique allows for precise control of the density.

CTRAP, as BTRAP before it, uses a custom sinusoidal waveform generator with 4 ports. As implied above, signals from these ports are offset in phase by  $\pi/2$ . Amplitudes of 1 to 9 Vpp and frequencies from 100 kHz to 10 MHz are typical as parameters to far outside this range lead to distorted signals. (Settings outside these intervals have been tried, but they not been found to outperform optimal values within these ranges.) Studies done during BTRAP's operation established the optimal well configuration for maximal lossless plasma compression in the lower electrode stack, which is identical to CTRAP's [41]. In both traps, radius reduction occurs over a wide range of frequencies, which is similar to the strong-drive behavior. However, the final radius is largely insensitive to the drive frequency, which contrasts with expectations if we're truly in the this regime.

## 6.4 Electron Loading

We now move from general techniques, applicable to many plasma species, to specific procedures used during the antihydrogen synthesis process.

Despite not being part of the antihydrogen atom, electrons play an important role in the sequence of steps we use to produce these atoms. Both positrons and antiprotons enter CTRAP at high energies. While positrons can radiatively cool into a potential well, collisional cooling with electrons speeds this process, increasing the loading rate. Antiprotons, on the other hand, take much longer to cool on their own, so electrons are essential for capturing them.

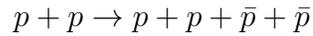
CTRAP loads electrons via photoemission from the degrader. The photon source is a KrF excimer laser (GAM Lasers EX-5), which produces 20 mJ, 10 ns pulses at 248 nm. In practice, we operate the laser at a pulse energy of 3 mJ, which lengthens its lifetime and loads electrons slowly enough to allow for more control over the final number captured. The 248 nm light, which corresponds to 5.0 eV/photon, can easily liberate electrons from the beryllium degrader, which has a work function of 3.6 eV [77].

To load electrons, a 1 V deep potential well is first constructed near the degrader. Small negative biases on electrodes above and below this well prevent the electrons from continuing up the stack and reflecting back to the degrader, respectively. A trigger is then sent to both the laser and the front door electrode (between the well and the degrader). When the laser is triggered, a series of optics mounted both inside and outside a vacuum cube above the hat send the pulse down toward the XY stage. The beam passes through the positron hole, striking the degrader. The other trigger, which is delayed by 1.8  $\mu$ s, opens the front door for 600 ns, allowing electrons to enter, but preventing most of them from exiting. The potential well is deepened by 1V, and after 1 s, another pair of triggers are sent. With this procedure, it takes 20

to 30 pulses to load  $10^8$  electrons.

## 6.5 Antiproton Loading

CERN's Antiproton Decelerator (AD) facility boasts the best source of low-energy antiprotons in the world. Protons from the nearby Proton Synchrotron with are sent to the AD, where they strike an iridium target [78]. Antiprotons are formed in the reaction



with kinetic energies around 2.7 GeV. Stochastic and electron cooling steps are used to reduce the energies to 5.3 MeV. Approximately every 90 s, 30 to 37 million antiprotons are sent through magnetic beamlines to one of several experimental zones, where they arrive in a 200 ns interval.

The antiproton bunches that travel toward CTRAP still need some adjustment before entering the trap. First, even when the antiprotons are centered on all of the beamline detectors leading to our zone, they often arrive in the trap off-center. Second, the energies are still too high for our kV-sized well. Both of these problems are addressed by the stack of foils and detectors below the entrance to the lower stack.

Fig. 6.5 shows this setup. About 1 m below this is a 5-segment silicon detector that can be moved in an out of the antiproton beam for coarse steering. The PPAC (Parallel Plate Avalanche Counter) then allows for fine-tuning the steering right before the beam enters the trap. It consists of two orthogonal sets of 5 aluminum strips each, all 2 mm wide, separated by 0.5 mm, and with biases of 150 V. The cell is filled with argon gas, which is ionized as antiprotons pass through. The freed electrons are drawn

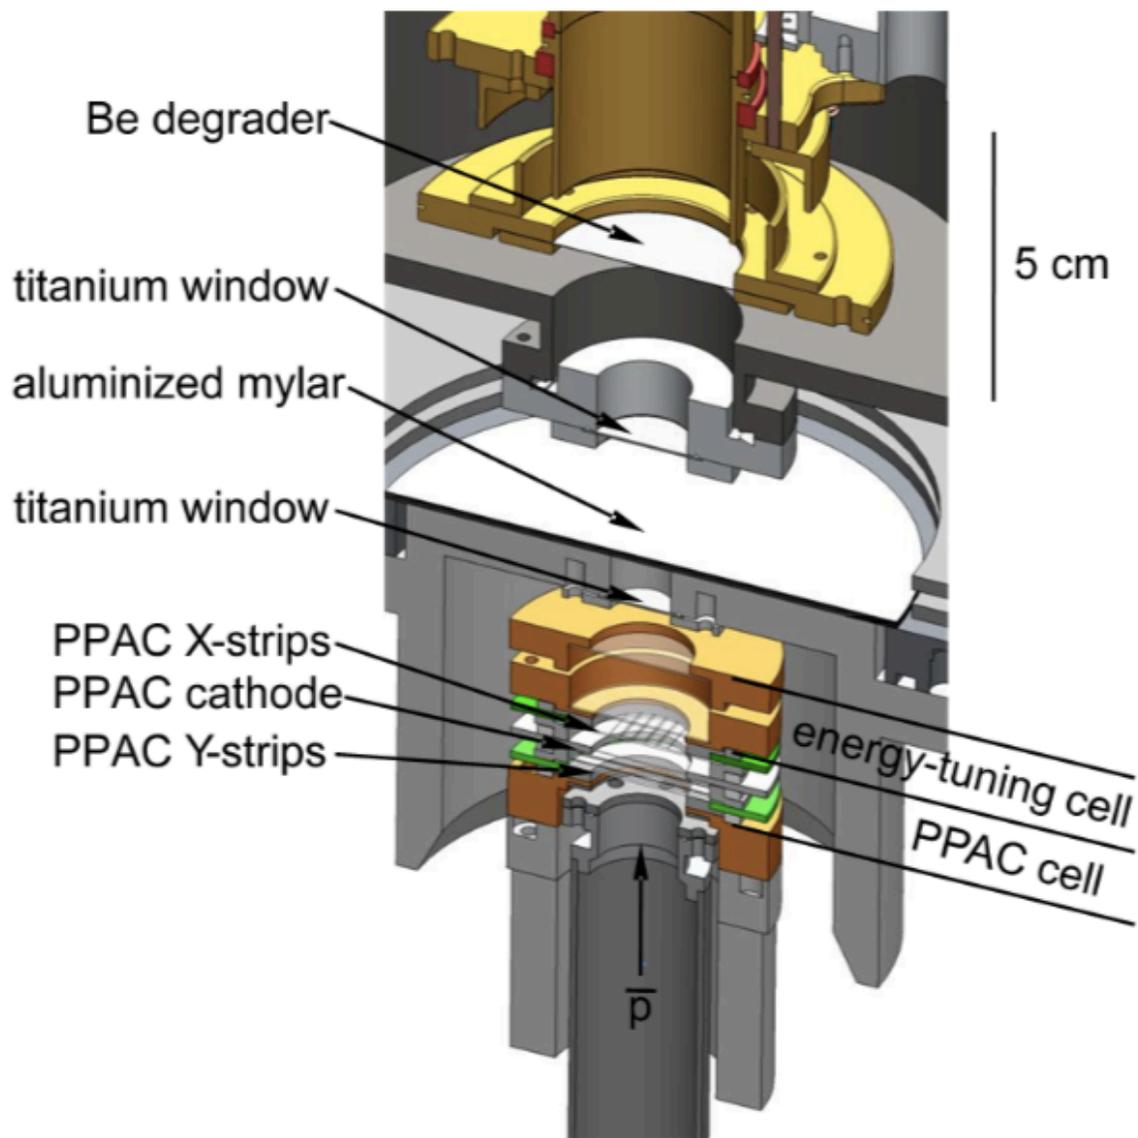


Figure 6.5: Steering PPACs and foil layers leading to CTRAP (adapted from [41])

to nearby strips, allowing us to determine the position of the beam. For a well-steered beam, the majority of the signal is seen on the central strip in each direction.

To slow the antiprotons, CTRAP is equipped with an energy tuning cell, two 10  $\mu\text{m}$  titanium windows, six sheets of 6.4  $\mu\text{m}$  aluminized mylar foil, and a 100  $\mu\text{m}$  beryllium degrader. Most of the slowing occurs in the degrader, whose thickness is chosen to be half the range of 5.3 MeV protons in beryllium [79]. The main purposes of the titanium windows and aluminized mylar are to separate spaces at different pressures and to insulate insert dewar layers at different temperatures. While these components don't contribute insignificantly to slowing antiprotons, they are chosen to be as thin as possible. Finally, the energy tuning cell is included for fine adjustments to antiproton energy loss. A mixture of He and SF<sub>6</sub> fill the cell, with a higher fraction of SF<sub>6</sub> corresponding to more slowing. This allows for the adjustment of incoming antiproton energies by up to 0.5 MeV. In the end, about 0.5% of antiprotons enter the trap with energies of less than 5 keV.

Prior to accepting antiprotons, the lower stack must be prepared for trapping. First, the magnetic field is locally increased to 3.7 T with the field-boosting solenoid surrounding the lower stack. This reduces the cyclotron radius of incoming antiprotons, leading to an increased trapping rate by a factor of  $\sim 5$  as fewer particles collide with trap walls. The electron cyclotron cooling rates also increase by a factor of 14, leading to quicker collisional cooling. Biases of -5 kV and 600 V are put on the HV electrode and the degrader, respectively. Between these, a  $10^8$ -electron cloud of radius 6 mm occupies a 100 V well.

Each ejection of an antiproton bunch from the AD is preceded by a trigger 4

$\mu\text{s}$  in advance. This trigger is used to precisely capture detector and PPAC output as antiprotons enter our trap. More critically, a custom-built switch changing the degrader's bias from 600 V to -5 kV performs the transition at a fixed offset from the warning trigger. The goal is to raise this potential barrier after most of the antiprotons have passed the degrader, but before many have reflected off the upper end of the well and returned to the degrader. Once the antiprotons have been captured, the -5 kV bias remains on the degrader for 60 s before it returns to 600 V to prepare for the next AD ejection. During this time, electrons cool the antiprotons into the 100 V well. With this procedure, we typically catch 150000 antiprotons in the 5 kV well, with 90% of these cooling into the 100 V well. In previous years, we have often accumulated  $10^6$  antiprotons. However, poor AD performance forced us to target 360000 antiprotons in 2018. In either case, once we have the desired number of antiprotons, the degrader and HV electrode return to ground and a beam block is inserted to prevent more ejections from entering the trap.

From this point, two goals need to be met in order to prepare antiprotons for antihydrogen synthesis. First, the cloud's radius must be reduced as much as possible. By doing so, we both avoid expansion into the walls when antiprotons are moved to the 1 T field of the upper stack and confine antihydrogen production close to the axis, where atomic potential energies are low. Second, electrons must be removed. Besides interfering with antihydrogen creation when antiprotons and positrons are merged, large numbers of electrons have been observed to centrifugally separate from mixed-in antiprotons, forcing the latter off-axis [65, 66].

We start by applying a pair of 100 ns pulses to the electrode holding the antiproton-

electron plasma. These are designed to temporarily bring the well depth to zero. Electrons and antiprotons have bounce periods of about 29 ns and 1.3  $\mu$ s, respectively, in the original well. Consequently, electrons are able to escape while antiprotons are left behind. The lack of detector hits during the process shows this to be the case. A ramp toward the degrader ensures that electrons leave the lower stack, and the charge-sensitive preamplifier counts the number ejected.

With about 15 million electrons remaining, the antiproton-electron plasma is spun for 400 s until it has a radius of about 0.5 mm. The factor of 4 reduction in size from the plasmas prepared in BTRAP [41] is a direct result of compressing with fewer electrons. Once the rotating wall is finished, most of the remaining electrons are pulsed out, leaving an estimated  $10^3$ - $10^4$  for cooling the antiprotons [45]. The antiprotons are then moved to the upper electrode stack, and the field-boosting solenoid is discharged.

## 6.6 Positron Loading

The positron loading procedure mirrors that for antiprotons, but it is quite a bit simpler. First, a 150 million electron plasma with radius 4 mm is prepared in the lower stack for slowing incoming particles. As an additional measure for slowing positrons, which arrive in the upper stack with kinetic energies between 60 eV and 70 eV, the bias on each electrode in CTRAP is increased by 58 V. On top of this shift, the top and bottom electrodes of the stack are set to 140 V to create a long well, and a pair of negative-voltage wells are constructed around the electrons.

As noted in a previous section, bunches of positrons are sent to CTRAP three times per AD cycle. Similar to the warning triggers from the AD for antiprotons,

the positron system sends a trigger prior to each positron shot. After a fixed delay, a negative-voltage pulse is applied to the top electrode. Again, the timing and length of the pulse is calibrated so that most incoming particles make it into the electrode stack and few are reflected back out. The admitted positrons are trapped in the long well, collisionally cooling with the electron cloud into the adjacent wells, until the next trigger arrives.

Once the desired number of positron shots has been collected, the 58 V offset is returned to zero and the barriers at either end of the electrode stack are removed. The electrons are also ramped to the degrader in a single motion that combines the contents of the two negative-voltage wells with minimal positron loss. Finally, a rotating wall is applied, leaving a plasma of radius 1 mm with 20 million positrons. The particles are moved to the upper stack for antihydrogen

It ought to be noted that positrons were caught in the upper stack during BTRAP's operation [41], allowing for simultaneous loading of antihydrogen's antimatter constituents. Because a reliable rotating wall scheme in CTRAP's upper stack has yet to be found, positrons must be compressed in the lower stack. Attempts have been made to trap antiprotons and positrons at the same time, then transfer positrons to the lower stack for compression, but these have so far always resulted in losses and/or plasma expansion. A future redesign to the upper stack rotating wall appears to be the only way to make concurrent loading feasible.

## 6.7 Antihydrogen Synthesis and Trapping

Once positron and antiproton clouds have been loaded, prepared, and moved to the upper stack, antihydrogen can be produced. Three principle mechanisms for creating the atoms have been well-explored in the literature. In each scheme, a particle carries away the energy and momentum needed to bind a free antiproton to a free positron. Radiative recombination takes this particle to be a photon [80], charge exchange uses an electron [81], and three-body recombination (TBR) ensures conservation with a scattered positron [82].

Radiative recombination has been demonstrated to be an unimportant production method, even when stimulated by laser light [83]. Charge exchange was first used to make antihydrogen in 2004 [11]. Improvements were later made to the process, and 3600 atoms were produced from 5 million antiprotons over 20 minutes [84]. In contrast, TBR can convert upwards of 40% of antiprotons to antihydrogen in seconds [2]. The vast majority of anti-atoms caught in magnetic traps have so far been produced by TBR, but developments in the other techniques may end up yielding more *trappable* antihydrogen.

Two methods for combining antiprotons and positrons have been used to create trapped antihydrogen. Early demonstrations produced atoms by exciting antiprotons into a cloud of positrons [1, 14]. A simulation studying different techniques for mixing the species later suggested that merging by potential manipulations would yield more trapped anti-atoms [85]. Improved trapping rates with this technique seem to bolster that claim [2].

In 2018, we implemented a similar technique for making and trapping antihydro-

gen. For each trial, antiproton and positron clouds that had been brought to the upper stack were placed in a nested well structure, centered near the octupole trap’s minimum. The Ioffe trap was then energized. Next, over 10 seconds, the antiproton well depth was lowered until it was no longer confining. Electric fields of 5 V/cm were then used to clear the trap of charged particles, including mirror-trapped antiprotons [86]. Finally, a trigger sent to discharge the trap also initiated recording of detector signals by the MCS, allowing time resolution for antihydrogen candidates.

Between the time taken by plasma studies, AD malfunction, and modifications to our antihydrogen synthesis routine, few full mixing trials were completed. However, the results of the nine we finished are shown in Fig. 6.6. (Trial-by-trial outcomes are given in Table 6.1.) Part (a) shows the sum over all trials of the detector counts in the first 50 ms after the atom trap was discharged, as well as those in the following 50 ms intervals. Part (b) shows simulated antihydrogen losses after trap discharged, justifying the assumption that the 0 to 50 ms bin holds anti-atom annihilations while the other two bins are almost exclusively background.

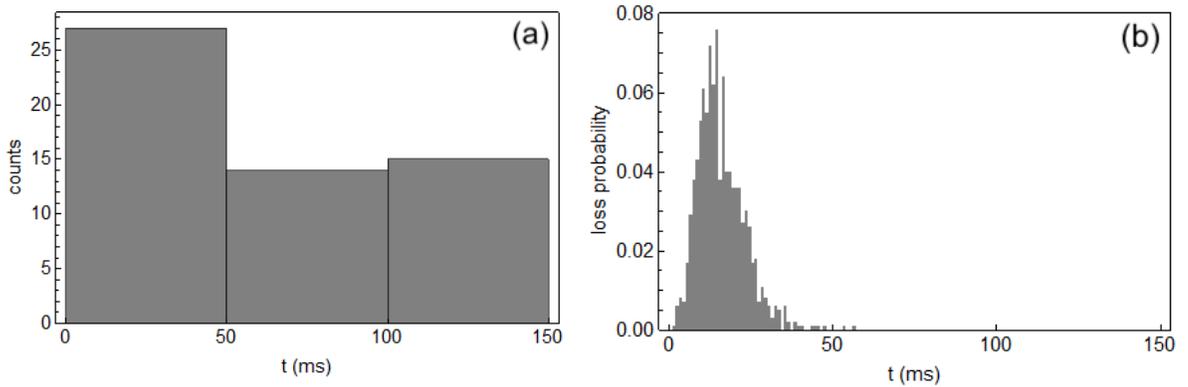


Figure 6.6: (a) Summed detector counts over antihydrogen trials and (b) simulated anti-atom losses during trap discharge

trial	counts		
	0 to 50 ms	50 to 100 ms	100 to 150 ms
1	2	3	2
2	4	0	3
3	5	2	2
4	2	0	0
5	3	2	4
6	6	0	0
7	3	3	0
8	0	3	1
9	2	1	3
sum	27	14	15

Table 6.1: Trial-by-trial detector counts are binned in 50 ms intervals after magnet discharged. Note that antihydrogen detection efficiency was about 26.7%.

In total, we observed  $5 \pm 2$  antihydrogen atoms per trial, with a probability of about 0.2% that the excess in the first 50 ms was due to background fluctuations. It's noteworthy that BTRAP trapped  $5 \pm 1$  atoms per antihydrogen trial, but with three times the number of antiprotons as we used in 2018 [1]. We also dealt with a signal-to-background ratio about 20 times lower than expected due to detector electronics issues. Resolving this problem should allow for faster iteration on antihydrogen synthesis schemes and better time resolution for determining when atoms leave the trap.

## 6.8 Summary

In this chapter, we discussed the methods used to produce and trap antihydrogen in 2018. To start, general techniques for both counting the number of particles in a plasma and understanding and manipulating a plasma's geometry were introduced. We then moved on to particle loading for electrons, antiprotons, and positrons. Finally, an antihydrogen synthesis routine was summarized. The 5 atom-per-trial yield we obtained using this routine was reported.

The first trapping of antihydrogen in the CTRAP apparatus represents a big step forward. The yield, however, is a bit disappointing. While the same number of anti-atoms were made per positron-antiproton mixing as in 2011, the number of antiprotons used was reduced by a factor of three. This indicates progress due to some combination of smaller plasma radii, a gentler merging procedure, and a deeper Ioffe trap. Additionally, the apparent size and amorphous shape of our antiproton clouds suggests that significant improvement can still be made. With better control over the geometry of these clouds, we may be able to soon produce large numbers of antihydrogen atoms, allowing for even better resolution of the 1S-2S line once spectroscopic studies commence.

# Chapter 7

## Atomic Theory

In order to simulate laser cooling and spectroscopy, it's necessary to work out the relevant atomic theory. We assume that hydrogen and antihydrogen are described by the same physics, save for the reverse signs of their constituents. In this chapter, we first study the lowest energy levels of the antihydrogen atom. It is especially important to account for small corrections to the 1S and 2S levels since these are the subjects of our intended precision measurement. We then move on to the 1S-2P transition, which we use for laser cooling. Lastly, the 1S-2S spectroscopy transition is explored.

### 7.1 Atomic States

The starting point of our analysis will be the nonrelativistic Hamiltonian

$$\begin{aligned} H &= H_0 - \vec{B} \cdot \vec{\mu} \\ &= H_0 - \vec{B} \cdot \left( \frac{g_L \mu_B}{\hbar} \vec{L} + \frac{g_e \mu_B}{\hbar} \vec{S} - \frac{g_p \mu_N}{\hbar} \vec{I} \right), \end{aligned} \tag{7.1}$$

where  $H_0$  is the Hamiltonian of an antihydrogen atom in its rest frame. The eigenstates of  $H_0$  are well-known, and each can be approximately described by a principal quantum number  $n$  and a hyperfine eigenstate (eigenstate of  $L$ ,  $S$ ,  $I$ ,  $J$ ,  $F$ , and  $m_F$ ). Since  $H$ 's field-dependent term does not couple states of different  $n$  and  $L$ , the problem of diagonalizing the operator is equivalent to that of diagonalizing it on each subspace of fixed  $n$  and  $L$ .

We concern ourselves with two transitions. Laser cooling works by driving the 1S-2P transition. Spectroscopy will probe the 1S-2S transition. While other states contribute to the physics we observe, especially in 2-photon spectroscopy, contributions from these states are folded into parameters like Rabi frequencies and ionization rates. No explicit computation of their properties is performed in this work. Consequently, we confine our study to the 1S, 2S, and 2P levels. In each case, we take  $H_0$  to be traceless on the subspace.

### 7.1.1 $L = 0$ States

The projection of  $H$  onto the subspace spanned by the  $nS$  states gives rise to a familiar eigenvalue problem, which is solved by the Breit-Rabi formula. We work in the  $|m_S, m_I\rangle$  basis, where  $\hat{z}$  is the direction of  $\vec{B}$ . Eigenvectors are labeled by  $|\psi_{s,i}^{(nS)}\rangle$ , where  $s = \lim_{B \rightarrow \infty} \text{sgn}(\langle S_z \rangle)$  and  $i = \lim_{B \rightarrow \infty} \text{sgn}(\langle I_z \rangle)$ . With the definitions

$$\mu_{\pm} = \frac{1}{2}(g_e \mu_B \pm g_p \mu_N) \quad (7.2)$$

$$\xi_{\pm} = \frac{2\mu_{\pm} B}{\Delta E_{HFS}} \quad (7.3)$$

$$\tan \theta = \frac{1}{\sqrt{1 + \xi_+^2 + \xi_-^2}} \quad (7.4)$$

we write the eigenvalues of  $H$  as

$$\frac{E_{\pm,\pm}^{(nS)}}{\Delta E_{HFS}} = \frac{1}{4} \mp \frac{\xi_-}{2} \quad (7.5)$$

$$\frac{E_{\pm,\mp}^{(nS)}}{\Delta E_{HFS}} = -\frac{1}{4} \mp \frac{1}{2} \sqrt{1 + \xi_+^2} \quad (7.6)$$

and the eigenstates as

$$\left| \psi_{\pm,\pm}^{(nS)} \right\rangle = \left| \pm \frac{1}{2}, \pm \frac{1}{2} \right\rangle \quad (7.7)$$

$$\left| \psi_{\pm,\mp}^{(nS)} \right\rangle = \cos \theta \left| \pm \frac{1}{2}, \mp \frac{1}{2} \right\rangle \mp \sin \theta \left| \mp \frac{1}{2}, \pm \frac{1}{2} \right\rangle. \quad (7.8)$$

It's worthwhile to pause here and take a look at the magnitudes of the various quantities for atoms in our trap. For the 1S states,  $\xi_{\pm} = (20 T^{-1}) B$  while for the 2S states,  $\xi_{\pm} = (158 T^{-1}) B$ , where we've set the precision so both signs round to the same coefficient. Since the fields in our trap are never less than 0.9 T,  $\xi_{\pm} \gg 1$  in all cases of interest. The energies are thus nearly linear with field and the energy eigenstates are approximately eigenstates of  $S_z$  and  $I_z$ . The linear coefficients defining energy (divided by Planck's constant) as a function of field are approximately  $\pm 14 GHz/T$ . Note that two states,  $\left| \psi_{-,\pm}^{nS} \right\rangle$ , have energies that increase with field, making them trappable in a magnetic minimum trap. The other two states,  $\left| \psi_{+,\pm}^{nS} \right\rangle$ , are not trappable in a magnetic minimum since their energies decrease with field. Intuitively, this makes sense since  $\left| \psi_{s_1, s_2}^{nS} \right\rangle \approx \left| \frac{s_1}{2}, \frac{s_2}{2} \right\rangle$  and the positrons's spin magnetic moment is about 660 times that of the antiproton's. These features can be seen in Fig. 7.1, where the 1S energies are plotted as a function of magnetic field.

The model in Eqn. 7.1 is approximately correct, but it clearly has shortcomings. Corrections accounting for the Pauli approximation, the classical treatment of electromagnetic fields, and the neglect of nuclear recoil and structure promise to sharpen our

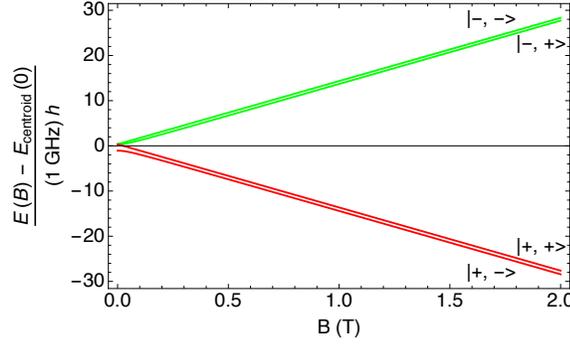


Figure 7.1: Energies of trappable (green) and untrappable (red) 1S states relative to the 1S zero-field centroid

energy estimates. Even so, corrections to nS states can be easily accounted for [87,88].

Eqns. 7.5-7.8 hold with two modifications. First, the g-factors take the values [3].

$$g_{\bar{e}} = g_{\bar{e}}^{free} \left( 1 - \frac{\alpha^2}{3n^2} + \frac{\alpha^4}{2n^3} \left( \frac{1}{2n} - \frac{2}{3} \right) + \frac{\alpha^3}{4\pi n^2} + \frac{\alpha^2}{2n^2} \frac{m_{\bar{e}}}{m_{\bar{p}}} \right) \quad (7.9)$$

$$g_{\bar{p}} = g_{\bar{p}}^{free} \left( 1 - \frac{\alpha^2}{3n^2} + \frac{\alpha^2}{6n^2} \frac{m_{\bar{e}}}{m_{\bar{p}}} \left( \frac{3 + 4a_{\bar{p}}}{1 + a_{\bar{p}}} \right) \right), \quad (7.10)$$

where  $a_{\bar{p}}$  is the anomaly. Second, each level is shifted by a diamagnetic term, quadratic in  $B$ .

$$\delta E_{diam,n} = \langle nS | \frac{e^2}{8m_{\bar{e}}} (\vec{B} \times \vec{r})^2 | nS \rangle \quad (7.11)$$

Let's try and understand the magnitudes of these corrections. The diamagnetic term is easy to calculate for  $n = 1$  and 2.

$$\delta E_{diam,1} = \frac{q^2 a_0^2}{4m_e} B^2 = (29772 \frac{Hz}{T^2}) hB^2 \quad (7.12)$$

$$\delta E_{diam,2} = \frac{7q^2 a_0^2}{2m_e} B^2 = (416818 \frac{Hz}{T^2}) hB^2 \quad (7.13)$$

To determine the effects of the g-factors, note that to three significant figures,

$$g_{1S} = g^{free} (1 - 1.77 * 10^{-5}) \quad (7.14)$$

$$g_{2S} = g^{free} (1 - 4.43 * 10^{-6}) \quad (7.15)$$

for both spin g-factors. (This agreement is spoiled at higher precision by differing higher-order terms.) Energy corrections may then be approximated by

$$\delta E^{(1S)}/h = (-248 \frac{kHz}{T})B \quad (7.16)$$

$$\delta E^{(2S)}/h = (-62 \frac{kHz}{T})B, \quad (7.17)$$

where  $\frac{\xi_+}{\sqrt{1+\xi_+^2}} \approx 1$  has been used.

For calculating particle trajectories and the 1S-2P transitions frequency, the 2S state does not matter. We expect to see fields between 0.9 T and 1.6 T, which yields no more than a 500 kHz shift. Kinematically, this corresponds to a 0.016 mK change in energy, which is negligible for even the cooled particles in our trap. Relative to the decay width of the 2P state, which is 99.58 MHz [44], 1S shifts are also unimportant.

On the other hand, 1S-2S spectroscopy is highly sensitive to these shifts. The most precise measurement of hydrogen's 1S-2S line has an uncertainty of 10 Hz [89] while the corresponding measurement for antihydrogen has an uncertainty of 5.4 kHz [3]. While the latter error may be optimistic given the complex environment of the atoms under study, it's worth noting that the width of observed line was 100 kHz, which is smaller than many of the shifts in Eqns. 7.12, 7.13, 7.16, and 7.17.

It's necessary to point out that even with the corrections above, we still may not have enough precision to determine antihydrogen's field-free 1S-2S frequencies to the same precision as those for hydrogen. Unless the Zeeman shift for each level is understood to better than a part in  $10^9$ , it's possible that yet uncalculated corrections or higher-order terms will push the theoretical uncertainty beyond the 10 Hz uncertainty for hydrogen. Possible resolutions include performing spectroscopy in a weaker field or performing spectroscopy on hydrogen in the same trap and comparing the results.

However, these considerations shouldn't noticeably affect spectral lineshapes. Thus, we will be content with the current level of precision for the simulations in this work.

The Zeeman effect not only moves the centers of lines observed during spectroscopy, it also expands their widths. For this reason, it's useful to get an idea of how 1S-2S transition frequencies vary near trap minima. Fig. 7.2 illustrates why we care about the minima. While peaks corresponding to the on-axis local maxima

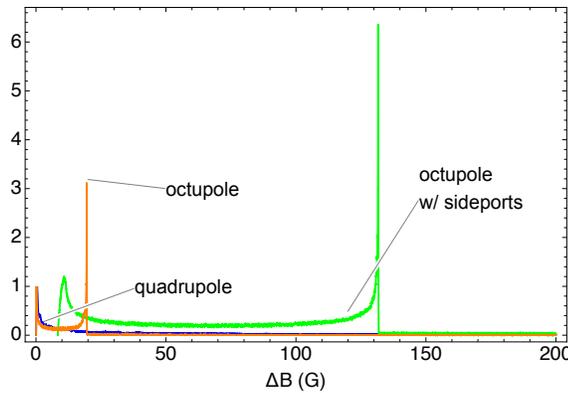


Figure 7.2: These histograms of fields seen by atoms when passing through a  $250 \mu\text{m}$  radial beam are taken with samples of 1000 uncooled atoms over 10 minutes. Bin heights are normalized so peaks near zero have a maximum value of 1.  $\Delta B$ 's zero is chosen to occur at the field minimum along the laser's axis in each trap.

of the octupole traps show up for the uncooled atomic samples used here, they disappear when atoms are cooled. On the other hand, the occupancy peaks near the trap minima are dependably present for all atomic distributions we consider.

At 1 T, the frequency of the  $|\psi_{-,-}^{1S}\rangle - |\psi_{-,-}^{2S}\rangle$  transition changes at about 961 kHz/T while the  $|\psi_{-,+}^{1S}\rangle - |\psi_{-,+}^{2S}\rangle$  transition goes at 18.6 MHz/T. While the difference between these rates shrinks with field, the latter is still more than 5 times the former at 1.6 T. It's clear that the line generated by the  $|\psi_{-,-}^{1S}\rangle - |\psi_{-,-}^{2S}\rangle$  transition is spread less by variation in the magnetic field. As a consequence, our focus going forward will be on

the  $|\psi_{-,-}^{1S}\rangle$  ground state.

### 7.1.2 L = 1 States

The 2P manifold is more complicated than the nS manifolds as the positron has nonzero orbital angular momentum. To start, let's ignore nuclear spin. The fine-structure splitting is denoted by  $\Delta E_{FS}$ . The following definitions are convenient.

$$\chi = \mu_B B / \Delta E_{FS} \quad (7.18)$$

$$\gamma_{\pm} = \sqrt{1 \mp \frac{2}{3}\chi + \chi^2} \quad (7.19)$$

$$\tan \theta_{\pm} = \frac{1 \mp 3\chi + 3\gamma_{\pm}}{2\sqrt{2}} \quad (7.20)$$

This time, the eigenvectors and eigenvalues are labeled by the fine-structure states  $|J, m_J\rangle$  they approach in the field-free limit. The energies are as follows.

$$\frac{E_{3/2, \pm 3/2}}{\Delta E_{FS}} = \frac{1}{3} \mp 2\chi \quad (7.21)$$

$$\frac{E_{3/2, \pm 1/2}}{\Delta E_{FS}} = -\frac{1}{6} \mp \frac{\chi}{2} + \frac{\gamma_{\pm}}{2} \quad (7.22)$$

$$\frac{E_{1/2, \pm 1/2}}{\Delta E_{FS}} = -\frac{1}{6} \mp \frac{\chi}{2} - \frac{\gamma_{\pm}}{2} \quad (7.23)$$

The corresponding eigenvectors, in the  $|m_L, m_S\rangle$  basis, are

$$|\psi_{3/2, \pm 3/2}\rangle = \left| \pm 1, \pm \frac{1}{2} \right\rangle \quad (7.24)$$

$$|\psi_{3/2, \pm 1/2}\rangle = \cos \theta_{\pm} \left| \pm 1, \mp \frac{1}{2} \right\rangle + \sin \theta_{\pm} \left| 0, \pm \frac{1}{2} \right\rangle \quad (7.25)$$

$$|\psi_{1/2, \pm 1/2}\rangle = -\sin \theta_{\pm} \left| \pm 1, \mp \frac{1}{2} \right\rangle + \cos \theta_{\pm} \left| 0, \pm \frac{1}{2} \right\rangle. \quad (7.26)$$

It's worth noting that  $\chi \approx (1.28T^{-1})B$ . For the fields present in our traps, 2P states are more a mix of their high-field limits than are the nS states.

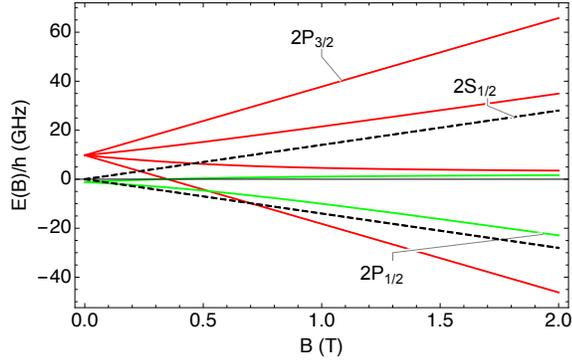


Figure 7.3: Energies of the  $2S_{1/2}$ ,  $2P_{1/2}$ , and  $2P_{3/2}$  states relative to the 2S zero-field centroid

Finally, we reintroduce the nuclear spin. Products of eigenstates in Eqns. 7.24-7.26 with eigenstates of the nuclear spin operator serve as an unperturbed basis. The hyperfine interaction, which splits energies at zero field by tens of MHz, then acts as a perturbation given the fields present in our trap. Fig. 7.3 makes the smallness of this perturbation obvious. Pairs of states that differ only by nuclear spin in the unperturbed limit appear to overlap. The 2S manifold is also shown to give a sense of its proximity for a later discussion on electric field induced 2S-2P mixing.

## 7.2 1S-2P Equations of Motion

### 7.2.1 Atom-Laser Interaction

We loosely follow the development of [44] to construct equations of motion for atoms driven by a laser tuned near a 1S-2P transition frequency. Under the electric

dipole approximation, the Schrödinger equation in the interaction picture is

$$\dot{c}_e(t) = \frac{i}{\hbar} \langle e | e\vec{r} \cdot \vec{E}(t) | g \rangle e^{i\omega_0 t} c_{e,g}(t) \quad (7.27)$$

$$\dot{c}_g(t) = \frac{i}{\hbar} \langle g | e\vec{r} \cdot \vec{E}(t) | e \rangle e^{-i\omega_0 t} c_e(t), \quad (7.28)$$

where  $\vec{E}(t)$  is the oscillating electric field of the laser radiation,  $\omega_0 = (E_e - E_g)/\hbar$ , and  $e > 0$  is the elementary charge. Simulations feature up to three beam paths at once, with each beam contributing a term to  $\vec{E}(t)$ . The frequency of each contribution, which is calculated in the atom's rest frame, depends on the atom's Doppler shift, which may differ between beams. Consequently, we write

$$\vec{E}(t) = \sum_{b \in \text{beams}} \frac{E_{0,b}}{2} \hat{\epsilon}_b (e^{i\omega_b t} + e^{-i\omega_b t}). \quad (7.29)$$

Here,  $E_{0,b}$  is the amplitude of the oscillating electric field due to beam  $b$ ,  $\epsilon_b$  is the polarization, and  $\omega_b$  is the frequency in the atomic rest frame. Thus, Eqns. 7.27-7.28 become

$$\dot{c}_e(t) = \frac{i}{\hbar} \sum_{b \in \text{beams}} \frac{eE_{0,b}}{2} \langle e | \vec{r} \cdot \hat{\epsilon}_b | g \rangle \left( e^{i(\omega_b + \omega_0)t} + e^{-i(\omega_b - \omega_0)t} \right) c_{e,g}(t) \quad (7.30)$$

$$\dot{c}_g(t) = \frac{i}{\hbar} \sum_{b \in \text{beams}} \frac{eE_{0,b}}{2} \langle g | \vec{r} \cdot \hat{\epsilon}_b | e \rangle \left( e^{i(\omega_b - \omega_0)t} + e^{-i(\omega_b + \omega_0)t} \right) c_e(t). \quad (7.31)$$

Applying the rotating wave approximation and defining

$$\Omega(t) = \sum_{b \in \text{beams}} \frac{eE_{0,b}}{\hbar} \langle e | \vec{r} \cdot \hat{\epsilon}_b | g \rangle e^{-i(\omega_b - \omega_0)t}, \quad (7.32)$$

our equations of motion simplify to

$$\dot{c}_e(t) = \frac{i}{2} \Omega(t) c_{e,g}(t) \quad (7.33)$$

$$\dot{c}_g(t) = \frac{i}{2} \bar{\Omega}(t) c_e(t). \quad (7.34)$$

It's convenient to write  $\Omega(t)$  in terms of laser intensities  $I_b$ .

$$\Omega(t) = \sum_{b \in \text{beams}} \frac{e\sqrt{2\mu_0 c I_b}}{\hbar} \langle e | \vec{r} \cdot \hat{\epsilon}_b | g \rangle e^{-i(\omega_b - \omega_0)t}, \quad (7.35)$$

## 7.2.2 Decay

In addition to the drive provided by the laser, spontaneous emission also affects the internal states of atoms. The decay rate from  $|e\rangle$  to  $|g\rangle$ , where a photon of polarization  $\hat{\epsilon}$  is emitted into a differential solid angle  $d\Omega$ , is given by

$$\Gamma_{e,g}(\hat{\epsilon})d\Omega = \frac{\alpha\omega_{e,g}^3}{2\pi c^2} |\langle e | \hat{\epsilon} \cdot \vec{r} | g \rangle|^2 d\Omega. \quad (7.36)$$

To understand the distribution of emissions, temporarily define  $\vec{v}_{e,g} = \langle e | \vec{r} | g \rangle$ . Then,

$$\begin{aligned} \sum_{i=1,2} |\langle e | \hat{\epsilon}_i \cdot \vec{r} | g \rangle|^2 &= \sum_{i=1,2} \vec{v}_{e,g}^T \hat{\epsilon}_i \hat{\epsilon}_i^\dagger \vec{v}_{e,g} \\ &= \vec{v}_{e,g}^T (1 - \hat{k} \hat{k}^T) \vec{v}_{e,g}, \end{aligned} \quad (7.37)$$

where  $\hat{\epsilon}_1$ ,  $\hat{\epsilon}_2$ , and  $\hat{k}$  are assumed to form an orthonormal basis.

It's necessary now to take an aside to think about the energy eigenstates. The 2P kets in Eqns. 7.24-7.26 are each eigenvectors of  $J_z$ . Before adding in hyperfine splitting, the basis of energy eigenstates consists of products of these states with eigenstates of  $I_z$ . Thus, elements of this basis are eigenvectors of  $F_z$ . Among eigenstates of  $F_z$ , the hyperfine term only couples those which agree in eigenvalue. Therefore, the energy eigenstates of the full Hamiltonian are also eigenstates of  $F_z$ .

Now, let's use the notation

$$|g\rangle = \psi_{1,0,0}(\vec{r}) |g_{SI}\rangle \quad (7.38)$$

to decompose ground states into their spatial and spin parts. In this section,  $\psi_{n,l,m_l}(\vec{r})$  refers to the eigenvector of the Coulomb Hamiltonian with principle quantum number  $n$  and orbital angular momentum described by  $l$  and  $m_l$ . Note that the spatial dependence is implicit in Eqns. 7.7-7.8 since it is the same for all states. Observe also that  $|g_{SI}\rangle$  forms a complete, orthonormal basis in the product space of positron and antiproton spins. Thus, an excited state may be written

$$|e\rangle = \sum_g c_{e,g} \psi_{2,1,m_L^{(e,g)}}(\vec{r}) |g_{SI}\rangle. \quad (7.39)$$

Since each  $|g_{SI}\rangle$  is an eigenstate of  $F_z$ , the following definitions make sense.

$$F_z |e\rangle = m_F^{(e)} |e\rangle \quad (7.40)$$

$$F_z |g_{SI}\rangle = m_F^{(g)} |g_{SI}\rangle \quad (7.41)$$

It's easy to see that

$$m_L^{(e,g)} = m_F^{(e)} - m_F^{(g)}. \quad (7.42)$$

In the notation just developed, the vector  $\vec{v}_{e,g}$  may be written

$$\vec{v}_{e,g} = \bar{c}_{e,g} \int dV \bar{\psi}_{2,1,m_L^{(e,g)}}(\vec{r}) \psi_{1,0,0}(\vec{r}) \vec{r} \quad (7.43)$$

Substituting the appropriate wavefunctions and integrating, we arrive at

$$\vec{v}_{e,g} = \bar{c}_{e,g} \frac{128}{243} a_0 (-(\hat{x} - i\hat{y})\delta_{1,m_L^{(e,g)}} + (\hat{x} + i\hat{y})\delta_{-1,m_L^{(e,g)}} + \sqrt{2}\hat{z}\delta_{0,m_L^{(e,g)}}), \quad (7.44)$$

where the coordinate system is constructed with  $\hat{z}$  parallel to the magnetic field and  $a_0$  is the Bohr radius.

Returning to the discussion of decays, substituting 7.44 into 7.37 yields

$$\sum_{i=1,2} |\langle e | \hat{\epsilon}_i \cdot \vec{r} | g \rangle|^2 = \frac{2^{14}}{3^{10}} |c_{e,g}|^2 a_0^2 \left[ (1 + \cos^2(\theta)) \delta_{|m_L^{(e,g)}|,1} + 2 \sin^2(\theta) \delta_{m_L^{(e,g)},0} \right], \quad (7.45)$$

where  $\theta$  is the angle between the local magnetic field and  $\hat{k}$ . Plugging this expression into 7.36 gives

$$\Gamma_{e,g}(\Omega)d\Omega = |c_{e,g}|^2 \frac{2^{13}}{3^{10}} \frac{\alpha a_0^2 \omega_{e,g}^3}{\pi c^2} \left[ (1 + \cos^2(\theta)) \delta_{|m_L^{(e,g)}|,1} + 2 \sin^2(\theta) \delta_{m_L^{(e,g)},0} \right] d\Omega. \quad (7.46)$$

For excited and ground states whose eigenvalues for  $F_z/\hbar$  differ by one, emission is biased to be parallel to the magnetic field. For those with equal eigenvalues, emitted photons tend to travel perpendicular to the field. These emission distributions are used to randomly choose emission angles during the simulation.

To find the rate of decay  $|e\rangle \rightarrow |g\rangle$ , Eqn. 7.46 must be integrated. Both angular terms integrate to  $16\pi/3$ . Thus

$$\Gamma_{e,g} = |c_{e,g}|^2 \frac{2^{17}}{3^{11}} \frac{\alpha a_0^2 \omega_{e,g}^3}{c^2}, \quad (7.47)$$

which deviates from the zero-field decay rate of  $\Gamma/2\pi = 99.6$  MHz only from the variations in  $\omega_{e,g}$  with field and the presence of  $c_{e,g}$ .

Finally, the total decay rate of  $|e\rangle$ ,  $\Gamma_e$ , must be calculated. This is the sum of  $\Gamma_{e,g}$  over the ground states.  $\Gamma_{e,g}$  depends on  $|g\rangle$  only through  $\omega_{e,g}$  and  $c_{e,g}$ . From Fig. 7.1 and 7.3, it's clear that  $\omega_{e,g}/2\pi$  varies with field by no more than 100 MHz, which corresponds to a fractional change of  $4 \times 10^{-8}$ . Since the sum over  $|c_{e,g}|^2$  is 1 (Eqn. 7.39)

$$\Gamma_e \approx \frac{2^{17}}{3^{11}} \frac{\alpha a_0^2 \omega_{1S-2P}^3}{c^2}, \quad (7.48)$$

where  $\hbar\omega_{1S-2P}$  is the energy difference between the 1S and 2P centroids. From Eqns. 7.47 and 7.48, the branching ratio of  $|e\rangle \rightarrow |g\rangle$  is  $|c_{e,g}|^2$ .

### 7.2.3 Target Transitions

There are two types of transitions  $|g\rangle \rightarrow |e\rangle$  we plan to drive. For laser cooling, we want transitions for which  $|e\rangle$  decays to  $|g\rangle$  with almost unit probability. Later, when we want to eject atoms from the trap, we will want transitions where  $|e\rangle$  has significant coupling to  $|g\rangle$  and a significant decay channel to an untrapped state. From Eqns. 7.44, 7.47, and 7.48 and the fact that there's minimal variation between different  $\omega_{e,g}$ , it's clear that  $\Gamma_{e,g}/\Gamma[e]$  is approximately proportional to  $|\langle e|\vec{r}|g\rangle|^2$  with a constant independent of  $|g\rangle$  and  $|e\rangle$ . Eq. 7.35 also shows that the Rabi frequency's depends on the connected states entirely through  $\langle e|\vec{r}|g\rangle$ . Thus, in addition to the probability for an excited state to decay to a specific ground state, a branching ratio gives information on the degree to which laser radiation can couple the states. This information is used to choose good transitions for cooling and atom ejection.

To find these transitions, consider Table 7.1. For the discussion in this subsection, 2P energy eigenstates are labeled by their hyperfine-free limits. The  $|\psi_{-, \pm}^{1S}\rangle \rightarrow |\psi_{-, -}^{2P}\rangle |\pm\frac{1}{2}\rangle$  transitions are almost perfectly closed, and are thus ideal for laser cooling. One might consider using the  $|\psi_{-, \pm}^{1S}\rangle \rightarrow |\psi_{0, -}^{2P}\rangle |\pm\frac{1}{2}\rangle$  transitions as well. Atoms in our traps can have velocities of up to  $\sim 100$  m/s. As each photon absorption or emission changes an atom's velocity by 3.3 m/s, laser cooling is expected to take tens of scatters. After ten  $|\psi_{-, \pm}^{1S}\rangle \rightarrow |\psi_{0, -}^{2P}\rangle |\pm\frac{1}{2}\rangle$  scatters, the atom has only a 48% chance of remaining trapped. Consequently, this (and any transition with a smaller branching rate to trapped states) will not work for laser cooling.

There appear to be several candidates for scattering an atom to an untrapped state. States  $|\psi_{0, -}^{2P}\rangle |\pm\frac{1}{2}\rangle$  and  $|\psi_{-, +}^{2P}\rangle |\pm\frac{1}{2}\rangle$  only couple to  $|\psi_{-, \pm}^{1S}\rangle$  through radiation

Excited State	BR( $ \psi_{-,-}^{1S}\rangle$ )	BR( $ \psi_{-,+}^{1S}\rangle$ )	BR(untrapped)
$ \psi_{-,-}^{2P}\rangle  -\frac{1}{2}\rangle$	1.00	0.00	0.00
$ \psi_{-,-}^{2P}\rangle  \frac{1}{2}\rangle$	0.00	1.00	0.00
$ \psi_{-,+}^{2P}\rangle  -\frac{1}{2}\rangle$	0.07	0.00	0.93
$ \psi_{-,+}^{2P}\rangle  \frac{1}{2}\rangle$	0.00	0.07	0.93
$ \psi_{0,-}^{2P}\rangle  -\frac{1}{2}\rangle$	0.93	0.00	0.07
$ \psi_{0,-}^{2P}\rangle  \frac{1}{2}\rangle$	0.00	0.93	0.07
$ \psi_{0,+}^{2P}\rangle  -\frac{1}{2}\rangle$	0.15	0.00	0.85
$ \psi_{0,+}^{2P}\rangle  \frac{1}{2}\rangle$	0.00	0.15	0.86
$ \psi_{+,-}^{2P}\rangle  -\frac{1}{2}\rangle$	0.85	0.00	0.15
$ \psi_{+,-}^{2P}\rangle  \frac{1}{2}\rangle$	0.00	0.85	0.15
$ \psi_{+,+}^{2P}\rangle  -\frac{1}{2}\rangle$	0.00	0.00	1.00
$ \psi_{+,+}^{2P}\rangle  \frac{1}{2}\rangle$	0.00	0.00	1.00

Table 7.1: Branching ratios from 2P states to 1S states at B = 1 T

polarized parallel to the magnetic field (Eqn. 7.44). In both our trap and ALPHA's, the field near the trap minimum is nearly parallel to the trap axis. Consequently, the transitions can only be appreciably excited by lasers traveling perpendicular to the trap axis. Additionally, our setup requires mirrors to send axially propagating cooling light along our trap's transverse directions, so any light entering our trap will be polarized transverse to the axis.

On the other hand, the  $|\psi_{-, \pm}^{1S}\rangle$  to  $|\psi_{0, +}^{2P}\rangle |\pm \frac{1}{2}\rangle$  or  $|\psi_{+, -}^{2P}\rangle |\pm \frac{1}{2}\rangle$  transitions couple to radiation transverse to the magnetic field. It is not obvious which transition will better transfer trapped states to untrapped states. We will address this in a later section.

Finally, a few brief comments are in order about whether we can think about the aforementioned transitions in isolation. Each curve in Fig. 7.4 consists of the four transitions involving a given  $|\psi^{2P}\rangle$ . Since transverse radiation can only couple states with  $\Delta m_F = \pm 1$  and each quartet of closely-spaced transitions involves two excited states differing in  $m_F$  by 1, the ground states each couple to only one excited state. The proximity of these four transitions is therefore only a worry when atoms in multiple trapped ground states are present. We may use microwaves, as in [3], to eliminate this problem by ridding the trap of atoms in unwanted states.

It is also worthwhile to note that the frequencies of the cooling transitions near the minima of the traps to be considered remain more than 14 GHz away from transitions to different  $|\psi^{2P}\rangle$  for the fields present in the traps. Thus, when we tune our laser to within a few hundred MHz of the cooling transition, Zeeman shifted to a trap's minimum, the distance to unwanted transitions will dwarf the 40 MHz laser linewidth

and the  $\sim 100$  MHz transition linewidth. Even with such a considerable detuning, it's possible that a large enough number of scatters might put the atoms being addressed into an unexpected ground state. On the other hand, a much closer approach between the transitions meant to clear the trap and other transitions may also be problematic. We will address both of these scenarios in the chapter on laser cooling simulation results.

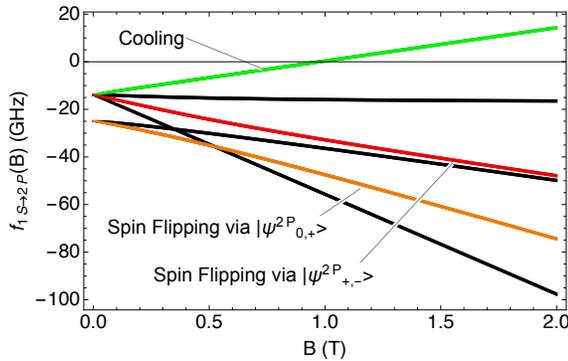


Figure 7.4:  $\bar{\text{H}}$  1S-2P transition frequencies with those used for cooling and emptying the trap (by spin flips) highlighted

## 7.2.4 Density Matrix

Before moving on, let's make a slight adjustment to the definitions of  $|g\rangle$  and  $|e\rangle$ . So far, these have only carried information about the atomic state. Treating the electromagnetic field as quantized, define the initial state of the electromagnetic field to be  $|\lambda_g\rangle$ . Let  $|\lambda_g\rangle$ , modified by the absorption of one photon from the laser, be  $|\lambda_e\rangle$ . In the following,  $|g\rangle$  is redefined to be the product of itself with  $|\lambda_g\rangle$ , and a similar redefinition is made for  $|e\rangle$ .

The effects of spontaneous emission may now be introduced to the model given by Eqns. 7.33-7.34. To do this, we switch our focus to the two-level density matrix.

The equations of motion may be written

$$\dot{\rho}_{gg} = \frac{i}{2}(\bar{\Omega}(t)\rho_{eg} - \Omega(t)\rho_{ge}) \quad (7.49)$$

$$\dot{\rho}_{ee} = -\Gamma\rho_{ee} - \frac{i}{2}(\bar{\Omega}(t)\rho_{eg} - \Omega(t)\rho_{ge}) \quad (7.50)$$

$$\dot{\rho}_{ge} = -\frac{\gamma}{2}\rho_{ge} + \frac{i\bar{\Omega}(t)}{2}(\rho_{ee} - \rho_{gg}). \quad (7.51)$$

Another way to represent these equations - in fact, the representation used in our simulations - is in terms of real variables. With the definitions

$$u = \rho_{ge} + \rho_{eg} \quad (7.52)$$

$$v = i\rho_{ge} - i\rho_{eg} \quad (7.53)$$

$$w = \rho_{ee} - \rho_{gg} \quad (7.54)$$

$$P = \rho_{ee} + \rho_{gg}, \quad (7.55)$$

Eqns. 7.49-7.51 become

$$\dot{u} = \Im(\Omega(t))w - \frac{\gamma}{2}u \quad (7.56)$$

$$\dot{v} = -\Re(\Omega(t))w - \frac{\gamma}{2}v \quad (7.57)$$

$$\dot{w} = \Re(\Omega(t))v - \Im(\Omega(t))u - \frac{\Gamma}{2}(w + P) \quad (7.58)$$

$$\dot{P} = -\frac{\Gamma}{2}(w + P). \quad (7.59)$$

Two unusual features of Eqns. 7.56-7.59 deserve explanation. First, the detuning between the frequency of the transition and the drive frequency is conspicuously absent. When the drive is harmonic, the detuning is extracted by redefining the phase of the excited state, which also makes  $\Omega$  constant. However, since  $\Omega(t)$  is not necessarily harmonic with more than one laser present (Eqn. 7.35), this cannot generally be done in a way that makes  $\Omega$  time-independent, so it is not done at all.

The second unusual feature is the decay of  $P$  in a system which appears to be closed. The complex Eqn. 7.49 usually has the term  $\Gamma\rho_{ee}$  to prevent this from happening. However, recall the redefinitions of  $|g\rangle$  and  $|e\rangle$  to include electromagnetic field states. Spontaneous emission couples  $|e\rangle$  to a continuum of states with the same atomic state as  $|g\rangle$ , differing only in the electromagnetic field state. While the decay rate from  $|e\rangle$  is finite, the decay rate to  $|g\rangle$ , with the field configuration  $|\lambda_g\rangle$ , is infinitesimal.

The quantity  $\gamma$  has not yet been defined. The lasers we use have finite linewidths. One way to introduce these into the density matrix equations relies on the phase-diffusion model [90]. Assuming the laser being used has an (angular) FWHM of  $\Gamma_{las}$ ,  $\gamma = \Gamma + \Gamma_{las}$ .

### 7.3 1S-2S Equations of Motion

The two-photon 1S-2S transition is handled differently from the single-photon 1S-2P transition. First, quantities like Rabi frequencies and decay rates are more complicated since they involve infinite sums. Instead of calculating these by hand, as was done in the previous section, results will be lifted from the literature. Second, the electric field generated by motion through the trap's magnetic field couple the 2S and 2P states. This leads to a slight shift in the energies of the 2S states, which can be important for precision spectroscopy. It also opens a one-photon decay channel, which may have a higher rate than the usual two-photon decay. Third, in addition to one and two-photon spontaneous emission, the 2S state suffers losses through single-photon ionization. Finally, momentum changes due to photon emissions are

neglected. This final point will be discussed in the chapter laying out results of spectroscopy simulations.

### 7.3.1 Equations of Motion and Two-Photon Quantities

From [91], we can write

$$\dot{\rho}_{gg} = \frac{i\Omega}{2}(\rho_{eg} - \rho_{ge}) \quad (7.60)$$

$$\dot{\rho}_{ee} = -\Gamma\rho_{ee} - \frac{i\Omega}{2}(\rho_{eg} - \rho_{ge}) \quad (7.61)$$

$$\dot{\rho}_{ge} = -\left(\frac{\gamma}{2} + i\delta\right)\rho_{ge} + \frac{i\Omega}{2}(\rho_{ee} - \rho_{gg}). \quad (7.62)$$

We assume the laser is introduced along a single beam path, allowing for the usual extraction of the angular detuning,  $\delta$ , and a time-independent Rabi frequency. The same comments as were made in the previous section hold regarding the absence of  $\Gamma\rho_{ee}$  in Eqn. 7.60 and the presence of the laser linewidth in  $\gamma$ . However, since the description is of a two-photon transition,  $\gamma = \Gamma + 2\Gamma_{las}$ .

For these equations, we still need a few parameters to be specified. To briefly summarize, [91] gives

$$\Omega_{1S-2S} = (9.252\text{e-}4 \text{ Hz}) \frac{I}{1\text{W}/\text{m}^2} \quad (7.63)$$

$$\Delta\nu_{1S-2S}^{\text{AC Stark}} = (3.334\text{e-}4 \text{ Hz}) \frac{I}{1\text{W}/\text{m}^2} \quad (7.64)$$

$$\Gamma_{ion}^{2S} = (1.511\text{e-}3 \text{ Hz}) \frac{I}{1\text{W}/\text{m}^2}. \quad (7.65)$$

On the other hand, two-photon decay is described in [92]. If  $\hbar\omega_0$  is the energy difference between the 1S and 2S states,  $\omega$  is the angular frequency of one of the emitted photons, and  $y = \omega/\omega_0$ , the distribution of emitted photon energies is  $y(1 -$

$y)\phi^2(y)$ . Here,  $\phi(y)$  is a function involving hypergeometric functions. The exact form of this expression is not very enlightening, but it is plotted in Fig. 7.5. The

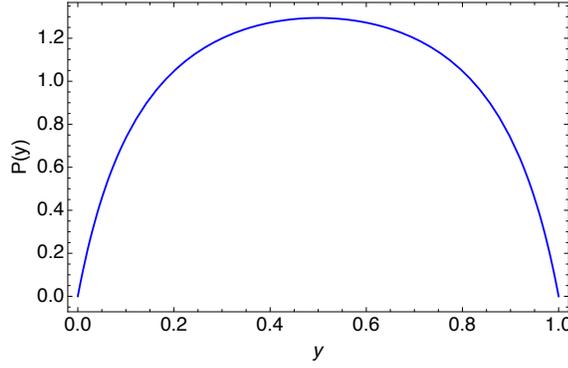


Figure 7.5: For two-photon decay, the distribution describing the fraction of the total emitted energy in one of the photons

distribution of angles between the two photons is much easier to describe, and is proportional to  $1 + \cos^2(\theta)$ . Finally, the decay rate is

$$\Gamma_{2-ph} = 8.2282Hz. \quad (7.66)$$

Two points may cause confusion. First, a precise definition of  $I$  is needed. We assume that atoms are excited by counter-propagating, identical Gaussian beams in a cavity.  $I$  is the time-averaged intensity of one of those beams at the location of the atom in question. The factor of two discrepancy between Eqns. 7.63-7.65 and the values listed in [91] is caused by two beams driving the transition rather than one.

Second, despite the earlier emphasis on the Zeeman effect, the quantities above do not depend on  $B$  at all. This is a simplification, but one that seems justified. Each quantity described by Eqns. 7.63-7.65 can be expressed as a linear combination of terms of the form

$$\langle e | r^i \frac{1}{E_g + \hbar\omega_L - H_0} r^j | g \rangle = \sum_{|\phi\rangle} \langle e | r^i | \phi \rangle \frac{1}{E_g + \hbar\omega_L - E_\phi} \langle \phi | r^j | g \rangle \quad (7.67)$$

or

$$\langle e | r^i \frac{1}{E_e \pm \hbar\omega_L - H_0} r^j | e \rangle = \sum_{|\phi\rangle} \langle e | r^i | \phi \rangle \frac{1}{E_e \pm \hbar\omega_L - E_\phi} \langle \phi | r^j | e \rangle, \quad (7.68)$$

where  $\omega_L$  is the frequency of the laser,  $H_0$  is the Hamiltonian in the absence of laser radiation, the sums on the right hand sides run over the complete set of eigenstates of  $H_0$ , and the indices  $i$  and  $j$  are left unspecified. Since  $|g\rangle$  and  $|e\rangle$  are spherically symmetric, only the nP terms in the above sums survive. Now, 1S and 2S energies change by about 14 GHz/T while nP energies change, at most, by around 28 GHz/T. Since  $2\hbar\omega_L$  is approximately the energy difference between the  $n = 1$  and  $n = 2$  levels, the magnitudes of denominators of surviving terms in Eqn. 7.67 are bounded below by about  $\hbar\omega = (1.2 \cdot 10^{15} \text{ Hz})h$ . Thus, the Zeeman effect negligibly effects this sum.

Similar statements hold for Eqn. 7.68 with a negative sign. With a positive sign, the denominator passes through zero at a positive energy of  $(4.1 \cdot 10^{14} \text{ Hz})h$ . Energy variations due to the Zeeman effect don't greatly affect the position of this relative to the lower limit of the continuum states, so the field-free version of Eqn. 7.68 should suffice.

From [93], two-photon decay rates can be expressed as sums of terms proportional to

$$\langle g | r^i \frac{1}{H_0 + \hbar\omega_\gamma - E_e} r^j | e \rangle = \sum_{|\phi\rangle} \langle g | r^i | \phi \rangle \frac{1}{E_\phi + \hbar\omega_\gamma - E_e} \langle \phi | r^j | e \rangle, \quad (7.69)$$

where  $\omega_\gamma$  is the frequency of one of the emitted photons. Zeeman shifts can be important when  $|\phi\rangle$  is a 2P level and  $\omega_\gamma/2\pi$  is tens of GHz. However, the probability for such a decay to occur is vanishingly small (Fig. 7.5). Quantitatively, the probability that a two-photon decay involves a photon of frequency 1 THz or less is about

$2 \cdot 10^{-12}$ . For intermediate nP states with  $n > 2$ , the minimum value of the denominator in Eqn. 7.69 is  $(4.6 \cdot 10^{14} \text{ Hz})h$ , so again Zeeman shifts don't matter.

Finally, it's important to understand the spins of states coupled in the above expressions. In each case, we can write

$$\begin{aligned} \langle \psi_a | r^i \frac{1}{\lambda - H_0} r^j | \psi_b \rangle &= \sum_{|\phi\rangle} \langle \psi_a | r^i \frac{1}{\lambda - H_0} |\phi\rangle \langle \phi | r^j | \psi_b \rangle \\ &= \sum_{n \geq 2} \langle \psi_a | r^i \sum_s \left( \frac{1}{\lambda - E_{nP_s}} |nP_s\rangle \langle nP_s| \right) r^j | \psi_b \rangle, \quad (7.70) \end{aligned}$$

where the sum with index  $s$  runs over the energy eigenstates in the nP manifold. The denominator  $(\lambda - E_{nP_s})$  has been shown to be much larger than the tens of GHz split between states of different  $s$ , except in rare two-photon emissions where one photon's frequency is less than a THz. Therefore, we can approximate the sum in Eqn. 7.70 as

$$\sum_s \left( \frac{1}{\lambda - E_{nP_s}} |nP_s\rangle \langle nP_s| \right) = \frac{1}{\lambda - E_{nP}} \left( \sum_{m_L} |nP_{m_L}\rangle \langle nP_{m_L}| \right) \otimes I_{spins}. \quad (7.71)$$

Here,  $|nP_{m_L}\rangle$  represents the spatial wavefunction  $\psi_{n,1,m_L}(\vec{r})$ ,  $E_{nP}$  is the approximate energy of the nP states, and  $I_{spins}$  is the identity matrix in the space of positron and antiproton spins. Since 1S and 2S states can be factored into a wavefunction  $\psi_{n,0,0}(\vec{r})$  and a vector in the space of positron and antiproton spin states, the quantities in Eqns. 7.63-7.65 and 7.66 are multiplied by the inner product of  $|\psi_a\rangle$  and  $|\psi_b\rangle$ 's spin components. For the states of particular interest,  $|\psi_{-,-}^{(1S)}\rangle$  and  $|\psi_{-,-}^{(2S)}\rangle$ , the spin components are exactly  $|\frac{-1}{2}, \frac{-1}{2}\rangle$  in the  $|m_S, m_I\rangle$  basis, making the additional factor exactly 1. This also means that a two-photon decay from  $|\psi_{-,-}^{(2S)}\rangle$  always results in an atom in the state  $|\psi_{-,-}^{(1S)}\rangle$ .

### 7.3.2 Stark Effects

As they travel through the trap, atoms experience small electric fields varying much more slowly than the optical frequencies of interest. These fields have two important consequences. First, the 1S-2S transition frequencies are shifted by small amounts. Second, the metastable ( $\tau = 0.12$  s) 2S states are coupled to the short-lived ( $\tau = 1.6$  ns) 2P states. This coupling decreases the lifetime of the 2S states while also increasing the probability of decays to untrapped states.

Patch potentials represent one source of slow-varying electric fields. Unfortunately, these fields are difficult to predict and may change from one spectroscopy experiment to another. Consequently, their effect is not included in our simulations.

An electric field is also generated in the rest frame of an atom moving through a magnetic field.

$$\vec{E}_{motion} = \vec{v} \times \vec{B} \quad (7.72)$$

The Lorentz factor has been dropped as it is negligibly different from one for trappable velocities. For the velocities ( $v < 100$  m/s) and magnetic fields ( $0.9$  T  $< B < 2$  T) in our trap, it can be shown that the interaction between the atom and the motional electric fields can be treated perturbatively.

#### Frequency Shift

To lowest order, a small electric field shifts the energy of  $|\psi\rangle$  by approximately

$$\Delta E_\psi = -\frac{1}{2}\alpha_\psi \vec{E}^2, \quad (7.73)$$

where  $\alpha_\psi$  is the polarizability of state  $\psi$ . The value of  $\alpha_\psi$  for a 1S state is  $\frac{9}{2}(4\pi\epsilon_0)a_0^3$  [94]. This leads to an energy shift of  $5.6 \times 10^{-8}$  Hz/(V/m)<sup>2</sup>. This shift is not important

for the expected electric fields in our trap.

The 2S manifold is more strongly affected by electric fields due to the nearby 2P states. Zeeman shifts appreciably change the energy differences between these sets of states, so the polarizability is sensitive to  $B$ . For the 2S state of interest, the

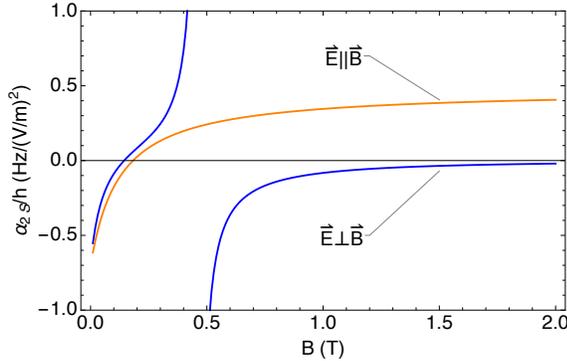


Figure 7.6:  $|\psi_{-,-}^{2S}\rangle$  polarizability as a function of field

polarizability as a function of field is shown in Fig. 7.6. We see that motional fields can give rise to shifts of a few hundred Hz near the trap minimum ( $B \sim 1$  T). The sign change near 0.5 T is due to an energy crossing between the 2S and 2P states. In this region,  $|\alpha_{2S}/h|$  never exceeds  $10 \text{ Hz}/(\text{V}/\text{m})^2$ .

### Decay Effects

Nearby 2P states also affect the decay rates and branching ratios of the 2S manifold. As the decay rates of 2P states exceed those of 2S states by more than seven orders of magnitude, only a tiny electric field is needed for single-photon decay channels to dominate the usual two-photon channel. This not only boosts the overall decay rate. It also opens previously absent electric dipole transitions to untrapped states. The effects for motional electric fields are shown in Fig. 7.7. Again we see the

effect of the 2S-2P crossing near 0.5 T.

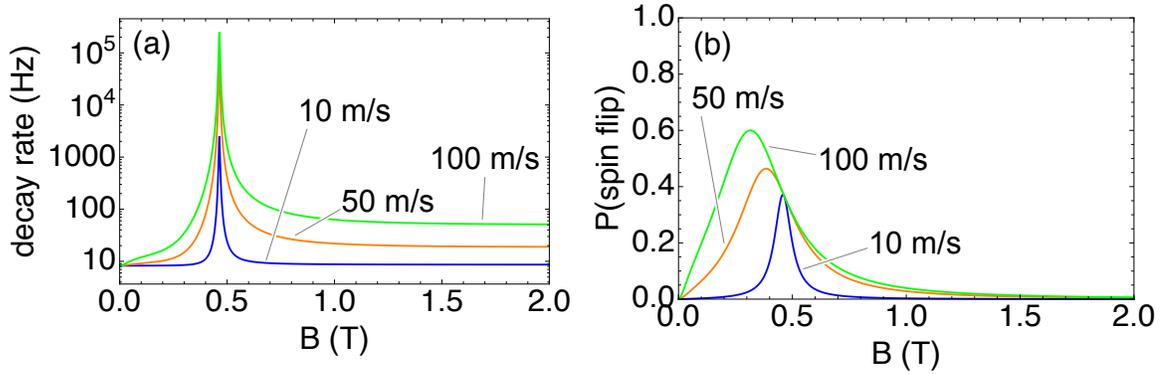


Figure 7.7: Decay rates and spin flip probabilities for  $|\psi_{-,-}^{2S}\rangle$  due to motional electric fields for various  $\bar{H}$  speeds.

## 7.4 Summary

In this chapter, we laid much of the theoretical groundwork needed to simulate the laser manipulations we intend to perform on antihydrogen. State energies, Rabi frequencies, decay rates, and other quantities calculated here appear in the simulations described in Ch. 8 and whose results are given in Chs. 9 and 10. We now move on to these chapters, reaping experimentally actionable insights from the basic atomic physics discussed here.

# Chapter 8

## Simulation Code

The complicated structure of our magnetic traps makes it difficult to calculate antihydrogen trajectories analytically. The situation is made even harder when lasers are introduced. Transition frequencies, decay rates, and Rabi frequencies depend on the trap and laser fields in an atom's vicinity, which change dramatically as it propagates. Adding to these issues is the inseparability of an atom's internal and kinematic variables during photon absorption and emission.

To tackle these problems, I developed and ran software on Harvard's Odyssey cluster to simulate antihydrogen in CTRAP. The programs in this package, written in C++, allow a user to track the internal and kinematic variable of atoms during propagation through a given trap environment. Trap fields are allowed to change, and one or more cooling or spectroscopy lasers may be added. Using these tools, I was able to show how atoms leave the trap during magnet dumps, optimize our cooling scheme, and understand CTRAP's spectroscopic abilities.

In this chapter, I begin by describing the foundational libraries on which the rest

of the code base is built. Next, I explain the methods used to calculate trap fields, derivatives, and other derived quantities. Following this, I detail how kinematic and internal degrees of freedom are evolved in the presence of trap fields and lasers. I then summarize the simulations I've built using the aforementioned building blocks. Finally, I lay out the code used to efficiently manage a large number of simulations in a cluster environment. The modularity of my design makes it simple to create programs addressing new situations and problems.

## 8.1 Foundations

### 8.1.1 Filesystem Operations

All filesystem calls are performed by *fstream* member functions or by Boost's Filesystem, Iostreams, or Interprocess library. The class *fstream* is the natural choice for managing reads and writes in C++. On the other hand, the Boost libraries replace hundreds of lines of Unix-specific custom code for managing directory structures, memory maps, and file locks with a few functions that work cross-platform.

As useful as these functions are on their own, almost all filesystem calls are wrapped by functions appearing in `File_IO.*`. (From here on, `filename.*` refers collectively to `filename.h`, `filename.hpp`, and `filename.cpp`.) The following list gives an idea of the functionality introduced by `File_IO.*`.

- `data_len` gives the number of numbers in a file.
- `print_to_file` prints the numbers following a pointer to a file.

- `read_file` loads a subset of a file.
- `map_file` constructs a memory mapping of a file.
- Several functions take a few parameters as input and return a standardized file name.

It's worth noting that for reading, writing, and determining file lengths, a file's contents can be interpreted as plain text or binary. The decision on how to interpret the file is made by examining the extension. File names ending in ".txt" are interpreted as plain text while those ending in ".bin" are assumed to be binary. This allows the user the flexibility to specify input and output formats without recompiling code.

### 8.1.2 Data Arrays

Sequential data is used in many parts of the simulation. It represents flattened 3D grids of trap quantities, arrays of changing currents, sequences of atomic states, and more. While the standard library provides array and vector types, neither are sufficient for this project. Array sizes are resolved at compile time, which is unacceptable when the size needed is passed to the program as an option. Vector sizes can be changed at run time, but it's difficult to control the amount of underlying memory allocated. Neither container is compatible with memory mapping to handle large files. To offer the flexibility needed, I developed four general use classes.

The first class, `mem_arr`, allocates a fixed amount of storage upon construction. It has the following properties.

- Upon construction, a single value or a subset of a file may be chosen to fill the allocated memory. In the latter case, the amount of memory to reserve is determined from the file segment's size.
- Unary and binary operators corresponding to addition, subtraction, multiplication, and division are overloaded to allow for the composition of a *mem\_arr* with either a scalar or another equal-sized *mem\_arr*.
- The square root function is overloaded to perform a component-wise square root on a *mem\_arr*.
- The equality and inequality operators are overloaded to compare two equal-sized *mem\_arrs*. Each returns a single boolean specifying whether or not two *mem\_arrs* are exactly equal.
- A function is provided to print the contents of a *mem\_arr* to file.

This class is particularly useful for loading small ( $< 100$  MB) arrays and doing algebra on segments of large arrays. The class used to describe file segments is discussed in a different section.

Rather than allocating memory to hold an array, the second class, *mapped\_arr*, sets up a memory mapping for a file. The memory map is read-only, so no mathematical operations can be performed on the array. However, comparison with other *mapped\_arrs* and *mem\_arrs* is allowed. As it would amount to an inefficient file copy, writing *mapped\_arrs* to file is also prohibited.

A special feature belonging to *mapped\_arr* deserves special mention. Instead of allowing an instance to directly map a file, an argument may be passed to the

constructor so that the target file is copied and the copy is mapped. As several processes may simultaneously try to copy a single file to the same location, I defined a class to manage the copied file (*tmp\_file\_mgr*). An instance of this works with a file lock and a log file to ensure that modifications to the copied file are done atomically. It also ensures that the copy is deleted when all processes have finished working with the file.

This temporary copy feature may seem inessential, but it actually gives a significant boost to the rate at which simulations run on Harvard's cluster. Laser simulations each require about 40 GB in grid files. The numbers stored in these files correspond to points on a flattened 3D grid, so page faults are common as particles move about. Since Odyssey spans multiple data centers separated by 10s of miles, the added memory latency slows simulations by orders of magnitude. We could try requesting large blocks of RAM and loading the files into *mem\_arrs*, but most nodes on the cluster provide 4 GB of RAM per core. High-memory jobs would be scheduled less frequently and would quickly diminish future scheduling priority. By copying the files to a node's local drive, which generally has a size of 250 GB, we cut memory latency without reserving large amounts of RAM.

The remaining two classes are triples of the first two. The classes *mem\_arr\_vec* and *mapped\_arr\_vec* are simply three-dimensional standard library arrays of *mem\_arrs* and *mapped\_arrs*, respectively. While the behavior of their constituents is explained above, a few features are added to take advantage of these containers' three-dimensional structures.

- The file name passed to read and write functions does not correspond to that

of any actual file being addressed. Instead, the name is modified to give a directory name and the names of three files in this directory. For instance, if the input file name is " /Desktop/f.bin", the three components would be read from or written to "Desktop/f\_vec/f\_x.bin", "Desktop/f\_vec/f\_y.bin", and "Desktop/f\_vec/f\_z.bin".

- Arithmetic operators are overloaded to compose *mem\_arr\_vecs* with other *mem\_arr\_vecs*, *mem\_arrs*, three-dimensional arrays, and scalars. Objects that are composed are required to have compatible sizes, and any operation between objects with different dimensions is defined in the obvious way.
- Equality and inequality operators are overloaded so two containers of the same size are equal if and only if each of their three components is equal.
- Vector algebra is defined so inner products and cross products can be taken between *mem\_arr\_vecs* of the same size as well as between a *mem\_arr\_vec* and a three-dimensional array. Functions are also provided for *mem\_arr\_vecs* to find the norm as well as to normalize the array. Note that each of these operations is performed viewing a *mem\_arr\_vec* as an array of three-dimensional vectors, so the result is either another *mem\_arr\_vec* or a *mem\_arr*.

Finally, it's important to mention that each of these classes comes equipped with member functions common to many standard library containers. These functions include `size()`, `data()`, `(c)begin()`, `(c)end()`, and `operator[]()`. In addition to exposing an object's underlying data and size, these allow for iteration over its elements and compatibility with many of the algorithms in the standard library.

## 8.2 Grids of Physical Quantities

### 8.2.1 Coordinate Grids

We often want to work with a collection of points in three-dimensional space. There are two main problems that arise when working with these ensembles. First, we want to have a way of enumerating these points. We may want to perform a calculation for each element in the set, or there may be an array storing physical quantities associated with each point. In either case, it is convenient to be able to map some range of indices to the collection of points of interest. The second problem involves mapping *arbitrary* spatial positions to nearby points in our selected ensemble, or equivalently, to the indices corresponding to these points. This is necessary when performing interpolations, for instance. Two classes are defined to solve these problems.

The more important class, `cart_coord_grid`, represents a cartesian grid with uniform spacing along each dimension. Its constructor loads the grid's limits in each dimension, the number of points along each dimension, and the grid's radial and axial boundaries. Specification of the boundaries is necessary because the space being represented is often cylindrical, forcing a spanning cartesian grid to possess invalid points. Two member functions establish the maps from three-dimensional position space to a one-dimensional index and back. The map yielding an index can be written as

$$idx_{1D}(\vec{r}) = \sum_{k=0}^2 s_k \left[ \frac{r_k - r_k^{min}}{dr_k} \right], \quad (8.1)$$

where  $r_k^{min}$  is the grid's minimum  $k^{th}$  coordinate,  $dr_k$  is the spacing along the  $k^{th}$

axis, and  $s_k = \prod_{l < k} n_l$  with  $n_l$  the number of points along grid axis  $l$ . It's worth pointing out that the zero-based indexing common in C++ and other languages is being used rather than the one-based indexing preferred in physics. The opposite function's definition is a bit involved, but the operation it performs can be easily understood. A valid grid index can be written uniquely in the form

$$idx = \sum_{k=0}^2 s_k i_k(idx), \quad (8.2)$$

where  $i_k(idx) < n_k$ . The three-dimensional point corresponding to index  $idx$  then has coordinates

$$r_k(idx) = r_k^{min} + i_k(idx) dr_k. \quad (8.3)$$

Given an algorithm for calculating  $i_k(idx)$ , we have a map from a one-dimensional index to the grid points.

The simpler and less-used class `file_coord_grid` represents an arbitrary assortment of points in three-dimensional space. This is used, for instance, when checking the differences between a field's actual and interpolated values at random points in the trap. The constructor for this class simply maps the contents of a specified file. Points are then given by contiguous triples of numbers in the file. In other words, a valid  $idx$  maps to the three-vector with components equal to the entries  $3 * idx$ ,  $3 * idx + 1$ , and  $3 * idx + 2$  in the mapped file. Given the unstructured nature of the grids described by `file_coord_grid` objects, there's no member mapping three-dimensional points to grid indices.

In addition to the features already highlighted, each coordinate grid class also has a few basic functionalities. The `grid_size()` member returns the number of grid points described by an object. The member function `out_of_bds()` is overloaded to

decide whether an input grid index or coordinate triple represents a point beyond the boundary of an object's grid. Finally, equality and inequality operators are overloaded to compare coordinate grid objects belonging to the same class. This is important, for instance, when adding two fields, each specified at every point of a coordinate grid, to make sure corresponding indices refer to the same point in space.

## 8.2.2 Data Grids

To work with quantities defined on a coordinate grid, it's useful to bind the array describing the quantity values to one of the coordinate grid objects outlined in the previous section. In fact, a slightly more complicated class is even more convenient.

The template class *grid\_quantity\_template* is defined in `Grid_Quantity_Template.*`. It contains four underlying types.

- *In\_Params* is a class that holds information on the source specifying the coordinate grid, instructions for creating the array of grid quantities, and some other information that is useful to package with the grid.
- *Coord\_Grid* is either *cart\_coord\_grid* or *file\_coord\_grid*.
- *Data* is a container holding the array of quantities. It is also required to have members functions for loading its contents from a file and printing them to file.
- *Out\_Params* is a class specifying the file(s) to which the *Data* object's contents are written.

Note that each of these types is a template parameter, not a concrete class. The exact requirements for the types are specified in `Grid_Quantity_Template.h`.

One benefit to having such a class is the uniform interface it provides when dealing with different kinds of quantities on a grid. This interface includes the following operations.

- `fill_contents()` takes a file name and two numbers. The file initializes both the *In\_Params* and *Out\_Params*. The coordinate grid definition is then loaded from either the same file or from a file identified in the *In\_Params* member. The information now present in the *grid\_quantity\_template* is used to initialize the *Data* member. The two input numbers specify the segment of the coordinate grid handled by this object. The classes used to import parameters from a file and pick a segment of a coordinate grid are discussed in later sections.
- `write_opts_file_template()` writes an options file, with parameters unspecified, to the input file name. This function allows a user to properly set parameters without looking through the source code for guidance.
- `print_contents()` writes the contents of the *Data* member to the file path(s) given by *Out\_Params*.
- `print_out_mem()` gives the disk space taken by a call to `print_contents()` from a *grid\_quantity\_template* initialized by a certain options file. Arrays specifying the values of physical quantities on a grid are often several GB in size, so it's worthwhile to let the user easily discover the memory commitment being made.
- `output_files()` returns a vector containing the file paths created by a call to `print_contents()`. This is used to check whether a given instance of *grid\_quantity\_template* has successfully constructed its *Data* array and printed it to file.

While it may seem tedious to specify four classes for each instantiation of *grid\_quantity\_template*, a few facts are worth noting. First, the coordinate grids are already defined, so only three additional classes are required. However, as can be seen in `Common_Grid_Quantities.h`, the reuse of certain parameter classes and templated versions of data classes can dramatically reduce the number of new classes one must define. Another important observation is that we avoid repeating the code implementing the interface, making modification and error checking far easier. Finally, by constructing the same interface for many data types, the use of these classes becomes simpler.

### 8.2.3 Field Definitions

We've seen that CTRAP's Ioffe trap features four different coils. The two mirror coils may be described approximately as collections of current loops positioned symmetrically about the trap center. The two Ioffe coils, which provide radial confinement, are more complicated. Upon the fields generated by these coils is superposed the Penning trap's uniform background field. Now, let's discuss how fields and field gradients are calculated.

Each Ioffe coil is accompanied by a winding definition file. These text files are lists of triples of coordinates in the order the winding passes through them. Unfortunately, these definitions need some modification to be useful. `Fix_Coil_Definitions` performs two tasks. First, it reads each winding file, performs a linear transformation on each triple, and writes the result to file in binary form. The linear transformation puts the winding into the natural coordinate system of the trap while also changing units from millimeters to meters. Second, after the modified quadrupole and octupole

definitions have been written to file, `Fix_Coil_Definitions` load the transformed octupole winding, removes the sideport-accommodating bends, and writes the result to file. Consequently, we have, in the correct orientation and units, binary files holding winding data for the quadrupole coil, the octupole coil, and a straightened version of the octupole coil. The straightened coil is useful for separating idiosyncrasies of our octupole trap from general properties of a more ideal octupole trap.

With concrete coil models, we can use Biot-Savart's law to calculate the resulting fields and field gradients. For the Ioffe coils, we integrate over the individual winding segments to get the following.

$$\begin{aligned}
 \frac{1}{I_0} \vec{B} &= \frac{\mu_0}{4\pi} \oint \frac{d\vec{l} \times \vec{r}}{r^3} \\
 &= \frac{\mu_0}{4\pi} \sum_n \int_{\vec{d}_n}^{\vec{d}_{n+1}} \frac{d\vec{l} \times \vec{r}}{r^3} \\
 &= \frac{\mu_0}{4\pi} \sum_n \left( \vec{f}(\Delta\vec{d}_n, \vec{r}_n) - \vec{f}(\Delta\vec{d}_n, \vec{r}_{n+1}) \right)
 \end{aligned} \tag{8.4}$$

Here, the  $\vec{d}_n$  are the points listed in the winding data,  $\Delta\vec{d}_n = \vec{d}_{n+1} - \vec{d}_n$ ,  $\vec{r}_n = \vec{r} - \vec{d}_n$ ,  $\vec{r}$  is the point at which the calculated field is observed, and  $\vec{f}$  is defined by

$$\vec{f}(\vec{x}, \vec{y}) = \frac{\vec{x} \cdot \vec{y}}{(\vec{x} \times \vec{y})^2} \vec{x} \times \vec{y}. \tag{8.5}$$

The derivatives of  $\vec{f}$  with respect to its second argument are easily calculated as

$$\partial_{y_i} \vec{f}(\vec{x}, \vec{y}) = \vec{f}(\vec{x}, \vec{y}) \left( \frac{\vec{x} \cdot \hat{e}_i}{\vec{x} \cdot \vec{y}} - 2 \frac{(\vec{x} \times \vec{y}) \cdot (\vec{x} \cdot \hat{e}_i)}{(\vec{x} \times \vec{y})^2} - \frac{\vec{y} \cdot \hat{e}_i}{y^2} \right) + \frac{\vec{x} \cdot \vec{y}}{(\vec{x} \times \vec{y})^2} \vec{x} \times \hat{e}_i. \tag{8.6}$$

Using Eq. 8.4 along with Eq. 8.6, we can calculate the field gradients.

The mirror coils are much easier to model. The parameters used to construct them are described in "Harvard\_ATRAP\_Design\_2009\_Submitted.pdf", which is included in the project for completeness. Each coil can be approximated as two

blocks, one above the midplane of the trap and one below. Each block consists of (axially) stacked sets of concentric loops. In other words, a cross-section of the block in the r-z plane would feature two dimensional arrays of coil ends. With the dimensions given in the design, the fields of a set of current loops can be added to give the approximate field generated by each coil.

To find fields and field derivatives for the mirror coils, we start by calculating the vector potential of a current loop in the Coulomb gauge. Assuming the loop has a radius  $a$  and lies in the x-y plane, the vector potential in cylindrical coordinates can be expressed as

$$\begin{aligned} \frac{1}{I_0} \vec{A}^{(a)} &= \frac{\mu_0}{4\pi} \oint \frac{d\vec{l}}{r} \\ &= \frac{a\mu_0}{\pi\beta} \left( \frac{2-k^2}{k^2} K(k^2) - \frac{2}{k^2} E(k^2) \right) \hat{\theta}, \end{aligned} \quad (8.7)$$

where  $\beta^2 = (a+\rho)^2 + z^2$ ,  $k^2 = \frac{4a\rho}{\beta^2}$ , and  $K(k^2)$  and  $E(k^2)$  are complete elliptic integrals as a function of the modulus squared. Taking the curl of this equation, we get the magnetic field

$$\begin{aligned} \frac{1}{I_0} \vec{B}^{(a)} &= \frac{\mu_0}{2\pi\beta} \left[ \frac{z}{\rho} \left( \frac{a^2 + \rho^2 + z^2}{\beta^2 - 4a\rho} E(k^2) - K(k^2) \right) \hat{\rho} \right. \\ &\quad \left. + \left( \frac{a^2 - \rho^2 - z^2}{\beta^2 - 4a\rho} E(k^2) + K(k^2) \right) \hat{z} \right]. \end{aligned} \quad (8.8)$$

The field derivatives, which can be found in a straightforward manner, are not included here as they are messy and not illuminating. To find the fields and field derivatives due to one of the mirror coils, we only need to add up the contributions due to each loop in each block. For example,

$$\frac{1}{I_0} \vec{B} = \sum_{a, z_0} (\vec{B}^{(a)}(\vec{r} - z_0 \hat{z}) + \vec{B}^{(a)}(\vec{r} + z_0 \hat{z})), \quad (8.9)$$

where sums run over the radii and axial offset of each loop in a block. Note that in this discussion,  $I_0 > 0$  corresponds to current moving counterclockwise, with respect to  $\hat{z}$ , around each current loop.

`Calculate_Fields` computes the fields and derivatives as outlined above. The program reads information on the coil, coordinate grid, and output file name from an options file. For each point on the coordinate grid segment of interest, it calculates the field (in Tesla) and, if specified in the options file, the derivative (in Tesla per meter) for a 1 A current using one of the two methods above. Finally, the results and a copy of the options file are printed. This ensures that the parameters used to generate the data are recorded for future reference. It's worth noting that many of the basic tasks (parameter loading, printing data to file) are performed by member functions of the *grid\_quantity\_template* class discussed in the previous section.

### 8.2.4 Derived Grid Quantities

Fields and field derivatives due to individual coils are useful building blocks for the quantities we need to know during antihydrogen simulations. However, loading information from all four coils, combining it with the currents at each step of the simulation, and then calculating useful derived quantities is costly and often unnecessary. Since we generally use only two sets of current configurations, it is worthwhile to compute the quantities we will need on a grid for each current setting. The commonly used sets of currents are listed in Table 8.1. The quantities we need to find are  $\vec{B}$ ,  $|\vec{B}|$ , and  $\vec{a} = \vec{F}/m_{\bar{H}}$ . Note that to a good approximation, the force depends only on fields and their derivatives.

	Quadrupole Trap	Octupole Trap
Quadrupole	470.0	0.0
Octupole	0.0	680.0
Pinch	-310.0	-210.0
Bucking	264.1	179.0

Table 8.1: Usual coil currents for each trap, in Amperes

The problem of calculating these three quantities is best tackled in two steps. First, `Sum_Fields` computes the fields and derivatives due to a fixed set of currents in the Ioffe trap coils and a given uniform background field. The names of the directories storing the results of `Calculate_Fields` for each coil and the values of the corresponding currents (in Amperes) are loaded from an options file. The components (in Tesla) of the background field, which is usually taken to be  $(1T)(-\hat{z})$ , are also lifted from the file. A simple linear combination of the component fields and derivatives is taken for each grid point, added to the background field, and the results are printed to file.

In the second step, `Calculate_Derived_Quantities` uses options from a file to determine which `Sum_Fields` outputs to load, which quantities to calculate, and where to write the results. After following the instructions on how to perform these tasks, grids of quantities needed for our simulations are available. Unfortunately, they are only defined on a finite, discrete set of points. This shortcoming will be remedied in the next section.

Once again, we mention the role of the `grid_quantity_template` class in each of these steps. The tedious tasks of locating and loading outputs from other programs,

determining the coordinate grid to be used, ensuring that the outputs and coordinate grid are all compatible with one another, and printing the final results to file are all handled by instantiations of this class. The power of the uniform interface is also apparent as a developer can use many of the same member functions for versions of *grid\_quantity\_template* with different template parameters, significantly easing the learning curve.

### 8.2.5 Interpolation

The structures we've discussed so far give simulated atoms access to relevant physical quantities on a sparse set of points. To extend this information to the entire trap volume, we need to define an interpolation. In this section, we explore three different interpolating functions applied on grids of three different spacings. The errors will be analyzed and justification will be given for the interpolation used in what follows. To make the writing more concise, we take  $f$  to be the function for which we seek an interpolation.

First, let's discuss the candidate functions. For simplicity, imagine that space is scaled so each axis has unit grid spacing and that the origin is moved so the atom of interest is at position  $(x, y, z)$  with  $0 \leq x, y, z < 1$ . If we want an affine interpolating function, we need to solve for four coefficients. This can be done by selecting four nearby grid points and forcing the interpolator to perfectly match  $f$  at these points. Of course, there are multiple quartets of points that may be chosen, and none is better than another for the range of interest. We average over these choices to reduce directional bias.

When I began this project, it was not immediately obvious that an affine interpolation would suffice as an approximation to the complex fields of our Ioffe trap. I thus built in the flexibility to use one of four interpolating functions. The first approximates  $f$  on a unit cube by averaging  $f$  over the eight surrounding grid points. This zeroth-order approximation was to be used more as a baseline than as a useful interpolator. Second is the affine function outlined above, though it's referred to as linear within the code. Next, we have a trilinear interpolator. This approximation consists of a linear combination of terms  $x^i y^j z^k$ , where  $i$ ,  $j$ , and  $k$  are either 0 or 1. Finally, the quadratic interpolator uses a linear combination of terms  $x^i y^j z^k$  with  $i + j + k \leq 2$ . There are certainly more options than those I've listed here, but this array seemed diverse enough to hold at least one solution while simple enough to implement and debug quickly.

I relied on three criteria for choosing my interpolator: speed, errors on  $\vec{B}$  and  $\vec{a}$ , and low energy drift for propagating atoms. I immediately found that the speed of an interpolation calculation is approximately proportional to the number of coefficients that implementation uses. The linear, trilinear, and quadratic approximations use 4, 8, and 10 coefficients, respectively. As long as the the other criteria were met, linear interpolation promised to outpace the other forms by nearly a factor of two.

Errors on  $\vec{B}$  and  $\vec{a}$  can be easily calculated, but it's unclear what constitutes too large of an error. One useful metric to compare these with is the uncertainty due to coil placement. Looking at the precision to which coil points are specified in their defining files, it appears that a  $50 \mu m$  error would appropriately describe perturbations around ideal placements. Applying Gaussian noise with this standard

deviation to coil positions, we can find fields and accelerations due to perturbed coils. Fig. 8.1 compares perturbation errors with those due to interpolations on grids of various sizes. Comparisons are made for a single random point in each  $1 \text{ mm}^3$  cube inside the the trap volume.

A few patterns are readily apparent. First, interpolation errors in the quadrupole trap are much smaller than those in the octupole traps. They compare favorably, even for a 1 mm grid spacing, with the uncertainties due to coil positions. For the octupole traps, grid spacing is more important. While errors from an interpolation grid with 0.4 mm spacing are generally smaller than those due to perturbed coils when calculating the magnetic field, they become comparable for acceleration calculations. Consequently, we choose the 0.2 mm grid spacing.

It's also important to check that errors in Zeeman shifts due to interpolations don't affect our results too much. For laser cooling on the 1S-2P transition, the shift is approximately 14 GHz/T. For the 1S-2S transition we focus on, it is 961 kHz/T. Most interpolated fields are within 0.1 mT of their exact counterparts. This corresponds to a 1.4 MHz shift in the cooling frequency and a 96 Hz shift in the 1S-2S frequency. The cooling transition has a natural width of about 100 MHz, making the error in field negligible. On the other hand, the narrowest 1S-2S lineshape we observe in simulations has a FWHM of 600 Hz. It's inadvisable to ignore errors of 96 Hz in this situation.

Thankfully, we do not need to resort to a slower interpolation or a denser grid to fix this problem. The points probed in Fig. 8.1 extend to walls of the trap, where quickly changing fields lead to large interpolation errors. Lower energy atoms gener-

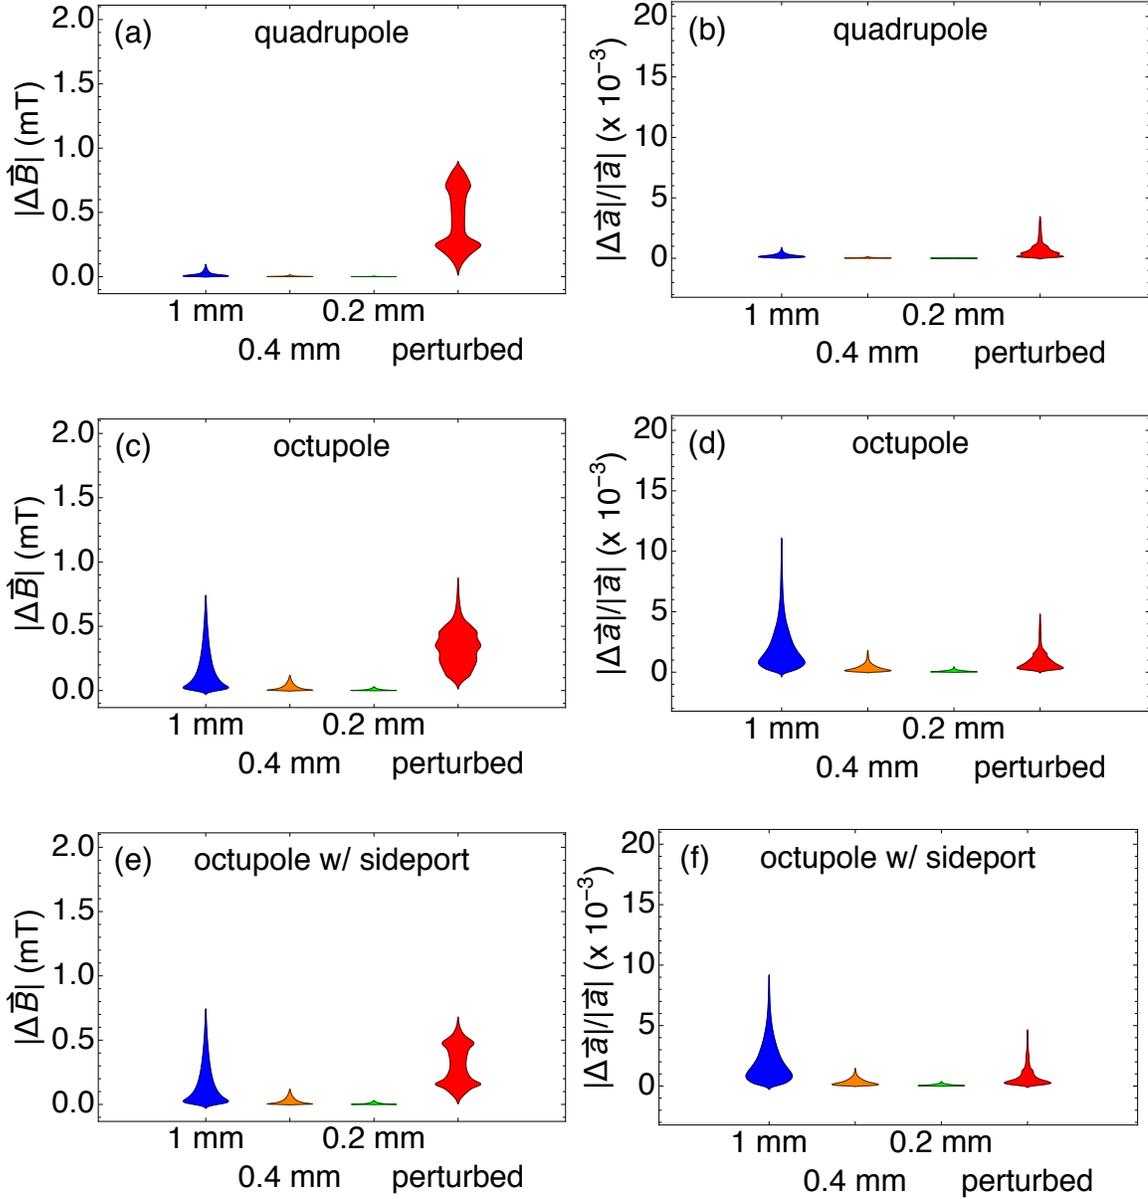


Figure 8.1: Errors due to interpolation on grids of different spacings as well as those due to perturbed coils are plotted for the magnetic field (a, c, e) and for the atomic acceleration (b, d, f). Results are shown for the (a, b) quadrupole, (c, d) octupole, and (e, f) octupole with sideports traps.

ally stay further from trap edges and yield narrower lineshapes. To account for these facts, we consider atoms with different energy distributions separately. As will be explained later, for each laser and trapping configuration, spectroscopy simulations were performed for atomic samples with three different energy distributions. In order of decreasing energy, these are labeled as uncooled, cooled, and Doppler limit (cooled). For each trap, distribution, and laser axis, we plot (Fig. 8.2)  $|\vec{B}|$ 's interpolation error, transformed to a Zeeman shift, for all in-beam points accessible to the most energetic member of the distribution, divided by the minimum linewidth observed in our spectroscopy simulations when varying other laser parameters.

Fig. 8.2 shows the Zeeman shift errors to be relatively minor. They may need to be reduced for a high-precision comparison between hydrogen and antihydrogen 1S-2S frequencies. However, the main goal of this work's spectroscopy simulations is to explore what the CTRAP apparatus is capable of. The observed errors don't appreciably change the lineshapes we generate.

The final criterion to check is energy stability. Since the calculations for fields and accelerations are largely independent, this gives a rough check that they yield compatible results. Of course, the result is also dependent upon our step size in time (1  $\mu$ s), so it's not a clean diagnostic. Fig. 8.3 shows that energy is largely conserved over long time scales.

To implement any of the interpolations mentioned above, I define the *interpolator* class in `Interpolators.*`. This is derived from *grid\_quantity\_template*, defined earlier. *Interpolator* bundles arrays of interpolation coefficients for all quantities of interest with instances of *cart\_coord\_grid* and *grid\_position*. The latter class handles the

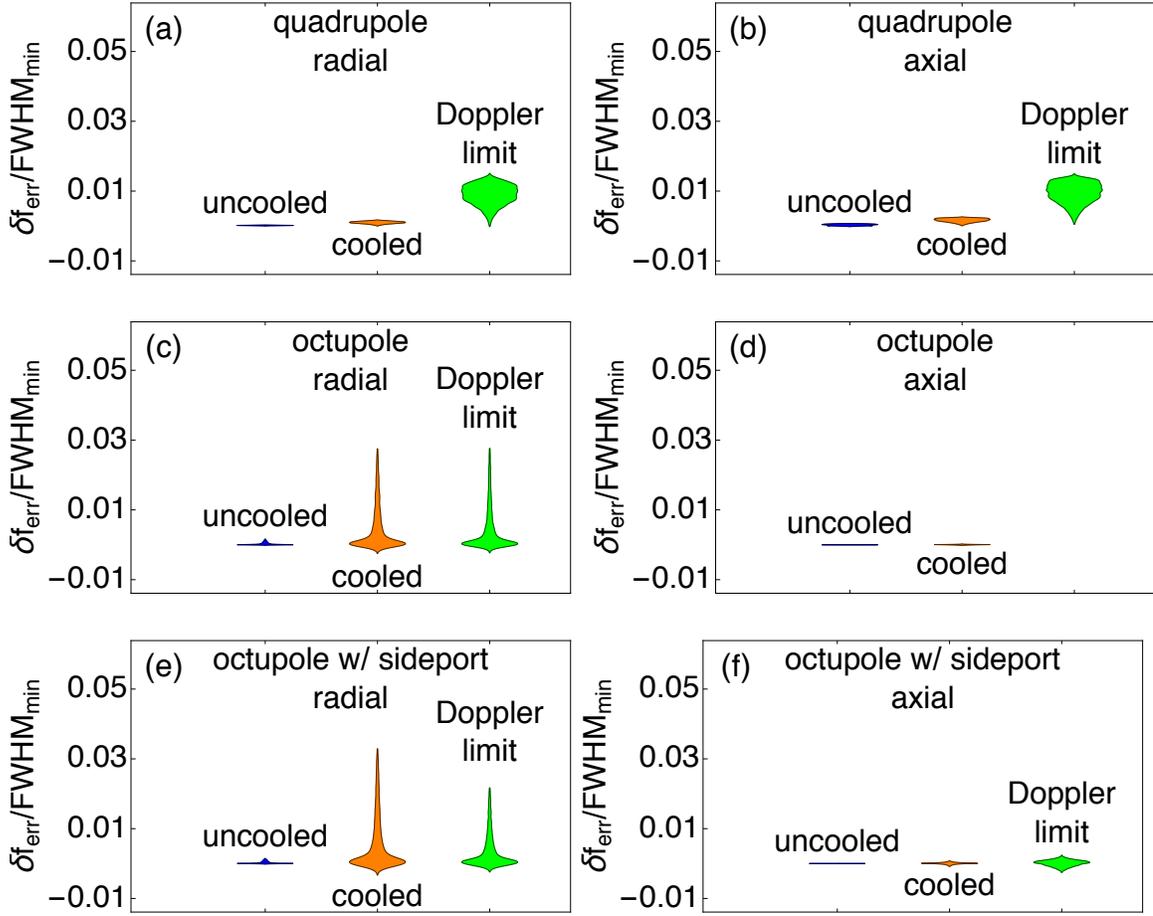


Figure 8.2: Errors due to interpolation are calculated along the transverse laser path (a, c, e) and along the axial path (b, d, f) for the (a, b) quadrupole, (c, d) octupole, and (e, f) octupole with sideports traps. Curves that fail to break the 30% loss threshold during an hour-long spectroscopy run are omitted from this analysis. For this reason, no data from a sample cooled to the Doppler limit is presented in (d).

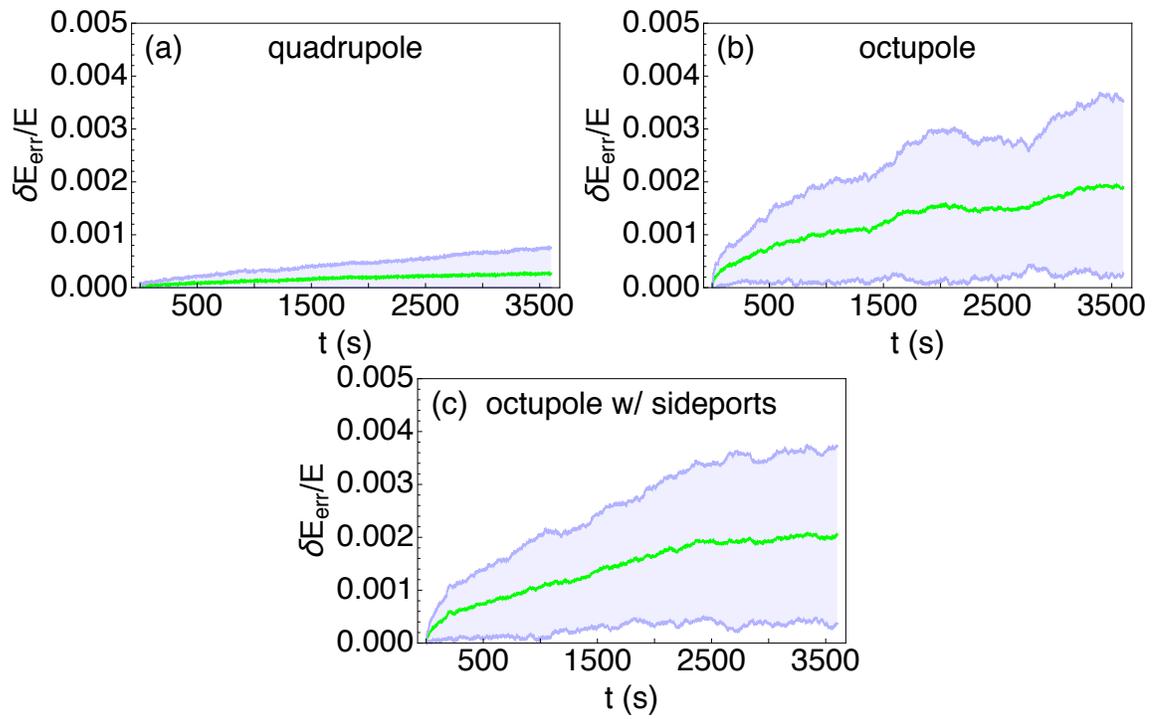


Figure 8.3: The mean (green) and one- $\sigma$  intervals (blue) are plotted for the absolute fractional change in energy for particles in the (a) quadrupole, (b) octupole, and (c) octupole with sideports traps.

tasks of determining when and where out-of-bounds queries have been made, storing positional information relative to the grid, and storing multinomials needed for interpolation calculations. These cut down on redundant function calls when calculating multiple quantity values at a single point. *Interpolator* objects also cache the last set of interpolation coefficients used for each quantity, reducing the frequency with which the large arrays of coefficients need to be read.

These classes, along with explicit expressions used for interpolation, are stored in `Position_Handlers.*`, `Interpolation_Data.*`, `Interpolation_Grids.*`, and `Interpolators.*`. Arrays of coefficients are produced by `Interpolate_Fields` using the output of `Calculate_Derived_Quantities`. The structure of these program is involved, and the range of capabilities is correspondingly wide. However, we use only linear interpolators in what follows. Still, for more precision, a higher-order approximation may be useful in the future.

## 8.3 Kinematic and Atomic Evolution

### 8.3.1 Kinematics

The class *kinematic\_handler* provides a general interface for calculating the evolution of kinematic variables. To recognize the power of its instances, its important to understand its parts.

*EM\_Quant\_Gen* is a template parameter giving the base class from which *kinematic\_handler* is derived. It is required to have a few basic components.

- `initialize()` takes the path to an options file, which is then used to set member

variables if needed.

- `EM_quant()` gives the field, field derivative, field magnitude, or acceleration at a point, depending on its template argument.
- `out_of_bounds()` returns a boolean indicating whether a quantity's value has been queried outside the trap volume.
- `reset_out_of_bnds()` resets the internal state of *EM\_Quant\_Gen* so future calls to `out_of_bounds()` return "true" until the next out-of-bounds query is made.
- Member `Bnorm_min` gives the minimum field magnitude in the trap if the trap fields are static and known during initialization.

Classes included in this package which are used as *EM\_Quant\_Gens* include the following.

- *grid\_EM\_quants* holds interpolation information for a fixed trap configuration.
- *grid\_EM\_quants\_changing\_currents* holds interpolation information for each coil's 1 A field and field derivatives along with time series information on each coil's current. This allows for the calculation of fields and accelerations for a changing trap configuration.
- *prompt\_EM\_quants* takes user input to determine output quantities.
- *one\_EM\_quants* returns "1" when the value of any quantity is queried.

The first of these is used in most simulations. The second is used when simulating trap dumps. The final two classes are used for testing purposes.

A stepper of (template) type *Step\_T* advances the kinematic state using the acceleration given by *EM\_Quant\_Gen*. Boost's Odeint library provides *Step\_T* for all non-test cases. For trap dumps, we use the 5th-order *runge\_kutta\_cash\_karp54* stepper. For for all other simulations, we use the symplectic, 4th-order *symplectic\_rkn\_sb3a\_mclachlan*. In each case, the stepper is called by the method *do\_step()*. The internal state, time, and step size are stored in the *kinematic\_handler* type, so no parameters are needed for this call.

An *energy\_handler* instance is included in each *kinematic\_handler*. It takes *EM\_Quant\_Gen* as a template parameter and uses it to determine the energy at each point in phase space by comparing the potential energy corresponding to the local  $|\vec{B}|$  to that of *EM\_Quant\_Gen*'s *Bnorm\_min* and adding kinetic energy. The result is returned by *get\_gr\_E()*. It's worth pointing out that while virtual inheritance is found nowhere in my code, *EM\_Quant\_Gen*'s *interpolator* members (if there are any) are static. Therefore, each mapped file, cached set of coefficients, and *grid\_position* occurs only once. This reduces the amount of memory needed and unifies the pool of cached information.

In addition to these fundamental pieces, *kinematic\_handler* stores several objects that simplify the importing and exporting of information from its instances. The options file read by the constructor may refer to a file holding a set of initial conditions. A portion of these initial conditions are stored inside the *kinematic\_handler* so they can each be addressed by the simulation. Alternately, template parameter *Init\_Num*

can be set to the number of initial conditions to be passed to the *kinematic\_handler* instance from another object. Objects also hold a vector of *state\_summary* instances, a type which is discussed in a later section, for describing intermediate and final states of atoms. A simple call to `print_contents()` prints the content of this vector to a file name loaded from the options file during construction. With these components, *kinematic\_handler* makes it easy to set up, run, and record the results of a simulation.

### 8.3.2 Atomic States

We are only really concerned with antihydrogen's properties for two principle quantum numbers. Atoms are assumed to start in a 1S state before being driven to a 2S or 2P state by our spectroscopy or Lyman- $\alpha$  laser, respectively. Other levels have effects on the evolution of an atom's internal state, especially for 1S-2S spectroscopy, but the fact that their contributions are far off-resonance allows us to ignore variations due to local quasi-static fields.

Even with such a limited scope of interest, we are concerned with a large number of states. For each spatial wavefunction, spin degrees of freedom give rise to a four-dimensional state space. Therefore, we must deal with four 1S, four 2S, and twelve 2P energy eigenstates. The class *ket* is designed to perform numerous operations requiring an understanding of the eigenstates of interest.

Instantiations of the *ket* template class represent specific eigenstates. The template parameters `N` and `Nm_Rng` establish the principle quantum number of the state and the valid set of states, respectively. The heart of each *ket* object is a map from strings to integers. The strings label integer linear combinations of spins (`F`, `mJ`,

2mI) while the mapped integer is the value of the linear combination for the state in question. Member functions include the following.

- The `fill()` function is overloaded to import a *ket*'s map from either an options file or directly from a map fed to the function.
- The static `basis()` function returns a vector of all *kets* with a specific principle quantum number and orbital angular momentum
- Comparison operators (`==`, `!=`, `<`, and `>`) are overloaded to determine whether two *kets* are the same and the order in which they appear in a call to `basis()`.
- The parentheses operator is overloaded to send strings to the integer values specified by *ket*'s map.

With this relatively simple class, programs are able to efficiently manipulate an query atomic states.

### 8.3.3 Laser Information

It is useful to have the ability to try different laser configurations by passing options to a simulation. In this spirit, three classes are defined: *laser\_properties*, *laser\_path*, and *path\_mngr*. The first of these characterizes fundamental properties of a source being used. The latter two give path-specific information.

The *laser\_properties* class loads data from an options file and stores it. This information includes the detuning of the laser and its FWHM. The detuning is generally taken relative to the 1S-2S resonance of interest plus the Zeeman shift corresponding to the trap minimum. However, an option exists allowing the zero-detuning frequency

to be adjusted. The *laser\_properties* class also determines whether the laser is pulsed or continuous. If it is pulse, the pulse length and repetition rate must be specified.

Class *laser\_path*'s purpose is to define a single laser path. Its contents, which are loaded from an options file, include the following.

- Member `one_way_P` gives the power transmitted through a unidirectional path or the one-way power propagating in a cavity oriented along the path.
- Member `unit_wave_vector` is a vector in the direction of the beam's wavevector. For the case of a build-up cavity, this vector may be oriented in either direction.
- Member `focus` sets the position of the focus for the path. In particular, this allows for beams parallel to, but not coinciding with, the axes of our trap coordinate system.
- Member `waist` gives the beam waist.
- Member `pol` specifies a complex polarization vector for the incoming radiation.

In addition to these values, two member functions exist to extract important information depending on the laser path. To determine whether a particle can be affected by the beam, `in_beam()` may be queried. This function returns true if the position passed to the function is within twice the waist plus a buffer. The buffer is usually set based on the time step and the trap depth for a particular simulation. It ensures that the beam is able to drive transitions in atoms which start a time step outside of the double-waist threshold, but may enter the beam during the step. Note that the intensity of a Gaussian beam is  $3 \cdot 10^{-4}$  times its maximum value at a radius of twice

the waist, so the chosen cutoff seems reasonable. A second method, `intensity()`, maps a point in three-space to the intensity of the beam at that point.

The final class, `path_mngr`, allows for easy management of several `laser_path` objects. The central component is a vector of `laser_paths`. This vector is filled on initialization of a `path_mngr` instance by reading a series of file paths from an options file. Each file path is passed as a separate options file to the constructor of a `laser_path`, and the resulting object is appended to the `path_mngr`'s vector. This allows for the reuse of files specifying beam paths. This class has only one member function, `in_beam()`, which effectively acts as an "or" for methods of the same name applied to each `laser_path` in its vector.

### 8.3.4 Atomic Physics

Many of the quantities needed to calculate the evolution of an atom's internal state depend on the local magnetic field. Often, these dependences are too complicated to calculate accurately and quickly. For this reason, we calculate them for a dense array of magnetic field values prior to running any simulations. This is done by the Mathematica scripts in the directory `Transition_Quantity_Generator`. The types of calculations done by these scripts can be found in the chapter on atomic physics. The command "make tq", run from the top directory of this project, executes the Mathematica scripts, generating the arrays of interest.

The class `arr_mngr` is used to handle these arrays. Each instance holds the difference in magnetic field,  $dB$ , between array points, a vector of atomic quantities at fields  $n \cdot dB$  for integers  $n \geq 0$ , and the most recent quantity returned. The `fill()`

method takes a string identifying the quantity to load and one or more *kets*, then loads the file holding the array of interest into the object's vector. The *arr\_mgr* class also has an overloaded parentheses operator. With no arguments, it returns the last quantity returned. When passed a field magnitude, it returns the quantity specified by the *arr\_mgr*'s vector, linearly interpolating between the two nearest points if the field fails to have the form  $n \cdot dB$ .

### 8.3.5 Internal State Evolution

This is the first of three subsections describing how we advance the internal states of our atoms. In each part, we explore classes for handling single atoms as well as classes for handling ensembles of atoms. The former is applied to laser cooling with a pulsed Lyman- $\alpha$  source while the latter is used for 1S-2S spectroscopy. The reason for this division is that an object from one of these classes evolves alongside a single *kinematic\_handler* instance. Consequently, an ensemble of internal states is only appropriate if the kinematic evolution is relatively decoupled from the internal evolution. This is obviously not true for laser cooling, but we will try to justify this point of view for spectroscopy.

The class *illumination\_handler\_base* acts as the interface between the external trap environment and the internal degrees of freedom. It is a template class that takes parameters *IT* and *EM\_Quant\_Gen*, which control the transition being probed (1S-2P vs. 1S-2S) and the method for determining local fields, respectively. The class derives from *local\_EM\_fields*, *laser\_properties*, and *path\_mgr*. Instances of *local\_EM\_fields* compute and cache local field values such as  $\vec{B}$ ,  $|\vec{B}|$ ,  $\vec{E}$ ,  $\vec{E}^2$ , the local

laser intensity, and the polarization relative to a  $\vec{B}$ -aligned coordinate system. In addition to the base classes mentioned above, objects of type *illumination\_handler\_base* also contain *arr\_mgr* objects corresponding to Rabi frequencies, decay and ionization rates, and different sources of frequency shifts. These are populated, and base classes initialized, by the *fill()* method. The *update()* member function computes new field values and fills *arr\_mgr* caches with corresponding values. It then combined these into total detunings, decay rates, and Rabi frequencies by calling *update\_derived\_quants()*.

Objects of type *illumination\_handler* manage the evolution of a single atom's internal state. Each instance contains three fundamental components.

- Member *d\_mat* is a four-dimensional real array tracking the probability amplitudes for the two states connected by the transition of interest.
- The vector *acc\_phases* tracks the phase of each time-dependent Rabi frequency. Doppler shifts cause beams to have different frequencies along different paths in the atom's rest frame, so this bookkeeping is necessary when adding their effects.
- Member *decay\_threshold* gives the level below which probability in our two-state system must fall to assume a decay has occurred.

These are initialized to reflect an atom in the ground or excited state by way of the *initialize()* method. The atomic state is moved forward in time by *integration\_step()*. It uses *illumination\_handler*'s base *illumination\_handler\_base* to calculate parameters for the state's equation of motion, the parentheses operator to calculate derivatives,

and Boost Odeint's default integrator to solve the equation of motion. The function `integration_step()` has two noteworthy properties. First, an "observer" (function, lambda, or callable class) may be passed in to check some quantity or accumulate statistics for every step of Boost's integrator. Second, integration stops and `integration_step()` returns the time elapsed whenever the probability in the ground and excited state drops below `decay_threshold`. This allows the simulation to handle a decay and, if the atom returns to its original ground state, resume its evolution.

For the evolution of an ensemble of atomic states, *illumination\_array\_handler* is used in place of *illumination\_handler*. The only difference between states in the ensemble is the frequency of the laser source being used. A vector holds the set of detunings relative to resonance at the trap minimum. Instead of a single real four-vector representing an atomic state, we use a pair of complex values. Of course, we have a vector of these, one for each detuning. As each member of the ensemble decays independently, a vector of booleans determines which members have yet to decay out of the original ground state or ionize. Another vector sets individual thresholds for decay. The `initialize()` and `integration_step()` methods perform similar functions to those explained above, again using the base *illumination\_handler\_base* to find parameter values based on local fields. Finally, `check_for_excitation()` populates the vector `decay_idxs`, which is then passed to the ensemble decay handler discussed in the next subsection.

### 8.3.6 Decays

It's possible for an atom to decay at any point when the probability in the excited state is nonzero. To manage these decays, two classes are defined, *decay\_handler* and *decay\_array\_handler*. The former works on single states and gives information on kinematic changes caused by decays. The latter works on ensembles of states and ignores absorption and emission recoils.

Objects of type *decay\_handler* are similar to those of type *illumination\_handler\_base*. They contain *arr\_mgrs* to keep track of decay and ionization rates, branching ratios, and transition frequencies. They have a *fill()* method to populate their *arr\_mgrs* as well as *update()* and *update\_derived\_quants()* member functions to keep values for local quantities current. In addition, four methods for carrying out decays are included.

- *choose\_channel()* decides whether a decay proceeds by ionization, single-photon emission, or double-photon emission. The choice is made randomly, weighted by the branching ratios.
- *execute\_one\_ph\_decay* and *execute\_two\_ph\_decay* each pick a ground state to decay to randomly, weighted by branching ratios. The recoil impulse(s) are chosen from the appropriate distribution(s), added to the absorption impulse(s), and stored for later addition to the atom's momentum.
- *execute\_decay\_chain* updates local quantities, chooses the decay channel, and runs the appropriate decay method.

Thus, *decay\_handlers* allow the entire decay process, with state and velocity changes,

to be carried out in a single function call.

Instances of *decay\_array\_handler* take *decay\_handler* as a base class. Each has a vector of ground states holding the results of decay, but velocity changes are ignored. The `fill()` method calls *decay\_handler*'s `fill()` and initializes a few of *decay\_array\_handler*'s arrays. The function `execute_vec_decay()` uses the methods of *decay\_handler* to determine how decays proceed for atoms corresponding to the indices in `decaying_idx`s.

In addition to calculating changes in momentum and ground state, *decay\_handler* and *decay\_array\_handler* also track decay statistics. For each atom under an instance's purview, the number of one-photon decays, the number of two-photon decays, and whether ionization has happened are recorded. These are conveniently packaged and printed, along with other useful information, by the *state\_summary* class, which is discussed later.

### 8.3.7 Evolution and Decay

Classes *atomic\_state\_evolution\_handler* and *atomic\_array\_evolution\_handler* conveniently package evolution handling and decay handling into a single object. This allows switches between internal state evolution and decays to happen automatically. It also lessens the burden on the user to remember the format of each function call. Each class contains objects of type *illumination\_(array\_)handler* and *decay\_(array\_)handler*. Also available are the `in_beam()` and `initialize()` methods, which call the corresponding members of *illumination\_(array\_)handler*, and the `fill()` method, which uses an options file to set the values of contained objects.

Despite these similarities, the two classes differ in how they execute time steps. When passed a point in phase space, the *atomic\_state\_evolution\_handler*'s `do_step()` method allows the internal state to evolve until it decays. Once the decay has been handled, if the atom has returned to the original ground state, the evolution continues. These processes alternate until the atom fails to return to the starting ground state or the allotted evolution time elapses. For a pulsed laser, the evolution time is equal to the pulse length. For a continuous laser, it is equal to the kinematic time step. Pulsed cooling simulations force the remaining excited state probability to decay immediately after `do_step()` is executed, managing any resulting transitions to the ground state.

The *atomic\_array\_evolution\_handler* class executes two different types of step sequences: `do_steps_in_beam()` and `do_steps_decay()`. The decisions to handle sequences of steps together, use different function for in-beam and out-of-beam propagation, and use a pair of complex numbers rather than a quartet of real numbers stem from optimization opportunities. The speedup is important since many of the colder atoms spend a considerable fraction of their time either in the laser beam or decaying after exposure. Indeed, before applying these changes, some atoms would take on the order of an hour to execute one second of simulation time!

Practically, `do_steps_in_beam()` carries out steps by first simulating an atom's propagation until it leaves the beam, hits a wall, or the total simulation time has elapsed. During this pass, an upper bound is placed on the probability in the excited state and a lower bound on the probability in the ground and excited states, each as a function of an ensemble member's initial state. Time and kinematic variables

are then reset to their values prior to the bounds-setting pass. Members whose ground + excited bound falls below their decay threshold or whose excited state bound exceeds one thousandth of their ground + excited probability are then treated with a full kinematic + atomic simulation. This allows the simulation to skip expensive calculations that fail to have much of an effect on atomic states.

Once atoms leave a beam, their evolution is handled by `do_steps_decay()`. Atoms with less than one thousandth of their ground + excited state probability in the excited state are ignored, reducing the computational cost of this function. Ensemble members that remain generally evolve by having their excited amplitudes scaled by  $\exp(-\Gamma dt/2)$  each time step, though a fuller calculation is possible if the Rabi frequency exceeds 1 mHz. Any members that cross the decay threshold have their decays managed by the `decay_array_handler` and are ignored for the rest of `do_steps_decay()`'s execution.

### 8.3.8 State Summary

Thorough tests involving large numbers of particles and a variety of laser and trap conditions require lots of computational time. After carrying out these simulations, unanticipated questions may arise. It is thus important to keep detailed records during simulation runs that retain as much information as possible. To standardize the information we keep along with the methods for printing and retrieving this data, we define the classes `state_summary` and `st_sum_arr_hand`.

The `state_summary` class contains 17 quantities we want to track during particle propagation.

- `pos_vel` is a pair of three-dimensional vectors representing a particle's position and velocity in SI units.
- `E` gives a particle's energy divided by Boltzmann's constant in Kelvin.
- `N`, `F`, and `mF` give an atom's principle quantum number and hyperfine state. We assume  $L = 0$ .
- `one_photon_decays` and `two_photon_decays` give the number of decays in each channel a particle has undergone during a simulation.
- `ionized` and `out_of_bounds` are booleans determining whether an atom has yet been ionized or reached a wall.
- `beam_passes` and `t_in_beam` give the number of passes a particle has made through the laser as well as the total time spent in the beam.

In addition, helper functions make it easy to move this data to and from a file system. The function `print_state_summary()` takes a vector of `state_summaries`, converts each element to a list of 17 number in the order listed above, and prints the results to file. On the other hand, `extract_state_summaries()` partitions files into blocks of 17 numbers, plugs each number into the corresponding `state_summary` field, and returns the resulting vector of `state_summaries`.

Class `st_sum_arr_hand` is provided to track atom ensembles. It contains a vector of `state_summaries` and a method for adding elements to the vector. Member function `finish_st_sum_arr()` takes a `kinematic_handler`, an `illumination_array_handler`, and a `decay_array_handler`. It adds a `state_summary` corresponding to each ensemble member's current state to the `st_sum_arr_hand`'s vector. An instance of

this type is included in each *atomic\_array\_evolution\_handler*, mainly to reduce the overhead of collecting data from the different handlers. It's perhaps worth adding a corresponding class for objects of type *atomic\_state\_evolution\_handler* for a similar reason, but this has not yet been done.

## 8.4 Simulations

The libraries outlined so far provide the tools necessary to simulate experiments in our trap with atoms and lasers. In the directory *Hbar\_Propagation*, many programs exist to run tests on these tools and to try out different simulation methods. However, only a few are used regularly. The directory *Submission\_Scripts* holds scripts that show the importance and frequency of use of different programs. More explicitly, the three most commonly run are as follows.

- *Pulsed\_LyAlpha\_Monitor* performs laser cooling on a given sample of atoms.
- *CW\_1S2S\_Arr\_gen* performs 1S-2S spectroscopy on trapped atoms.
- *Dump\_Trap* tracks particles as the trapping fields are lowered to zero.

In addition to these, a few programs are used to generate initial conditions. The most basic simulation aims to replicate the distribution of antihydrogen atoms produced and trapped when antiprotons and positrons are mixed. We start by assuming that the positrons occupy a uniform-density, uniform-temperature ellipsoid at the trap minimum. The density and dimensions are those targeted in recent experiments: a radius of 1 mm, a half-length of 1.5 cm, and a density of  $5.0 \cdot 10^7 \text{ cm}^{-3}$ . The temperature, based off measurements of typical plasmas in our trap [45], is taken to be 20 K. Each

atom is generated at a random point in the positron plasma. Its velocity is picked from a Boltzmann distribution at the prescribed temperature and added to the local rotational velocity of the plasma. The particle is allowed to propagate through the trap for one minute. If it remains confined for that time, *state\_summarys* reflecting the initial and final states are recorded. It is typically the latter which is used in subsequent simulations (spectroscopy, cooling, etc) since the random points chosen at the beginning of this simulation localize all atoms within the positron plasma. Atoms in random states are generated and allowed to propagate until a pre-determined number of trapped atoms have been found. The resulting set of particles is called the "un-cooled" distribution. Histograms of energies in this distribution for different trapping configurations are plotted in Fig. 8.4. It should be noted that this program is defined almost entirely in `Make_ICs.hpp` while options specific to each trap are fleshed out in `Make_(trap name)_ICs.cpp`.

A second class of initial conditions is desirable for considering the effects of laser cooling to the Doppler limit on an atomic sample. As will be shown, the cooling schemes we simulated were unable to reach this limit. It's thus worthwhile to check whether efforts to optimize cooling will yield much of an advantage for applications like 1S-2S spectroscopy. Since we cannot reach the Doppler limit using the cooling schemes we have developed, we need to cheat a bit. `Make_Doppler_Lim_ICs` takes as one of its inputs a file containing a list of *state\_summarys* representing the results of laser cooling on a sample of atoms. A scaling factor,  $E_{DL}/\bar{E} < 1$ , is computed and multiplied with each particle's energy to give a corresponding target energy. If each atom is cooled to its target, the mean energy will reach the Doppler limit. The

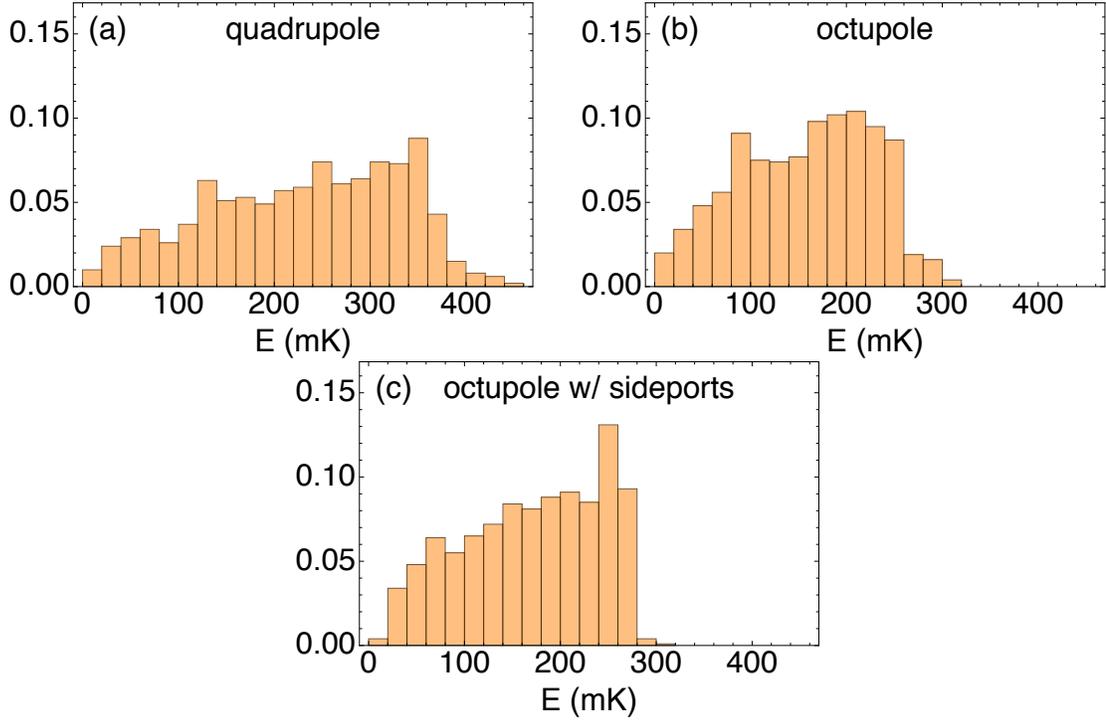


Figure 8.4: Energy distributions for uncooled initial conditions

pulsed cooling simulation is then run with the following modifications.

- If a scatter has not occurred in the past 60 seconds, a scatter is forced.
- Discard all scatters that increase an atom's energy or decrease it below its target value.
- Consider two scenarios.
  1. A scatter has been discarded and the difference between the atom's energy and the target energy is less than  $\frac{p_\gamma^2}{2m_H}$ , where  $p_\gamma$  is the momentum a Lyman- $\alpha$  photon.
  2. A scatter has been forced, but has not changed the atom's velocity.

If either occurs, the atom's velocity is scaled so its energy matches the target energy. If this is impossible, the velocity is scaled by a factor of 0.1.

Once the atom's energy is within a factor of  $10^{-6}$  of the target, the final state is recorded and the next atom is addressed.

## 8.5 Cluster Management

We've discussed the code needed to perform calculations prior to and during simulations. However, there are several practical difficulties that must be overcome. First, many of the programs outlined must be run in parallel. The serial runtimes would otherwise range from weeks to years. Second, running jobs on Harvard's Odyssey cluster, our choice for executing massive numbers of jobs, requires some human oversight. This leads to several difficulties that we will talk about below. Finally, each job requires several options to be specified. There are many awkward ways to do this, but we define an interface using Boost's ProgramOptions library to make an easy-to-use interface.

### 8.5.1 Partitions

Much of the code we've seen so far can be run in parallel in obvious ways. For instance, the calculation of fields on a grid can be performed at each grid point independently. Similarly, the behaviors of different atoms in a trap with a cooling or spectroscopy laser are uncoupled. Although it's easy to imagine how to parallelize operations on a case-by-case basis, it's more useful to have a unified way of breaking

up jobs.

Our tool for this is the *partition* class. Each instance contains global information on a certain partition as well as data specific to a certain segment. Global information is held in the following members.

- `tot_elems` gives the total number of units being partitioned. This could be points on a grid or particles being simulated.
- `tot_parts` is the total number partitions into which `tot_elems` is divided.
- `chunk_size` corresponds to the smallest number that must divide a partition's segment. For instance, if a file representing a vector of *state\_summaries* is partitioned, `chunk_size` would be 17.

Data local to an instance's segment is held in these members.

- `part_idx` is an index, which ranges from 0 to `tot_elems - 1`, identifying the segment of interest.
- `first_elem_idx` is the index, amongst all elements being partitioned, of the first element of the local segment.
- `elems_in_part` is the number of elements in the local segment.

Each *partition* object is defined by giving the three global variables and a function `txtpart_idx`. Checks are done to make sure the definitions make sense, with exceptions being raised if a test is failed. The quantity  $(\text{tot\_elems}/\text{chunk\_size})$  is then divided by `tot_parts` to give integer quotient  $q$  with remainder  $r$ . The first  $r$  seg-

ments have  $q + 1$  chunks while the remaining segments are given  $q$  chunks. These specifications make calculating `first_elem_idx` and `elems_in_part` straightforward.

There are two principle contexts in which *partition* objects are used. In both cases, instances are passed to a function, which selects its workload according to the argument. First, when loading or mapping files into data arrays, we often only want to work with a segment. For instance, when we're performing a pointwise mathematical operation on one or more quantities defined on a grid, it's often faster to break the grid into multiple pieces and perform the operation on each. The other case occurs when a simulation is performed on a collection of atoms. It's usually faster to partition the set and perform the simulations at the same time. The *partition* class provides a clean, unified way to break lengthy serial jobs into manageable parallel jobs.

## 8.5.2 Cluster Management

In the following, we use "task" to mean all of the work a program aims to accomplish while a "job" denotes the work corresponding to one partition segment. Computing the magnetic field at each point on a grid is an example of a task. Computing these fields for grid point 1000 through 1099 is a job.

As was already noted, our tool for running jobs in parallel is Harvard's Odyssey cluster. As useful as the cluster is, it also gives rise to a number of small difficulties that take away from user productivity. Examples include the following.

- Each resource request for completing a task requires a submission script. Often, we want to submit dozens of these scripts, making manual construction tedious and error-prone.

- When jobs fail, which often happens when tasks are divided into many jobs, it takes some work to figure out which failed and to write a script re-running those jobs.
- There are limits on the number of jobs that can be submitted at once, the number of characters in a script, and the maximum job index, making it important to be ever vigilant when writing scripts.
- Due to job number limits and failures, it's somewhat difficult to figure out when a task is complete. It's also time-consuming to concatenate the resulting files, so it's best to outsource this to a cluster processor.

The functions defined in `Job_Manager.*`, `File_IO.*`, and the `Cluster_Implementations` directory automate away these problems, allowing users to focus on development and analysis.

The function `pseudo_main()` is the entry point to our cluster management suite. It parses three command line options.

- `"opts_fn"` specifies an options file parameterizing the task of interest. If neither `"print_out_mem"` nor `"print_opts_fn"` is passed through the command line, the program executes this task.
- `"print_out_mem"` prints the expected storage space filled after successfully completing the task specified by `"opts_fn"`.
- `"print_opts_fn"` prints a blank options file giving the program's required inputs at the path given by `"opts_fn"`.

If none of these is passed through the command line, a list of possible command line options is printed with descriptions of their uses.

If "opts\_fn" is passed to the program without "print\_out\_mem" or "print\_opts\_fn", the function `exec_prgrm()` is run. This checks to see whether all output files the program is expected to generate exist. If they do, the task is presumed to have been completed and the program exits with a "Job complete" message. Otherwise, one of four flags determines the subsequent behavior.

- "loc" forces the entire task to run on the processor running the program. Once all jobs have been completed, the resulting files are concatenated.
- "loc\_part\_calc" runs a single job on the processor running the program. Additional command line arguments specify which job runs.
- "cluster" exports the task to the cluster, submitting it as an array of jobs to be run on separate processors. The number of jobs into which the task is divided is determined by a parameter from the options file. If all jobs have been successfully completed, file concatenation is carried out on a single cluster processor.
- "cluster\_single" is equivalent to allocating a single cluster processor and running the program with the "loc" flag.

Passing none of these flags is the same as passing the default flag, "loc".

There are three operations we've swept under the rug so far: file checking, file concatenation, and submission script writing. The function `clear_incomplete_file_set()` takes a vector of file names, checks that all of them exist, and, if at least one does

not, it deletes every existing file in the vector. The return value is true if a file was missing and false otherwise. To concatenate files, `cat_part_files()` is used. It takes a vector of file names to construct and the number of partial files that must be concatenated. It's important to note that `part_file_name()` is defined as a map from a concatenated file name and a *partition* object to the name of the file holding results from the corresponding job. This makes finding and combining these files easier.

Finally, we must examine how scripts are written. Harvard's cluster uses SLURM, so `Cluster_Implementations` contains `SLURM_Impl.hpp`. This is included in `Job_Manager.cpp`. To add functionality for a different cluster management program, code with similar functions should be placed in the same directory and included in place of `SLURM_Impl.hpp`. The primary function, `submit_SLURM_script()`, writes resource request scripts in the format accepted by SLURM. In addition, it takes command line options that set options like allocated memory and time. All management of job, job index, and character limits are performed automatically.

### 8.5.3 Program Options

Each program we've discussed requires several parameters in order to be fully defined. There are easy ways to introduce these parameters - compiling a new program each time a value changes, passing information by command line arguments - but these methods become messy when dozens of parameters need to be set or large sweeps over parameter space are performed. A better way of handling this problem is to input values using text files.

Boost's `ProgramOptions` library offers this functionality, but it's useful to wrap

some of it in custom code to bundle multiple calls that usually occur sequentially, to bundle related data, and to make its use a bit more intuitive. These definitions are made in `Program_Options_Wrapper.*`. There is one type defined, `mob_desc`, which is a tuple consisting of a string, a function returning a pointer to Boost's `value_semantic` type, and another string. An instance of this type is used to retrieve a single parameter's value. The first string gives the prompt or flag indicating where the value is written. A `value_semantic` object sets options for the parameter such as a default value, whether it's required, and where the value is to be stored. The final string gives a description of the parameter.

There are three main functions for finding parameter values.

- `get_cmd_opt()` retrieves values passed as flag arguments on the command line.
- `get_file_opt()` collects values from a specified options file.
- `get_file_opt_from_cmd()` reads the command line argument passed with the "opts\_fn" flag to determine an options file path. It then runs `get_file_opt()` using this path.

These functions are overloaded for three input types: a vector of `mob_descs`, a vector of vectors of `mob_descs`, and the Boost type `options_description`. We generally use only the first two. The return values of these functions are `variables_map` objects, which can be queried to find parameter values.

## 8.6 Summary

The material in this chapter touches on all important facets of the simulation code I developed. The descriptions give a relatively full picture of how operations are carried out and managed. However, it's important to understand that this discussion is by no means comprehensive. Extra details can easily be filled in using the raw code alongside the information in this chapter. The full code, along with a concise README, may be found on GitHub ([github.com/n09jones/Hbar](https://github.com/n09jones/Hbar)).

The results of simulations run using this code are laid out in Chs. 9 and 10. Without the work described here, these insights would either be inaccessible or would cost a great deal of valuable experiment time. In addition to what I have found, the structure of the code allows others to easily test new scenarios. The hope is that the framework I have constructed will make future progress easier.

# Chapter 9

## Laser Cooling Simulations

This chapter presents a series of simulations that either calculate the results of laser cooling with a pulsed laser or probe a sample of laser cooled atoms. To start, the steps used to simulate laser cooling are briefly laid out. Next, the outcomes of these simulations are discussed. Finally, the feasibility of two methods for determining the degree to which cooling has been effective are explored. One of the principle uses for cooled atoms, precision 1S-2S spectroscopy, is discussed in the next chapter

### 9.1 Methods

We start each simulation with an atom from the uncooled distribution (see 8.4) corresponding to the trapping fields being used. A symplectic integrator with a step size of  $1 \mu\text{s}$ , as discussed in 8.2.5 and 8.3.1, advances the kinematic state of the atom. If a point outside of the trapping volume is sampled by the stepper, the atom is assumed to have left the trap and the simulation ends. After each 1/(laser repetition

rate), we check whether the atom is in the laser beam, defined to extend to a radius of twice its waist (8.3.3).

If the atom is found to be in the beam, we integrate the density matrix equations (Eqns. 7.56-7.59) to determine whether the atom decays. More specifically, we follow these steps.

1. Randomly select  $\lambda_P, \lambda_e \in [0, 1)$ .
2. Use local quantities ( $\vec{B}$ , intensities, etc) to calculate  $\Omega(t)$ ,  $\Gamma$ , and  $\gamma$ .
3. With the atom starting in the ground state, integrate Eqns. 7.56-7.59 over the duration of the laser pulse.
4. If  $P < \lambda_P$  (the atom has decayed) before the end of the pulse is reached, simulate a decay and return to step 1. Step 3's integration will only cover the remaining time in the pulses.
5. If  $\rho_{ee}/P > \lambda_e$  (the atom ends in the excited state) after integration has finished, simulate a decay.

To simulate decays, we follow another sequence of steps.

1. Randomly select a ground state, weighted by the branching ratio from the excited state of the transition of interest.
2. Change the atom's momentum by that of the absorbed photon.
3. Randomly choose the emitted photon's direction from the distribution corresponding to the decay chosen in step 1.

4. Change the atom's momentum by that of the emitted photon.

If the atom decays to a different ground state than the one from which it was excited, the final state is recorded and the simulation ends.

As was hinted at earlier, we also estimate the frequency of unintended transitions to determine whether the single-transition calculation is a reasonable approximation. Every time the effects of an incident pulse are calculated for the intended transition, they are also calculated for every unintended transition. These extra calculations also treat their transitions in isolation. Consequently, the primary transition does not deplete the ground state or interfere in any way. The calculated frequency for unintended transitions thus represents an upper bound. It is also important to note that the internal and kinematic states of the atoms are not affected by the results of the extra calculations.

For each simulation, a few laser-related choices are made. The transition to be driven is chosen along with the detuning from this transition. The number and direction of the laser paths through the trap must also be selected. While these parameters vary between simulations, others remain constant. The quantities held constant, along with their values, are listed in table 9.1. These reflect the parameters of the pulsed Lyman- $\alpha$  source we constructed for laser cooling in CTRAP [33].

## 9.2 Laser Cooling Simulation Results

In this section, we lay out the results of the laser cooling simulations. Based on the discussion in 7.2.3 and our choice to focus on the  $|\psi_{-,-}^{1S}\rangle$  ground state, calculations in

Quantity	Default Value
Laser Repetition Rate	30 Hz
Laser FWHM	40 MHz
Pulse Length	16 ns
Pulse Energy	12 nJ
Beam Waist	3 mm
Beam Focus	$(0, 0, z_{B_{min}})$
Beam Polarization	$\hat{x}$ or $\hat{y}$

Table 9.1: Simulation default parameters

this section are done for the  $|\psi_{-,-}^{1S}\rangle \rightarrow |\psi_{-,-}^{2S}\rangle$  transition. Using the parameters introduced in Table 9.1, we perform each test in the quadrupole, octupole, and octupole with sideports traps. For each of these trap configurations, cooling is simulated with light propagating down the x-axis, the z-axis, and all three axes. Results are compared to find the optimal frequencies and laser configurations for maximal cooling in each trap. For these parameters, we extract cooling rates and asymptotic energies.

### 9.2.1 Spectra

For each trap and laser configuration described above, the simulation runs for an hour at a variety of laser frequencies. The resulting mean temperatures are plotted against the laser detuning. Here, detunings represent  $f_{laser} - f_{transition}$  for the field at the minimum of the trap in question. Particles heated out of the trap do not contribute to plotted mean temperatures.

We plot the spectra after 10 minutes of cooling (Fig. 9.1) as well as after an hour of cooling (Fig. 9.2). A few observations stand out.

- After an hour of cooling, the average atomic energy varies only slightly over a  $\sim 200$  MHz range around each spectrum's minimum. Cooling on this timescale is thus robust against errors of  $\sim 14$  mT when determining the trap's minimum field.
- For the octupole trap with sideports, the minimum for z-axis laser cooling is offset from the x-axis and 3-laser cooling minima. Fig. 3.7 explains this, showing that the minimum field in the axial laser's path (waist = 3 mm) corresponds to a Zeeman shift of about 150 MHz higher than the minimum field in the transverse laser's path.
- The quadrupole trap differs from the two octupole traps in that multiple lasers are required for substantial cooling.

The last point, regarding the differences between cooling in the octupole and quadrupole traps, deserves further discussion. From the three-laser simulation, it appears possible to cool antihydrogen in the quadrupole trap. Therefore, the problem is likely a lack of coupling between independent degrees of freedom. To test this, we plot the mean kinetic energy as a function of time for each degree of freedom. This is a coarse way to assign energies to degrees of freedom, but the effects of cooling are nevertheless apparent.

In Fig. 9.3 we see that, in the quadrupole trap, cooling along a single axis has only a small effect on other axes. For cooling along the x-axis, subfigure (a) shows

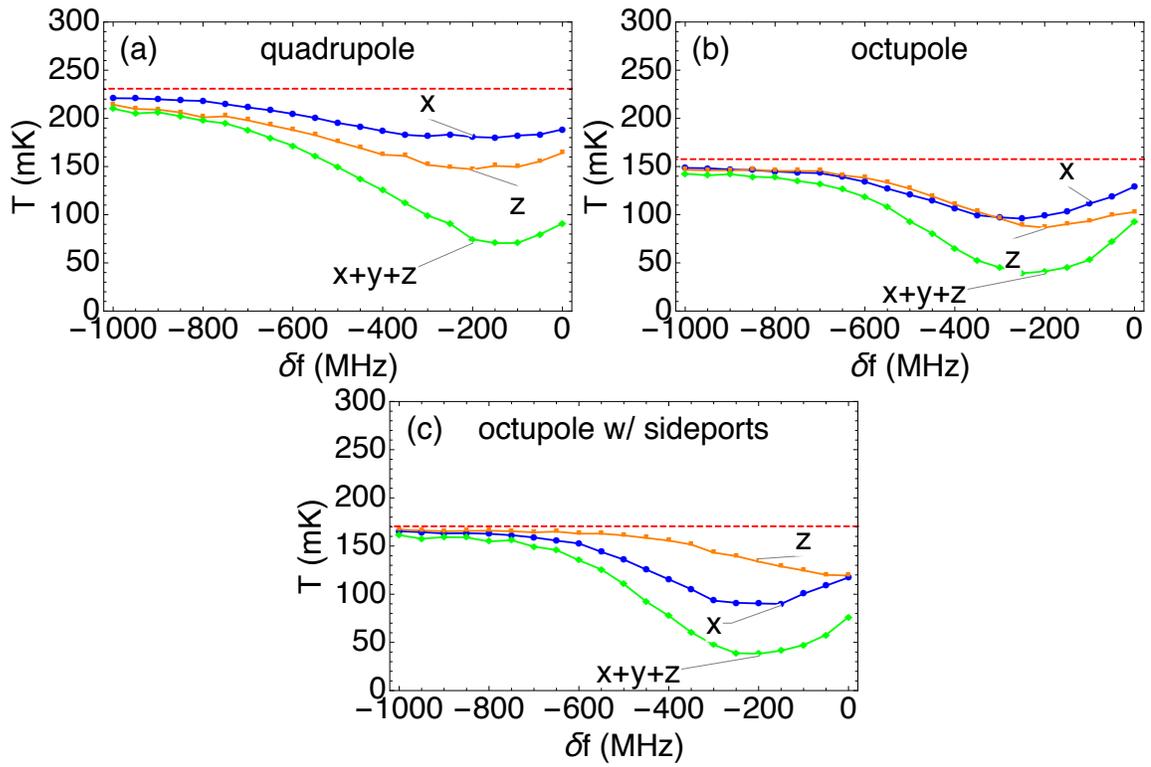


Figure 9.1: Mean temperatures vs. detuning (relative to resonance at the trap minimum) are shown after 10 minutes of cooling. The dashed red lines denote the initial mean energies of the atoms while the x, y, and x+y+z labels give the axes along which lasers propagated to generate each curve.

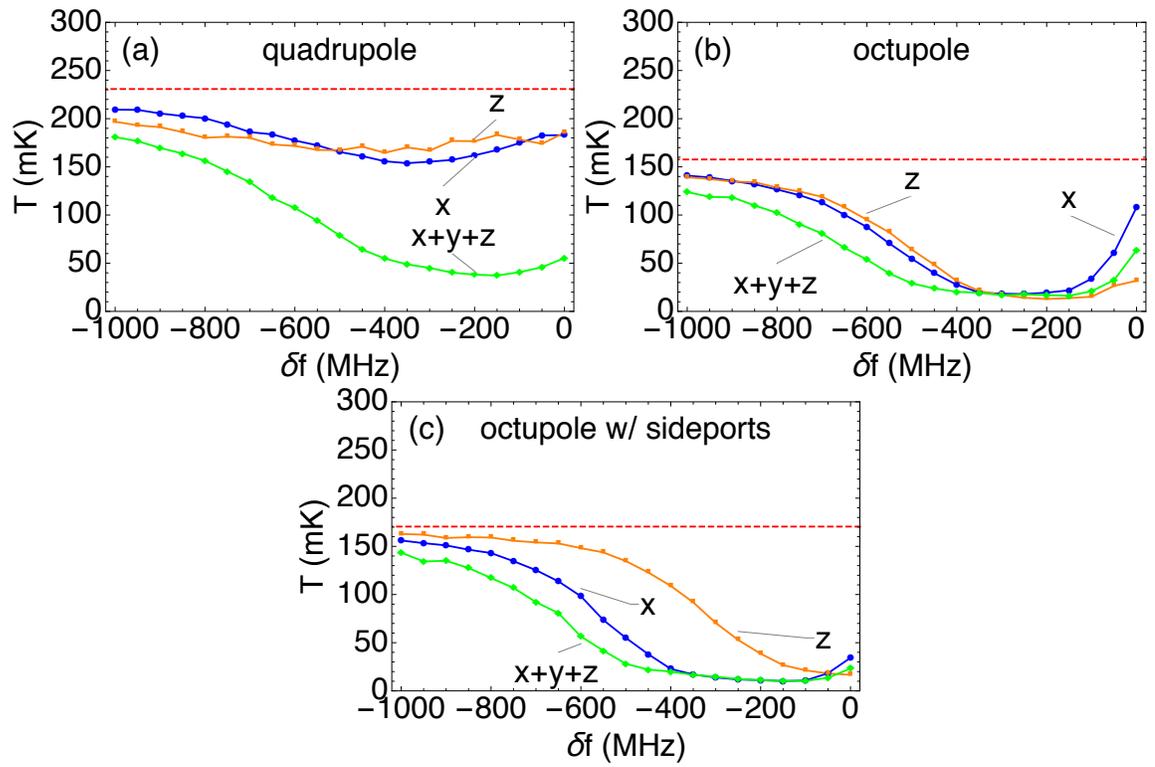


Figure 9.2: Mean temperatures vs. detuning (relative to resonance at the trap minimum) are shown after 60 minutes of cooling.

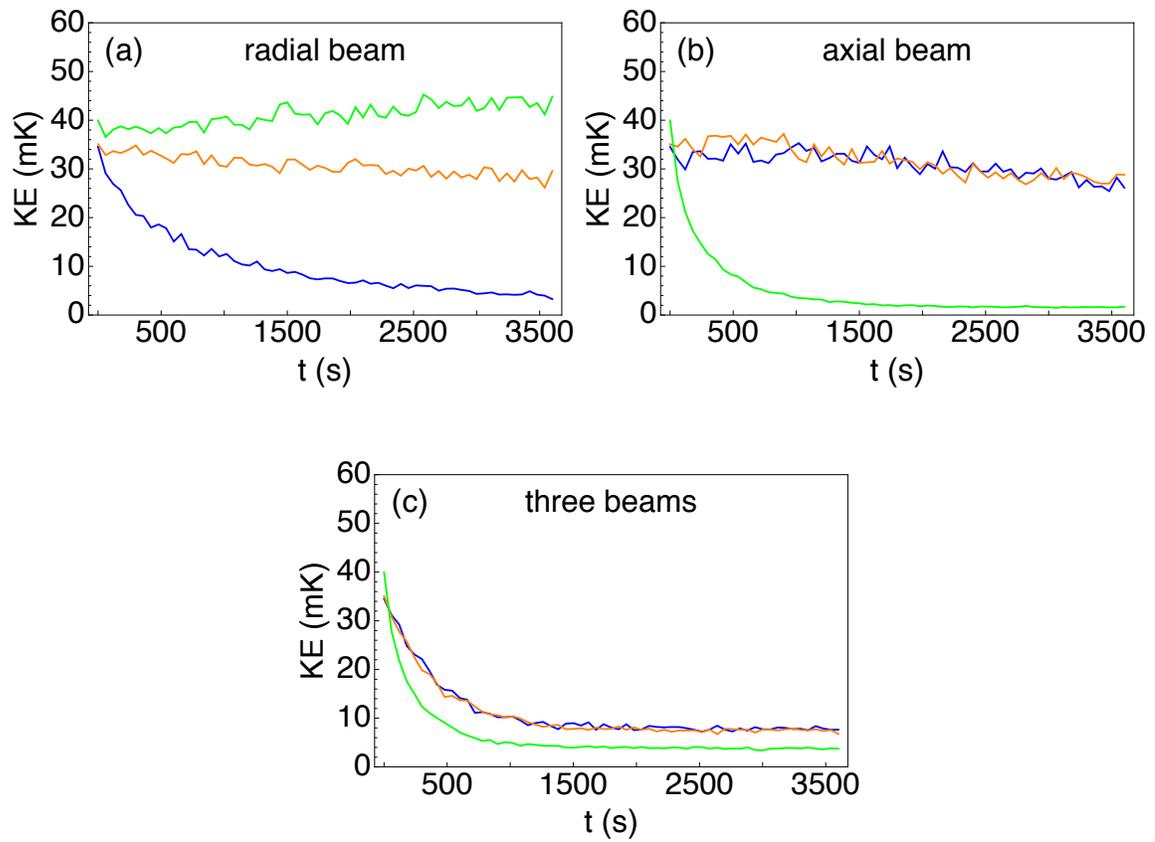


Figure 9.3: Cooling in a quadrupole trap (-150 MHz detuning) along different axes. Kinetic energy in the x direction is blue, y direction is orange, and z direction is green.

that the y degree of freedom is gently cooled and the z degree of freedom is slightly heated. Part (b) shows that z-axis cooling causes the x and y degrees of freedom to cool to a small degree. In each case, no orthogonal degree of freedom is cooled quickly enough for a one-axis cooling scheme to be useful. On the other hand, (c) shows clear cooling along all three degrees of freedom when three lasers are used.

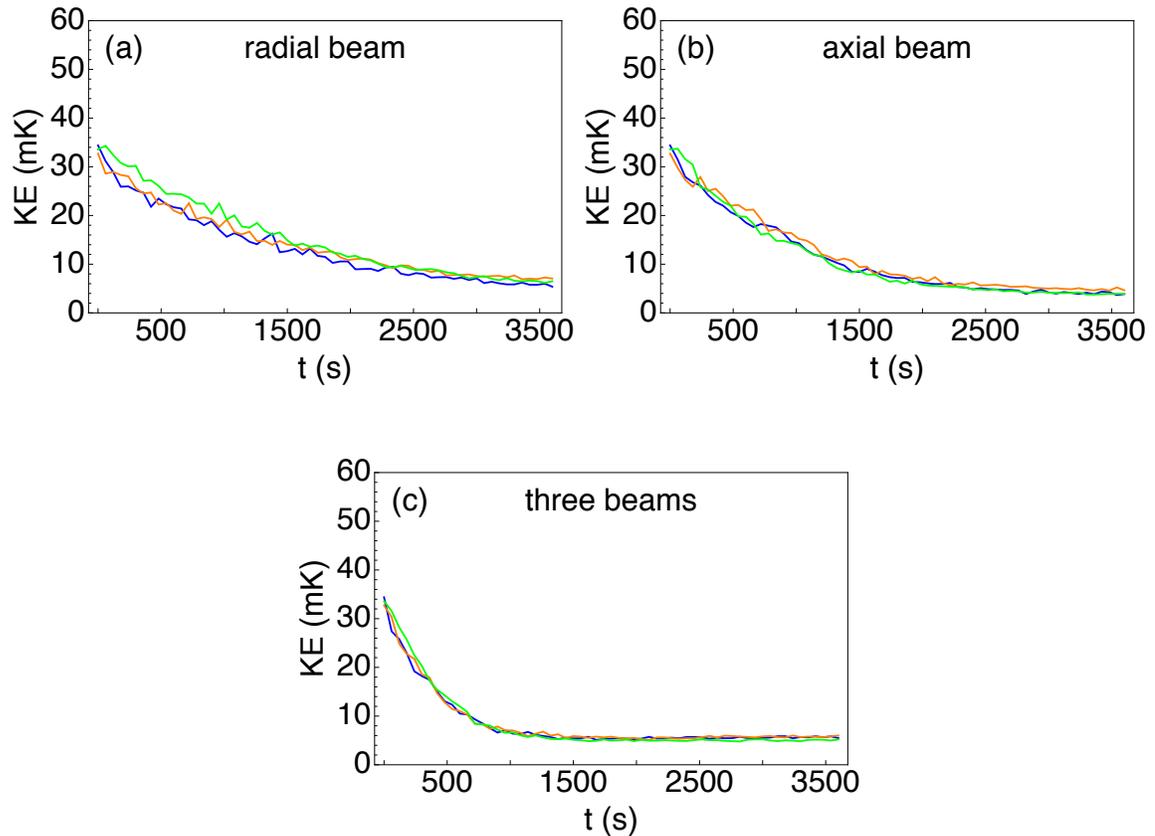


Figure 9.4: Cooling in a octupole trap (-150 MHz detuning) along different axes. Kinetic energy in the x direction is blue, y direction is orange, and z direction is green.

Fig. 9.4 shows the same data as Fig. 9.3, but for the octupole trap. The corresponding graphs for the octupole trap with sideports look similar. Both one-axis

and three-axis cooling effectively lower the energies for all three degrees of freedom. We see, then, why single-axis cooling is effective in the octupole traps but not in the quadrupole trap.

### 9.2.2 Asymptotic Behaviour

It is difficult to determine the best detuning for cooling from the plots in Figs. 9.1-9.2. The minima after 10 minutes of cooling appear to be slightly different than those after an hour of cooling. However, the flatness of the one hour spectra near their minima ensures that any reasonable choice of detuning should give near-optimal cooling. In what follows, we set the detuning to -150 MHz for all three traps.

In Fig. 9.5, we plot mean energies during two hours of laser cooling in several laser and trap configurations. Where cooling is successful, two hours appears to be sufficient time for reaching asymptotic energies. As noted before, we see that atoms in the quadrupole trap require three lasers to effectively cool while atoms in either of the octupole traps require only one laser. It is also interesting to note that x-axis and z-axis cooling in the octupole traps lead to different final energies. In both cases, three-laser cooling converges to the same asymptotic energy as does x-axis cooling.

It is useful to extract some numbers from these plots. We fit each curve to

$$T_f + A \exp -t/\tau. \tag{9.1}$$

Except for the data from simulations of single-axis cooling in the quadrupole trap, the exponentials match the curves well. We also take the average of the mean temperature over the final minute of each simulation, reporting the value as  $T_f^{avg}$ . Extending this

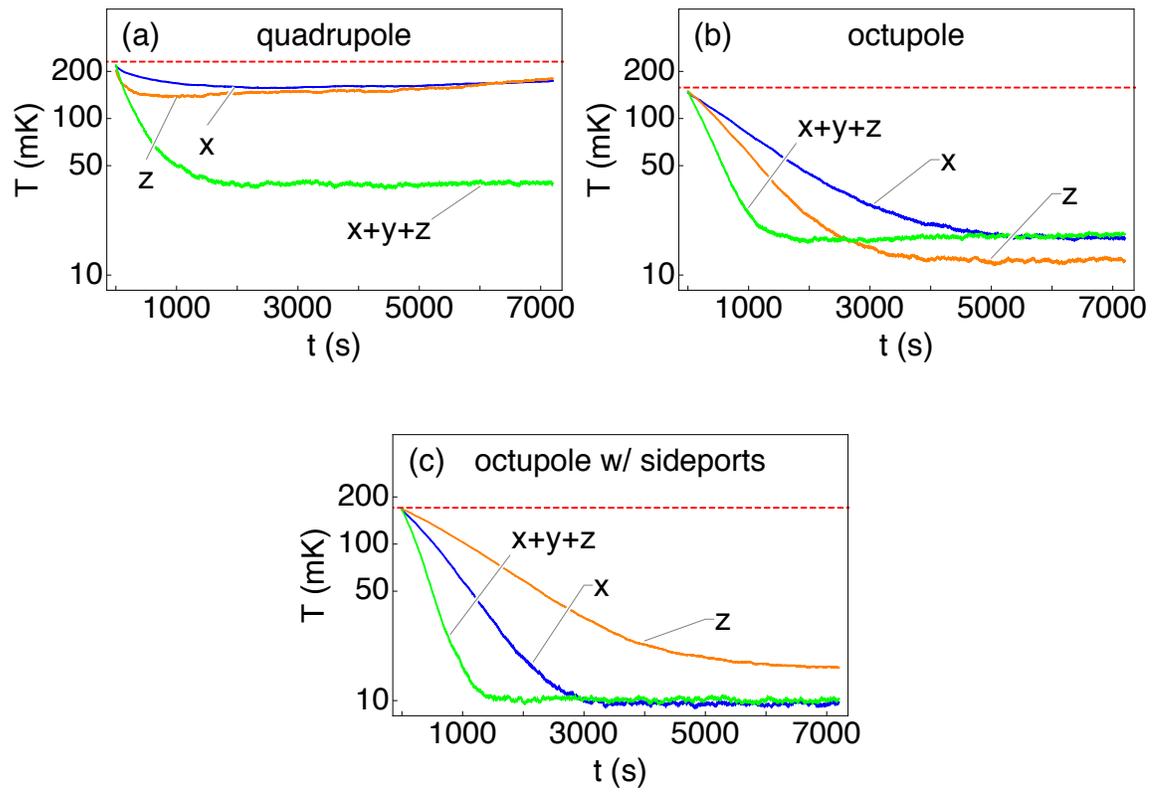


Figure 9.5: Laser cooling with a a detuning of -150 MHz. Labels indicate the laser axes used for cooling

time to 10 minutes changes the results by less than 0.2 mK. The averages and fit parameters are reported in Table 9.2.

Trap	Laser Axes	$T_f^{avg}$ (mK)	$T_f$ (mK)	$\tau$ (s)
Quadrupole	x	-	-	-
Quadrupole	z	-	-	-
Quadrupole	x+y+z	38.8	38.1	360
Octupole	x	17.2	16.0	1262
Octupole	z	12.5	12.0	805
Octupole	x+y+z	18.3	17.2	358
Octupole w/ SP	x	9.6	9.3	669
Octupole w/ SP	z	16.3	14.7	1388
Octupole w/ SP	x+y+z	10.1	10.0	308

Table 9.2: Fit parameters for Fig. 9.5

Finally, it is worthwhile to try to understand an earlier observation. As noted above, the asymptotic mean energy for three-laser cooling in either octupole trap is nearly the same as that for laser cooling along the x-axis. It differs sharply from the asymptotic mean energy for laser cooling along the z-axis. To understand why this might be the case, consider Fig. 9.6. We see that as cooling accelerates, atoms spend significantly more time in the path of a transverse laser than they do in the axial path.

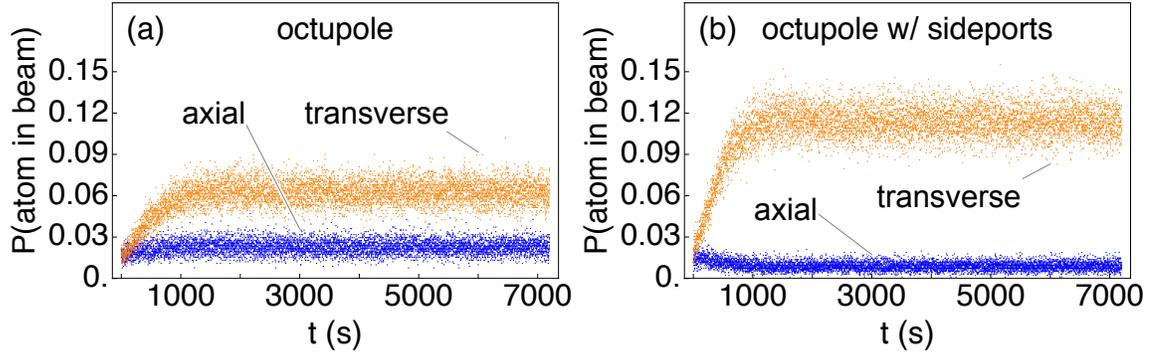


Figure 9.6: Time dependence of the fraction of antihydrogen atoms in the waists of the axial and transverse beams during three-laser cooling

### 9.2.3 Losses and Unintended Transitions

Ideally, a laser cooling experiment would consist of driving a single, closed transition until all trapped atoms have acceptably low energies. At least three effects prevent this picture from being perfectly realized.

1. During their random walk toward lower temperatures, some atoms will be heated out of the trap.
2. The cooling laser will occasionally drive an unintended transition.
3. The excited state in the cooling transition for  $|\psi_{-,+}^{1S}\rangle$  decays to untrapped ground states with a nonzero probability.

As we are just working with the  $|\psi_{-,-}^{1S}\rangle$  ground state, we address only the first two problems.

The fractions of the initial 1000 atoms lost by heating in the simulations above are shown in Table 9.3. As expected, single-axis cooling in the quadrupole trap drives out a large fraction of the atoms. It is disappointing to see losses around ten percent

for the working cooling schemes, but it's important to remember that many of the atoms considered have energies close to the depth of their confining trap.

Trap	Laser Axes	P(heating loss)
Quadrupole	x	0.184
Quadrupole	z	0.599
Quadrupole	x+y+z	0.084
Octupole	x	0.104
Octupole	z	0.071
Octupole	x+y+z	0.087
Octupole w/ SP	x	0.046
Octupole w/ SP	z	0.027
Octupole w/ SP	x+y+z	0.046

Table 9.3: Probabilities for particle loss due to stochastic heating during two hours of laser cooling

Earlier, we discussed a method for bounding the number of unintended transitions during a simulation. The bounds recorded in Table 9.4 were calculated for each of the cooling simulations shown in Fig. 9.5. It is worth noting that no simulated atom experienced more than three unintended transitions during its two hours in the trap.

### 9.3 Cooling Detection Strategies

One of the main motivations for cooling trapped antihydrogen atoms is to narrow the 1S-2S lineshape. However, we should be able to test whether cooling has

Trap	Laser Axes	P(non-cooling transition)
		Upper Bound
Quadrupole	x	0.005
Quadrupole	z	0.020
Quadrupole	x+y+z	0.076
Octupole	x	0.003
Octupole	z	0.003
Octupole	x+y+z	0.019
Octupole w/ SP	x	0.009
Octupole w/ SP	z	0.003
Octupole w/ SP	x+y+z	0.025

Table 9.4: Bounds on probabilities for at least one excitation to an undesired state during two hours of laser cooling

taken place without requiring an additional laser. Furthermore, having the ability to roughly measure the energy of the trapped atoms would be useful for ensuring that a cooling scheme is reproducible. This section explores two possible methods for realizing these goals: dumping the magnetic trap and inducing spin flips to untrapped states using the Lyman- $\alpha$  laser.

In order to study these methods, we need cooled and uncooled populations of trapped antihydrogen. The initial conditions for the cooling simulations give uncooled sets of atoms. For the cooled samples, we take the surviving atoms from our best cooling simulation in each trap. Concretely, for the experiments shown in Fig. 9.5, we use the final states of the z-axis laser cooling simulation for the octupole trap and the final states of the three-axis laser cooling simulations for the other two traps. Thus, for each trap we can compare an uncooled population of atoms to a near-optimally cooled population.

It's important to point out that neither of the methods presented here are yet developed to the level necessary for application to experimental data. Instead, the results of this section serve as proofs of principle for the techniques in question. Indeed, more simulations than time permitted are necessary to rigorously interpret data from cooled or uncooled atoms.

### **9.3.1 Dumping the Trap**

One way to confirm that laser cooling has been successfully carried out is to ramp down the trap and observe the timing of antihydrogen annihilations. Using the measured coil currents during a trap dump, we get the loss profiles shown in Fig. 9.7.

The clear separation between loss peaks makes it likely that cooled and uncooled samples can be distinguished, even with fewer than 1000 atoms.

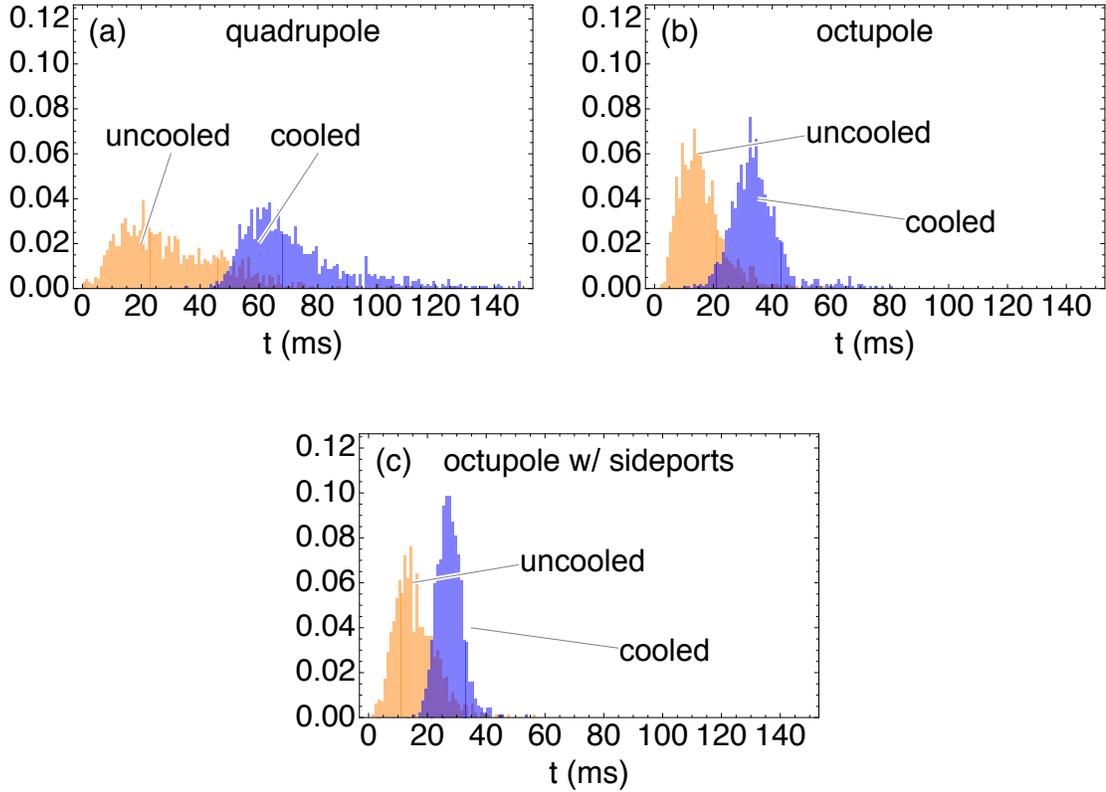


Figure 9.7: Comparisons between cooled and uncooled loss profiles during trap rampdowns

The separation of peaks makes it clear that one sample is cooler than the other. To have a more quantitative temperature diagnostic, we compare initial particle energies against the times they come out of the trap. Discrete data points are used to construct regions giving 90% confidence intervals for an atom’s initial energy as a function of the time at which it hits a trap wall. Plots are made (see Fig. 9.8) for normal CTRAP dumps ( $\tau_0$  of 38.7 ms for the quadrupole trap and 11.8 ms for the octupole traps) as well as dumps slowed by factors of 10 and 100.

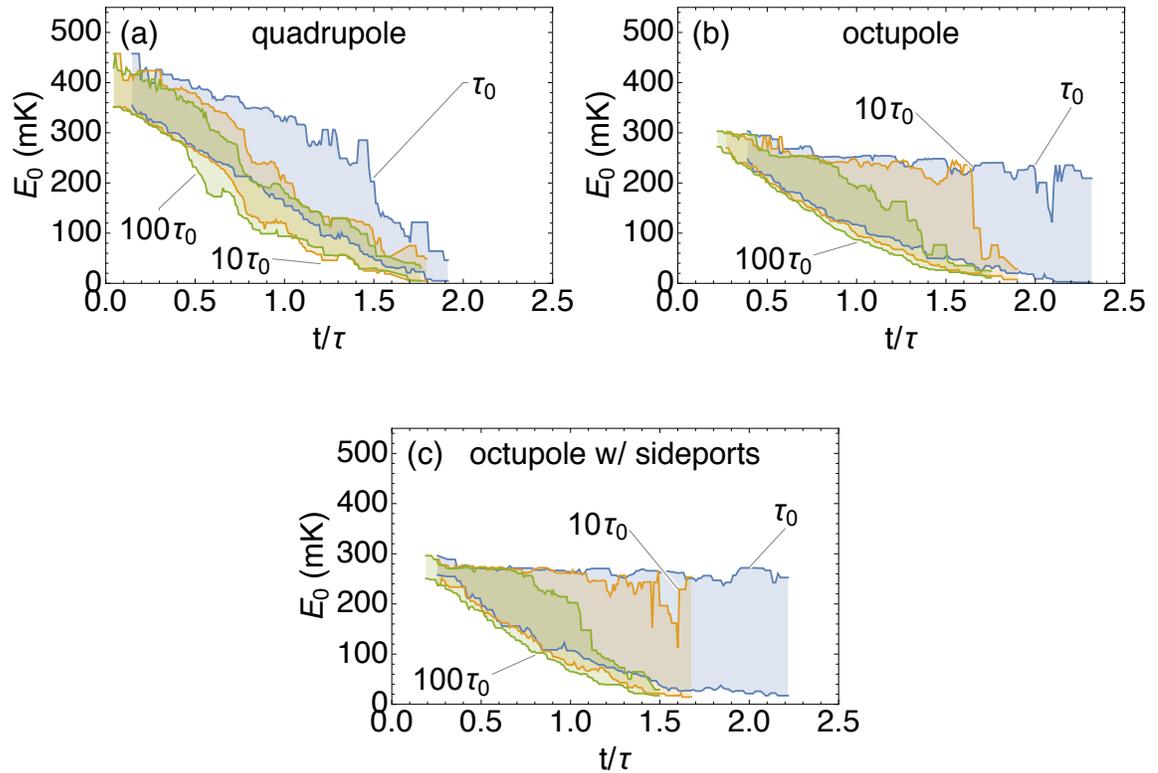


Figure 9.8: The regions shown in these plots give starting energy ranges (to 90% confidence) of particles leaving the trap as a function of time. Each region represents a different exponential decay time,  $\tau$ , for the radial confining field, which is expressed as a multiple of the fastest ramp time achievable,  $\tau_0$ .

It appears that it might be possible to translate losses vs. time in a trap dump to an energy spectrum. However, the widths of the confidence intervals promise to add significant uncertainty to any distribution constructed in this way. Slowing the ramp appears to tighten the error bounds, but they remain significant for the ramp rates we have plotted. Presumably, the errors would drop below some specified limit if we continued to slow the trap dumps. Unfortunately, the speed of the simulation prevents much more exploration in this direction.

### 9.3.2 Laser Removal

We mentioned in 7.2.3 that two transitions hold promise for removing antihydrogen atoms from the trap. In this subsection, we study both of these transitions and their potentials for differentiating between cooled and uncooled antihydrogen. As usual, we assume atoms begin in the state  $|\psi_{-,-}^{1S}\rangle$ .

To start, consider the probability of ejecting an atom from the trap as a function of the laser frequency. Since the goal is to quickly eject cooled atoms, cooled samples are used to construct each spectrum. The results of the simulations, which run for 20 s to limit saturation broadening of the peaks, are displayed in Figs. 9.9 and 9.10.

From the generated spectra, we choose frequencies that should efficiently clear the trap of cooled atoms. For the quadrupole trap, we choose a detuning from the trap minimum of -500 MHz. For the octupole traps we choose a detuning of -100 MHz, except in the case of excitation to  $|\psi_{+,-}^{2P}\rangle |-\frac{1}{2}\rangle$  in the trap with sideports, where we use -150 MHz. One might argue that these frequencies are not quite optimal, but the breadths of the peaks in Figs. 9.9 and 9.10 suggest that the results should not be too

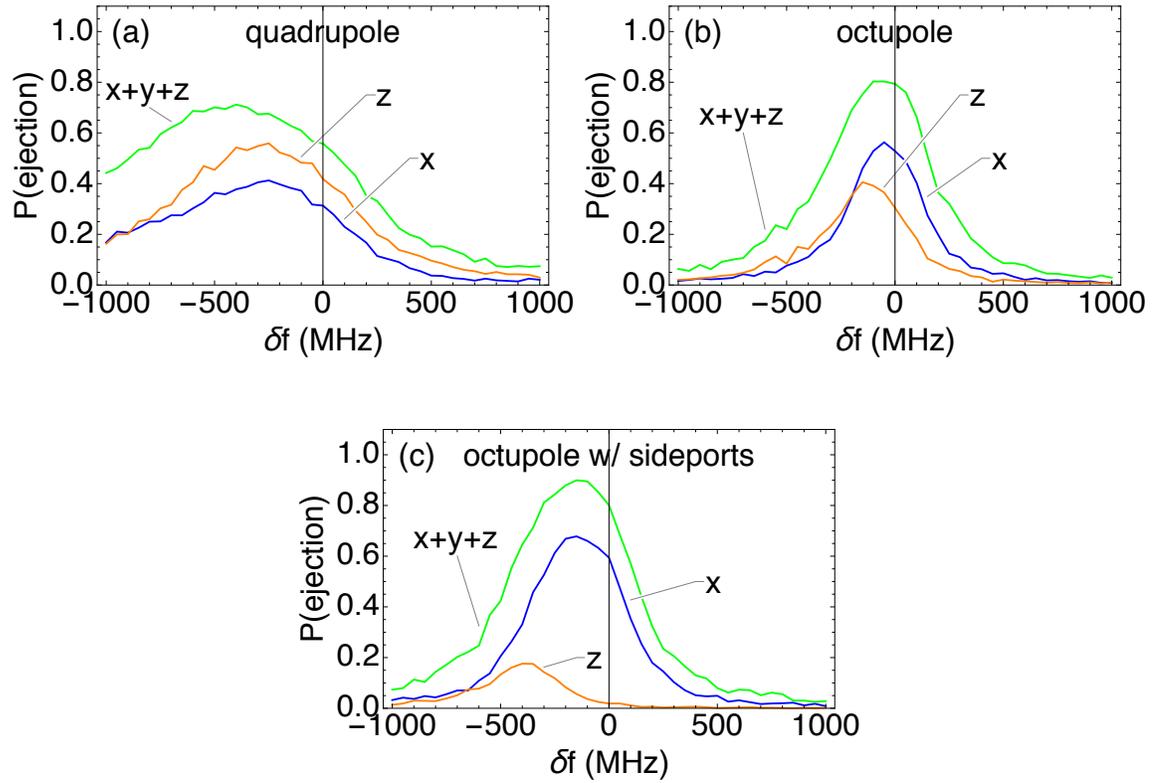


Figure 9.9: Ejection spectra for the  $|\psi_{-,-}^{1S}\rangle$  to  $|\psi_{0,+}^{2P}\rangle |-\frac{1}{2}\rangle$  transition

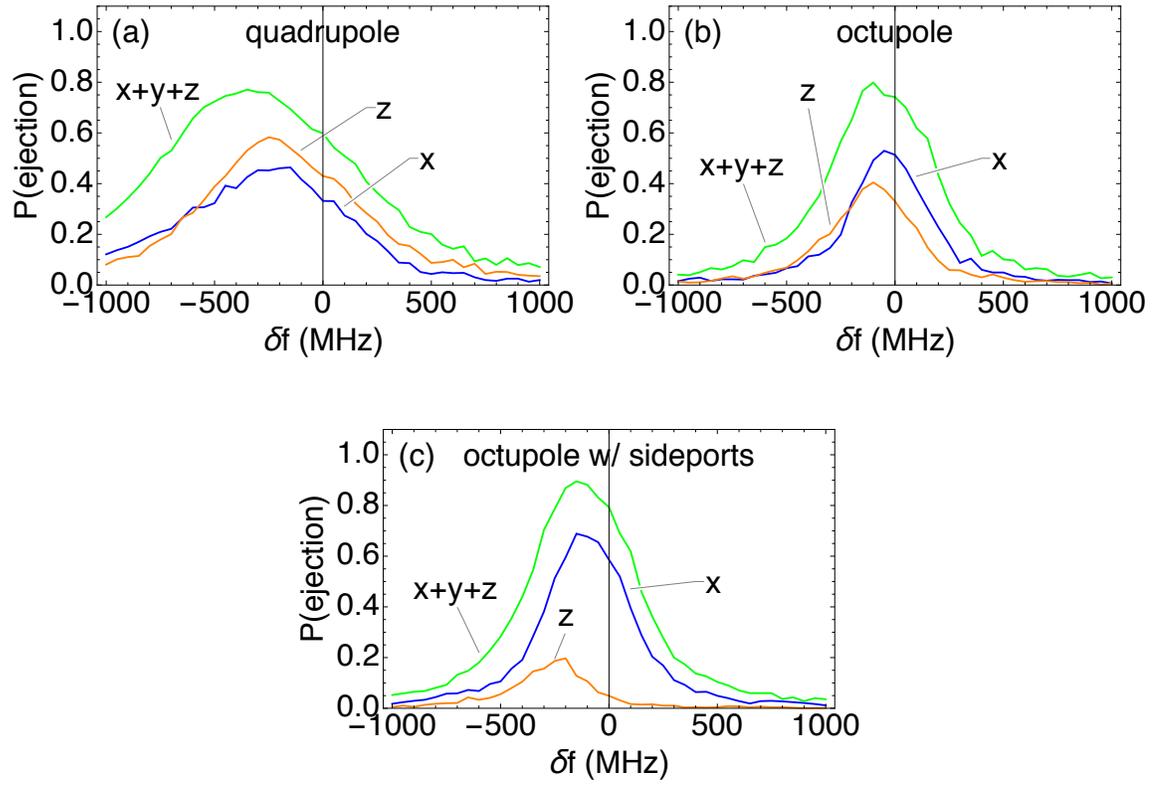


Figure 9.10: Ejection spectra for the  $|\psi_{-,-}^{1S}\rangle$  to  $|\psi_{+,-}^{2P}\rangle |-\frac{1}{2}\rangle$  transition

sensitive to the precise frequency choices.

With lasers tuned to these frequencies, we allow samples of cooled and uncooled atoms to propagate in each of the three traps for 600 s, keeping track of when atoms leave the traps. The results of simulations are shown in Figs. 9.11-9.13.

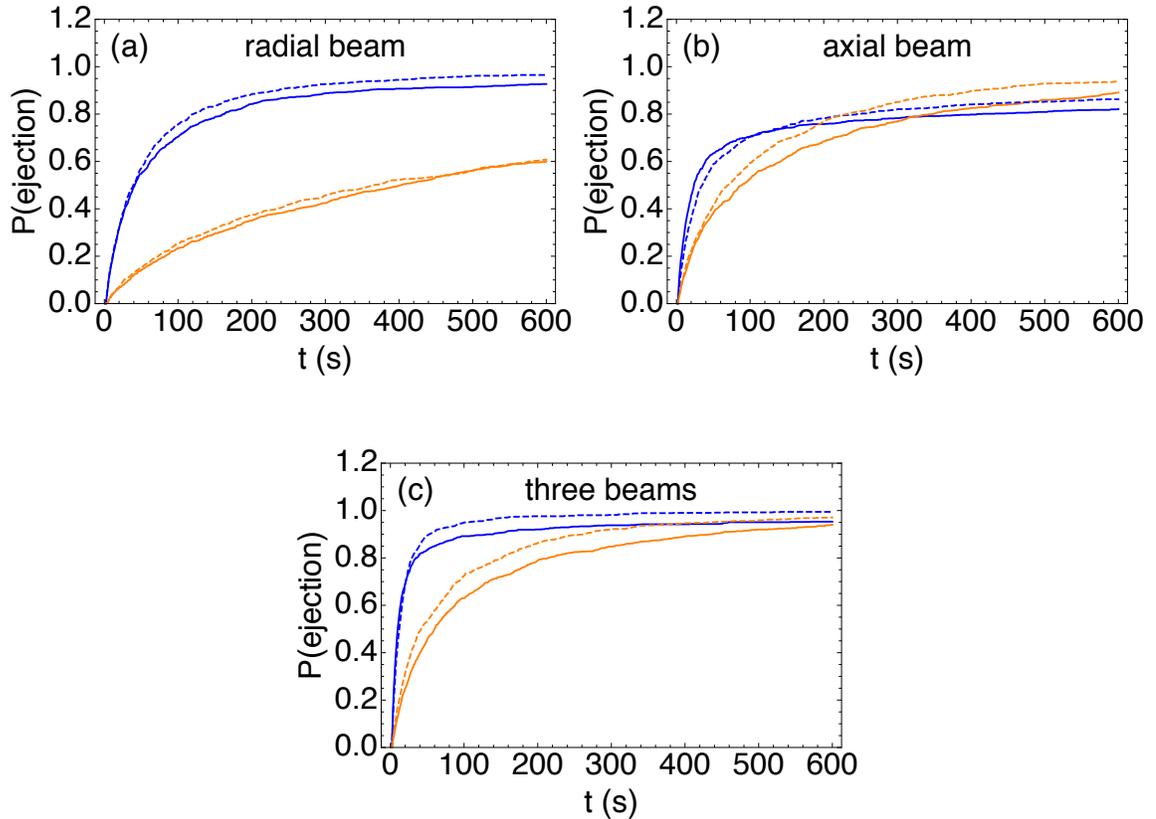


Figure 9.11: Spin flip probability via  $|\psi_{0,+}^{2P}\rangle$  (solid) or  $|\psi_{+,-}^{2P}\rangle$  (dashed) for cooled (blue) or uncooled (orange) atoms in the quadrupole trap

From these plots, we can draw a few conclusions.

- The choice of how long to illuminate the atoms is critical, and ought to depend on the trap and laser configurations.
- It appears that a single radial laser is the best configuration to use. Three laser

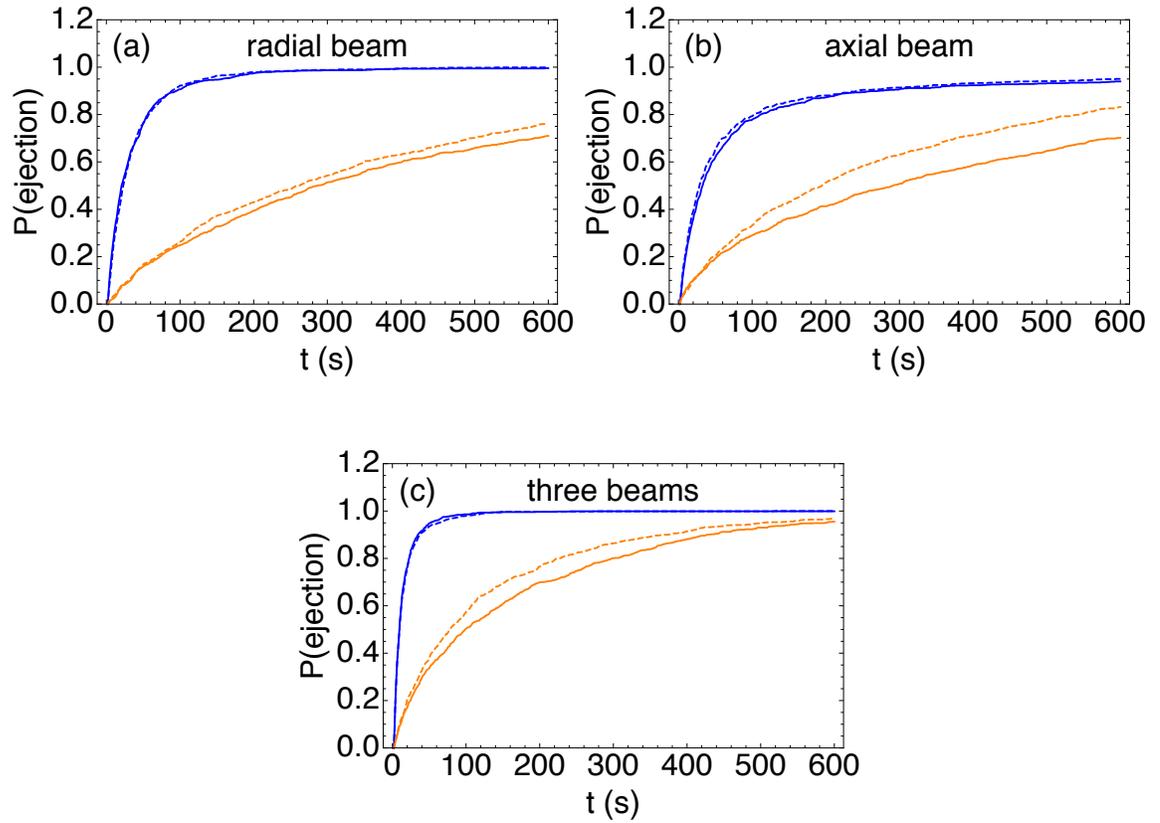


Figure 9.12: Spin flip probability via  $|\psi_{0,+}^{2P}\rangle$  (solid) or  $|\psi_{+,-}^{2P}\rangle$  (dashed) for cooled (blue) or uncooled (orange) atoms in the octupole trap

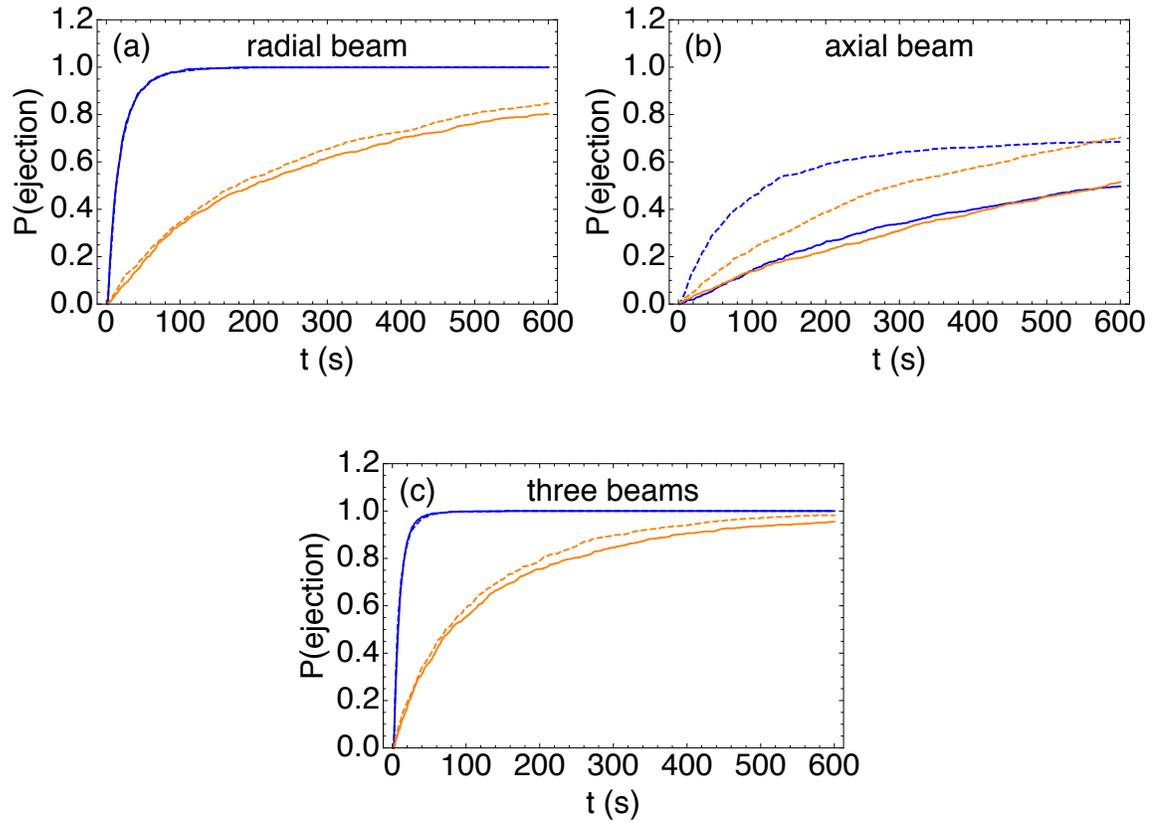


Figure 9.13: Spin flip probability via  $|\psi_{0,+}^{2P}\rangle$  (solid) or  $|\psi_{+,-}^{2P}\rangle$  (dashed) for cooled (blue) or uncooled (orange) atoms in the octupole trap with sideports

configurations close the gap between cooled and uncooled losses quickly while a single axial laser provides less contrast between cooled and uncooled particles than does a radial laser.

- If the experiment time and the laser configuration are chosen wisely, the choice of which transition to use doesn't seem to make much of a difference.

With detectors possessing high enough signal-to-noise to see individual annihilations as they happen, we can push a bit further and use the number of atoms lost over time to very roughly determine the temperature of our cooled sample. Fig. 9.14 shows a data set that might make this possible. The plots for the other transition look similar. While Fig. 9.14 orders loss curves of different energies as we might expect, it's important to recognize certain limitations when comparing these to an experimental loss curve. For instance, an experiment will generate a loss curve using atoms from an unknown energy distribution. Comparing this to loss curves generated by atoms in narrow energy ranges may hint at the mean and bounds for the unknown distribution, but firm conclusions cannot be established. Even so, this technique may prove to be a useful tool for comparing the efficacy of different cooling schemes.

Finally, we pointed out earlier that spin-flipping transitions might be difficult to model accurately because of other nearby transitions. However, we find across the simulations run for this study that at most 0.6 percent of the atoms would undergo an unintended transition. This is a small enough fraction to not affect the general conclusions here.

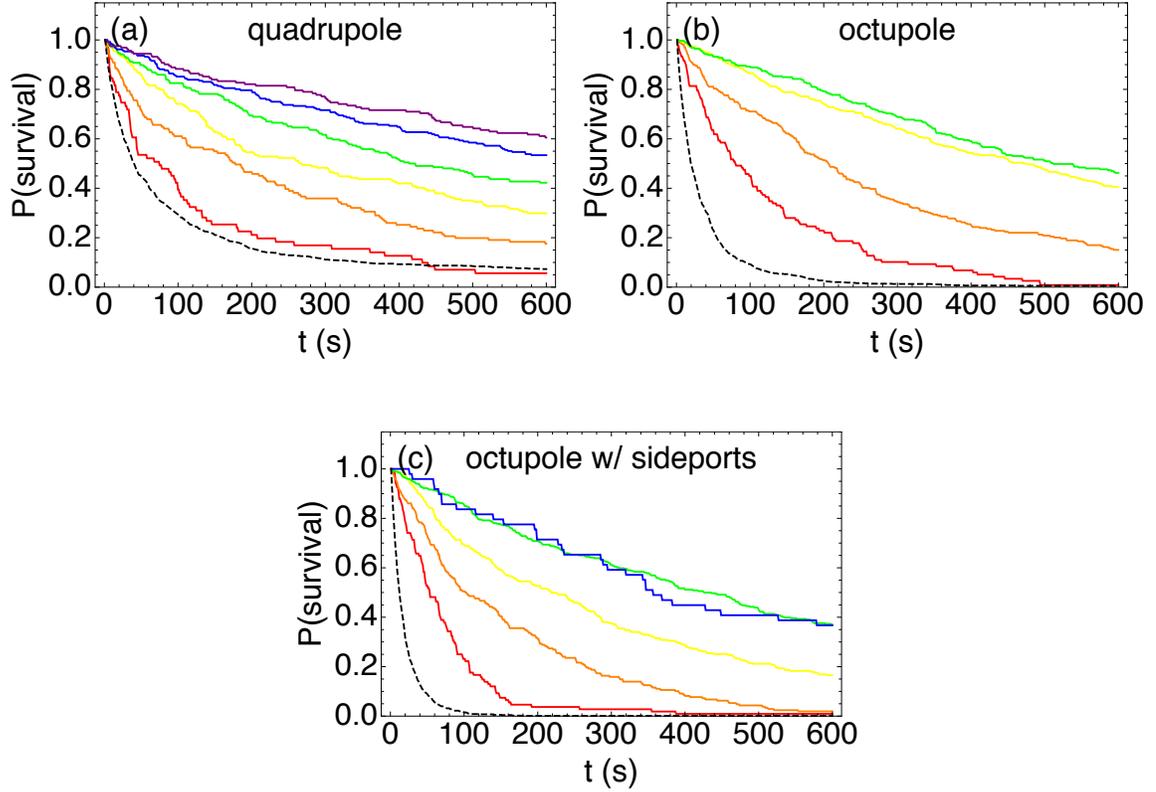


Figure 9.14: Shown are the losses over time for  $|\psi_{-,-}^{1S}\rangle$  to  $|\psi_{0,+}^{2P}\rangle |-\frac{1}{2}\rangle$  spin-flip experiments with the laser directed down the x-axis. Each solid line represents particles in a specified 100 mK energy range from an uncooled sample. The  $[0, 100)$  mK atoms are represented by red curves, and successive ranges are colored in rainbow order. Each dotted black line represents the particles in a cooled sample.

## 9.4 Summary

The simulations in this chapter shed a great deal of light on what's possible with magnetically trapped antihydrogen atoms and a pulsed Lyman- $\alpha$  source.

The primary result is the quality and speed of cooling possible in a magnetic trap. Cooling times, as defined in Table 9.2, range from 5 to 20 minutes. For three-laser cooling near the optimal frequency, distribution mean energies are approximately at their asymptotic values after 30 minutes of cooling. Asymptotic temperatures are, for the best laser path configurations, approximately 38 mK for the quadrupole trap and 10 mK for the octupole traps. The free-atom Doppler limit of 2.4 mK remains out of our reach, but this is expected given the non-negligible laser width and the significant effects potentials can have on cooling [95].

A second important result is that techniques exist, using only the trap and the cooling laser, for probing the energies of the atoms. The feasibility of two independent approaches has been demonstrated. Consequently, once a reproducible antihydrogen production and cooling method has been established, the two energy measurements can be compared. This cross-check can be used to build confidence in these techniques.

Finally, the parameter space exploration performed in the simulations described above is extremely valuable experimentally. While broad features of our results should generalize to other antihydrogen traps of similar shape, the information about cooling times, optimal cooling laser frequencies, and the paths that yield the lowest energies in each trap saves our experiment weeks to months of trial and error.

In addition to providing insight on laser cooling, the simulations described in this chapter are useful for generating initial conditions for 1S-2S spectroscopy. By

cooling atoms before performing spectroscopy, broadening due to finite transit times, the Zeeman effect, and other less important mechanisms is reduced. A narrower lineshape generally increases the precision of the measurement. Narrowing of the 1S-2S line due to cooling will be demonstrated in the next chapter.

# Chapter 10

## 1S-2S Spectroscopy Simulations

One of the ultimate goals of trapping antihydrogen atoms is to compare the 1S-2S transition frequency with that of H. The most accurate measurement of this quantity in hydrogen was carried out in a beam experiment, yielding an uncertainty of 10 Hz [18]. For antihydrogen, all 1S-2S spectroscopy experiments to date have been performed in magnetic traps, with the most recent quoting an uncertainty of 5.4 kHz [3].

One cause of the 540-fold uncertainty gap between hydrogen and antihydrogen measurements is the broad  $\sim 100$  kHz FWHM lineshape generated by trapped antihydrogen. This compares unfavorably to the 2 kHz linewidth seen in hydrogen beam experiments. Much of the width is due to transit time broadening, although the Zeeman shift, ionization width, and other mechanisms also contribute. Cooling the atoms, as discussed in the previous chapter, ought to reduce broadening due to transit-time and Zeeman effects, along with other less important mechanisms.

In this chapter, the results of simulations of 1S-2S spectroscopy in CTRAP's

magnetic trap are presented. In particular, we look at the lineshapes generated by cooled and uncooled atoms. Spectroscopy in the latter case has been simulated for trapped antihydrogen [96], but the effects of cooling on the spectra have not been previously addressed. As was the case in the previous chapter, these simulations don't just clarify what's possible with the CTRAP apparatus. They give insight on the laser parameters and illumination times we ought to use. Optimizing these based on simulations, especially those requiring atoms to be pre-cooled, should save weeks to months of experiment time.

## 10.1 Methods

Simulations of 1S-2S spectroscopy handle kinematic evolution similarly to their laser cooling counterparts. The position and velocity of an atom is advanced by a symplectic stepper with a  $1 \mu\text{s}$  step size, as in 9.1. However, because the spectroscopy laser is continuous rather than pulsed, a particle's position is checked each step to see whether it is within the region in which the beam is relevant.

Once in the beam, the time step is adjusted so the atom moves no more than one percent of the beam waist per step. This ensures that intensity-dependent quantities (Rabi frequency, ionization rate, etc) can be approximated as constants when solving equations of motion. In the beam, Eqns. 7.60-7.62 are effectively integrated to advance the internal state each kinematic step. After each integration,  $\rho_{gg} + \rho_{ee}$  is compared with a threshold, chosen either at the beginning of the simulation or after the last decay. If the threshold exceeds the probability that the atom is in the two-level system, the atom is assumed to decay or ionize and the threshold is reset to a random number

in  $[0, 1)$ . As long as the atom decays to the ground state in which it began, internal and kinematic evolution continues.

When an atom leaves the beam, the internal state evolution can be described by multiplying  $\rho_{eg}$  by a complex exponential and  $\rho_{ee}$  by the norm squared of the exponential. For efficiency, we track the sum of the exponent and only perform the exponentiation and multiplication once per millisecond, or when the atom re-enters the beam. For each multiplication, we compare the decay threshold with  $\rho_{gg} + \rho_{ee}$  to determine whether the atom has decayed, resetting its internal state and the threshold if it has.

The in-beam and out-of-beam propagation techniques described above are repeated until the simulated time reaches a predetermined limit or the particle hits a wall. It's important to point out that each simulation simultaneously addresses many laser frequencies. The internal state and kinematics are approximated as decoupled, so we can carry several internal states along with the particle, allowing each to evolve assuming different detunings. When one of these internal states is lost by decay or ionization, we record the time, location, and channel of the loss and prevent that internal state from evolving further.

While this description of the spectroscopy simulations gives a decent overview of how the physics is modeled, a few points deserve explanation: the initial conditions used in the simulations, time-saving simplifications to the equations of motion, and the zero-recoil approximation. Details regarding these subjects are given below.

### 10.1.1 Initial Conditions

Spectroscopy simulations begin by selecting an atom from one of the three distributions: uncooled, cooled, and Doppler limit cooled. The uncooled distribution is described in 8.4 and represents the trapped antihydrogen atoms generated by mixing antiprotons and positrons. The cooled distribution, used in 9.3, is a sample of uncooled atoms laser cooled under optimal conditions until an asymptotic mean temperature is reached. The final distribution, also discussed in 8.4, is constructed by cooling particles from the "cooled" distribution so the mean energy coincides with that in the free-particle Doppler limit. This may not be possible to realize with our Lyman- $\alpha$  cooling schemes, but it's worthwhile to have a sample representing a low-energy limit to compare with the cooling that we are able to accomplish. It's important to remember that each trapping field configuration has its own uncooled, cooled, and Doppler limit cooled distributions. The mean energies for each of the nine distributions used as initial conditions are presented in Table 10.1.

Trap	Uncooled	Cooled	Doppler Limit
Quadrupole	230.8	38.8	2.4
Octupole	157.7	12.5	2.4
Octupole w/ SP	170.5	10.1	2.4

Table 10.1: Mean energies (in mK)

### 10.1.2 Simplifying the Equations of Motion

Solving the equations of motion Eqns. 7.60-7.62 is as central to spectroscopy as it is time-consuming. As a useful simplification, we choose to neglect the laser linewidth. Sub-Hertz widths for the 972 nm precursor to the 243 nm light have been demonstrated [97], so this may be justified. However, even if  $\Gamma_{las}$  contributes significantly to  $\gamma$  in practice, it's still of interest to separate broadening mechanisms due to the laser's finite width from those due to the details of the performing spectroscopy in a magnetic trap.

The advantage of this approximation is that a solution to the simple equations

$$\dot{c}_g = \frac{i\Omega}{2}c_e \quad (10.1)$$

$$\dot{c}_e = (i\delta - \frac{\Gamma}{2})c_e + \frac{i\Omega}{2}c_g. \quad (10.2)$$

generates a solution to Eqns. 7.60-7.62 through the identification  $\rho_{ab} = c_a \bar{c}_b$ . These equations can be easily solved as follows.

$$\lambda = \frac{1}{2}(-\frac{\Gamma}{2} + i\delta) \quad (10.3)$$

$$\mu = (\lambda^2 - \frac{\Omega^2}{4})^{1/2} \quad (10.4)$$

$$c_g(t) = \frac{e^{\lambda t}}{\mu} [\mu c_g(0) \cosh \mu t + (-\lambda c_g(0) + \frac{i\Omega}{2}c_e(0)) \sinh \mu t] \quad (10.5)$$

$$c_e(t) = \frac{e^{\lambda t}}{\mu} [\mu c_e(0) \cosh \mu t + (\frac{i\Omega}{2}c_g(0) + \lambda c_e(0)) \sinh \mu t] \quad (10.6)$$

Evaluating this analytic form is considerably faster than numerical integration.

### 10.1.3 No-Recoil Approximation

Finally, let's examine the assumption that photon emissions do not change an atom's velocity. This approximation allows a single trajectory to carry multiple atomic states at once, all responding to lasers of different frequencies. As long as the calculations represented in Eqns. 10.3-10.6 take significantly less time than solving the kinematic equations, this significantly speeds the process of generating spectra.

To demonstrate that this is a reasonable approach, we perform Monte Carlo experiments using scatter data collected while running spectroscopy simulations. We choose data taken at the laser frequencies that maximize the number of one and two-photon decays. For each atom, we repeat the following steps 100 times.

1. Choose the number of one-photon decays assumed to happen during this iteration from the Poisson distribution defined by the rate of observed one-photon decays.
2. Do the same for two-photon decays.
3. For each two-photon decay, randomly choose the total momentum of the two photons using the distributions described in 7.3.1.
4. Take the atom's speed to be the maximum possible value given its energy.
5. For each one-photon decay, add the momentum of a randomly directed Lyman- $\alpha$  photon to the atom's velocity.
6. For each two-photon decay, add the randomly chosen total momentum, directed randomly, to the atomic velocity.

This process, repeated for each atom, builds a distribution approximating that generated by recoils during spectroscopy. Note the assumption that the atom is always traveling at maximum speed and the choice of laser frequencies maximizing scattering bias this test toward overestimating the effects of recoil.

The results of these experiments are collated for each trapping setup and each set of initial conditions. Figs. 10.1-10.3 compare the recoil-affected energy distributions to the unperturbed distributions. It appears that recoil fails to greatly change the energy distribution of the trapped atoms. While susceptible to some criticism, this experiment appears to justify the approximation made by ignoring recoil during spectroscopy simulations.

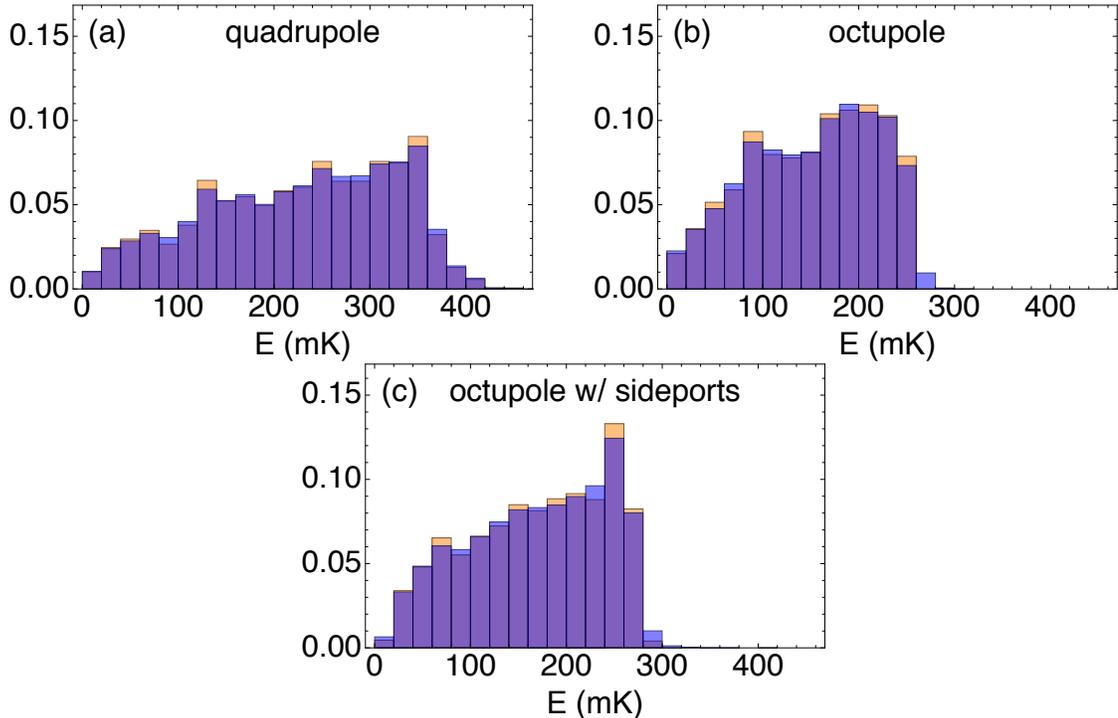


Figure 10.1: Effects of recoil on uncooled energy spectra, with blue representing the spectrum affected by recoil and the orange depicting the unperturbed spectrum

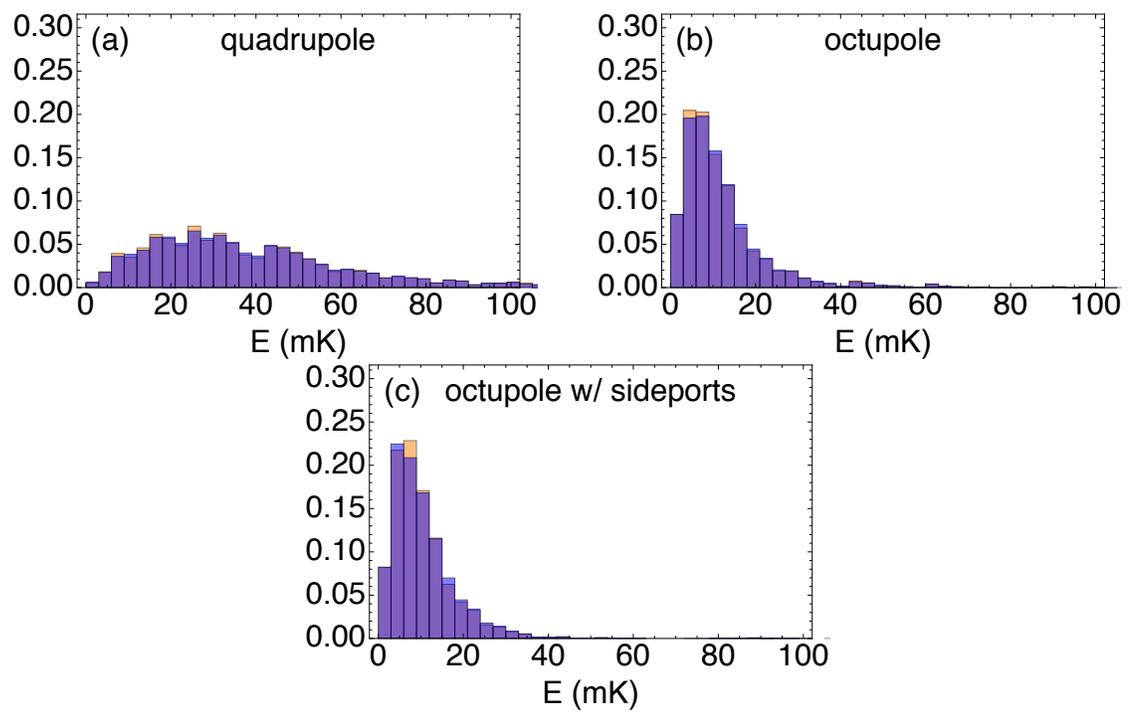


Figure 10.2: Effects of recoil on cooled energy spectra, with blue representing the spectrum affected by recoil and the orange depicting the unperturbed spectrum

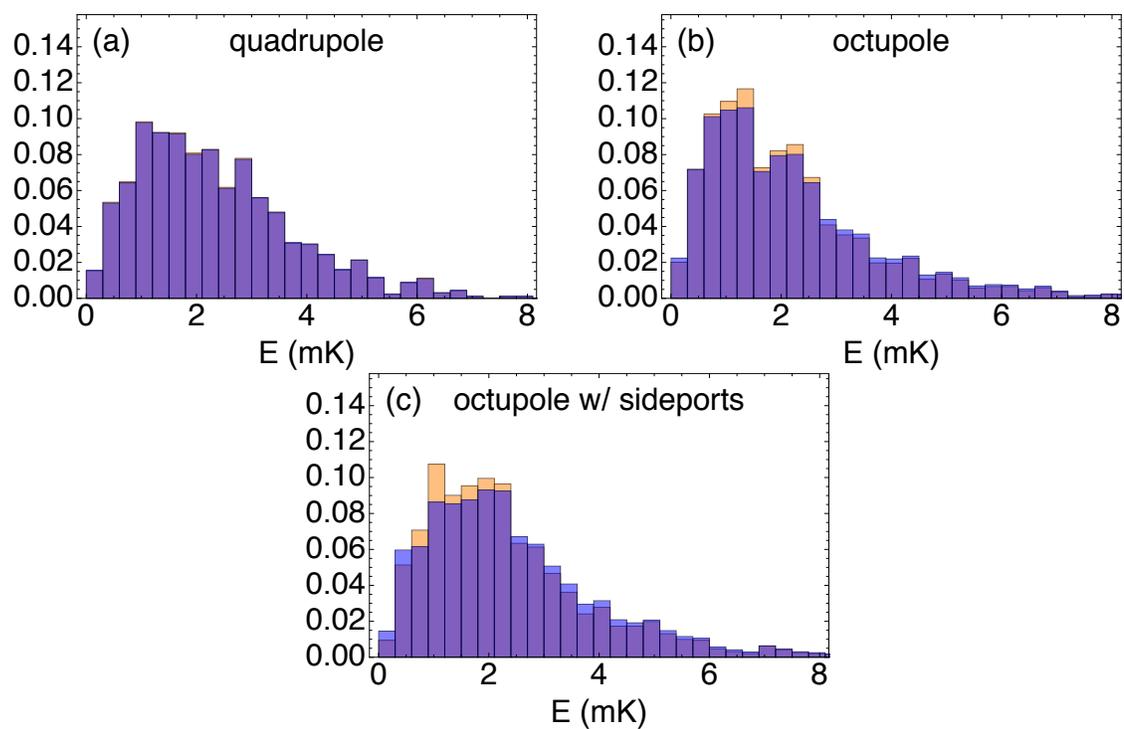


Figure 10.3: Effects of recoil on Doppler cooled energy spectra, with blue representing the spectrum affected by recoil and the orange depicting the unperturbed spectrum

## 10.2 Spectroscopic Signal

In this section, we present the results of 1S-2S spectroscopy simulations under a variety of conditions. For each trap, we vary the laser's waist size, the path taken by the beam, and the temperature of the atomic sample.

The spectra displayed here are produced in similar ways. For each trap, initial particle distribution, and laser state, the 1S-2S spectroscopy simulation runs for one hour of simulated time. In each case, the laser power in each of two counterpropagating beams is set to 1 W and the laser linewidth is taken to vanish. From each simulation, we are able to plot fractional particle losses as a function of time and frequency. Different conditions lead to different loss rates, so each spectrum is shown after a different amount of illumination time. In particular, we capture each spectrum soon after its peak passes 0.3. If no point on the spectrum passes this threshold during the hour-long simulation, we use its final state. Constructing the spectra in this way prevents saturation from broadening the lines and allows for easier comparison between spectroscopy experiments done under different conditions.

Figs. 10.4-10.6 display the results of performing 1S-2S spectroscopy in various traps. To more easily see how sample temperature, laser axis, and waist size affect the time it takes to collect data and the width of the lineshape, these values are plotted in Fig. 10.7. Several interesting observations may be made regarding the data we have collected.

First, it's important to assess what can be done without any cooling. ALPHA observed a  $\sim 50$  kHz FWHM (at 243 nm) using uncooled atoms illuminated for 300 s [3]. It appears that their peak fractional loss was between 0.45 and 0.55, which is

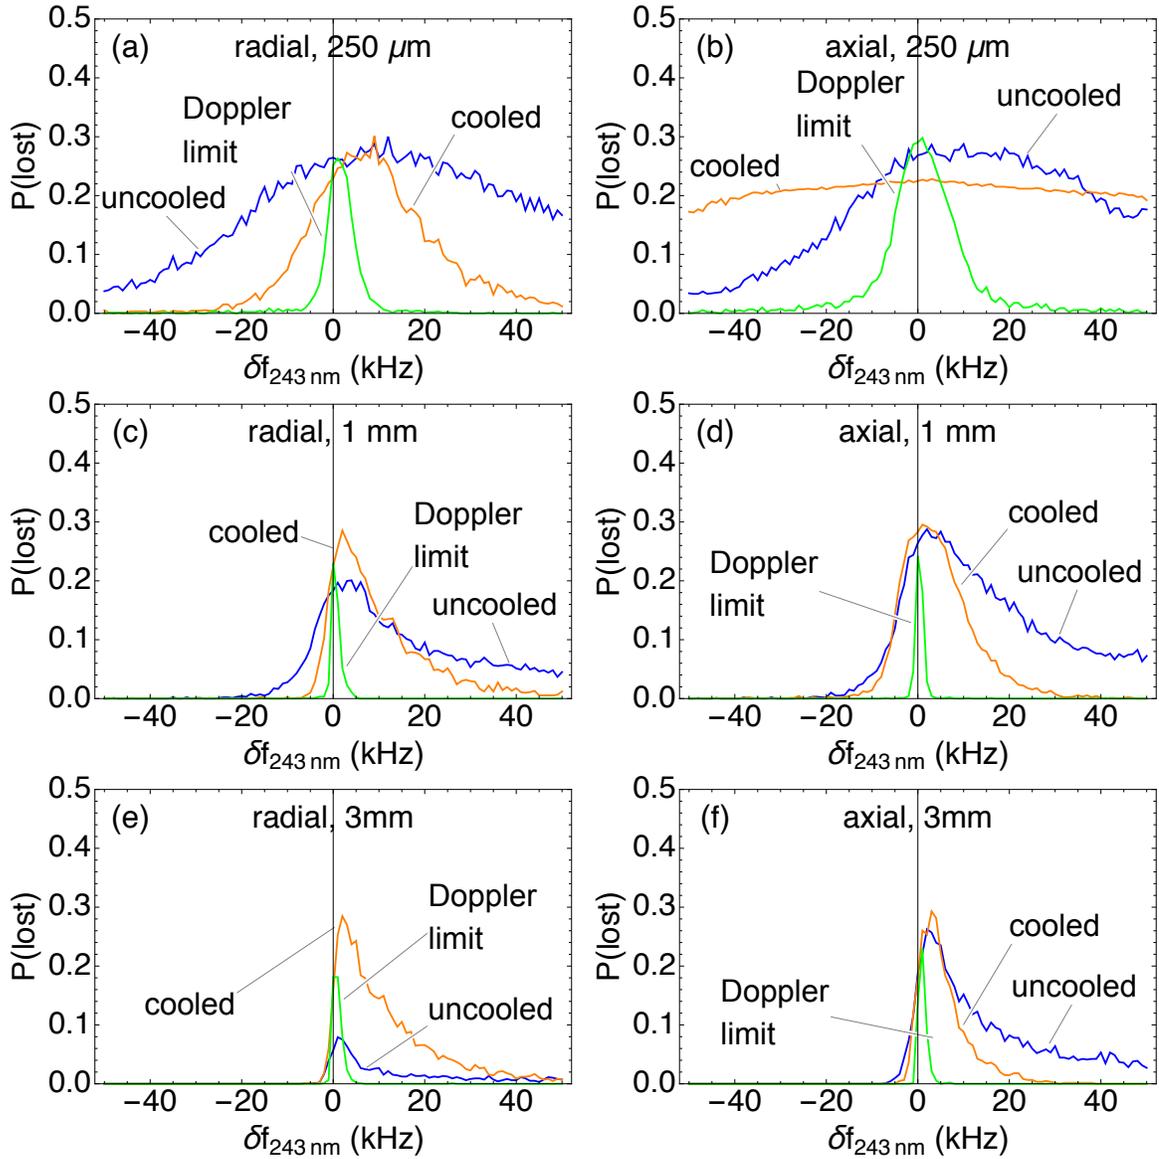


Figure 10.4: 1S-2S spectra in the quadrupole trap, with labels giving the laser's path and the beam's waist

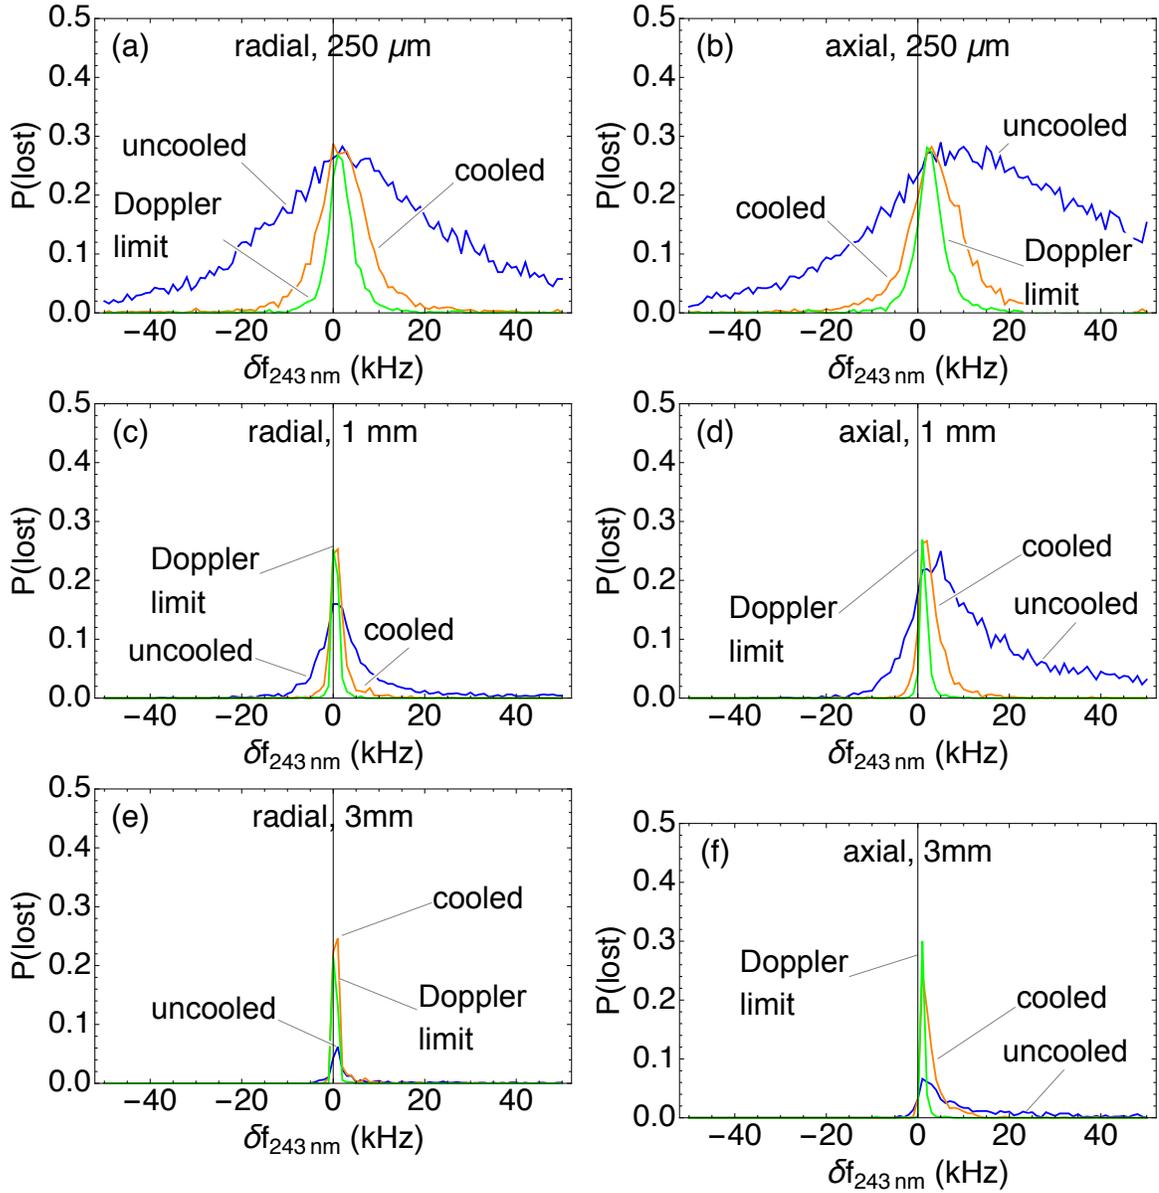


Figure 10.5: 1S-2S spectra in the octupole trap, with labels giving the laser's path and the beam's waist

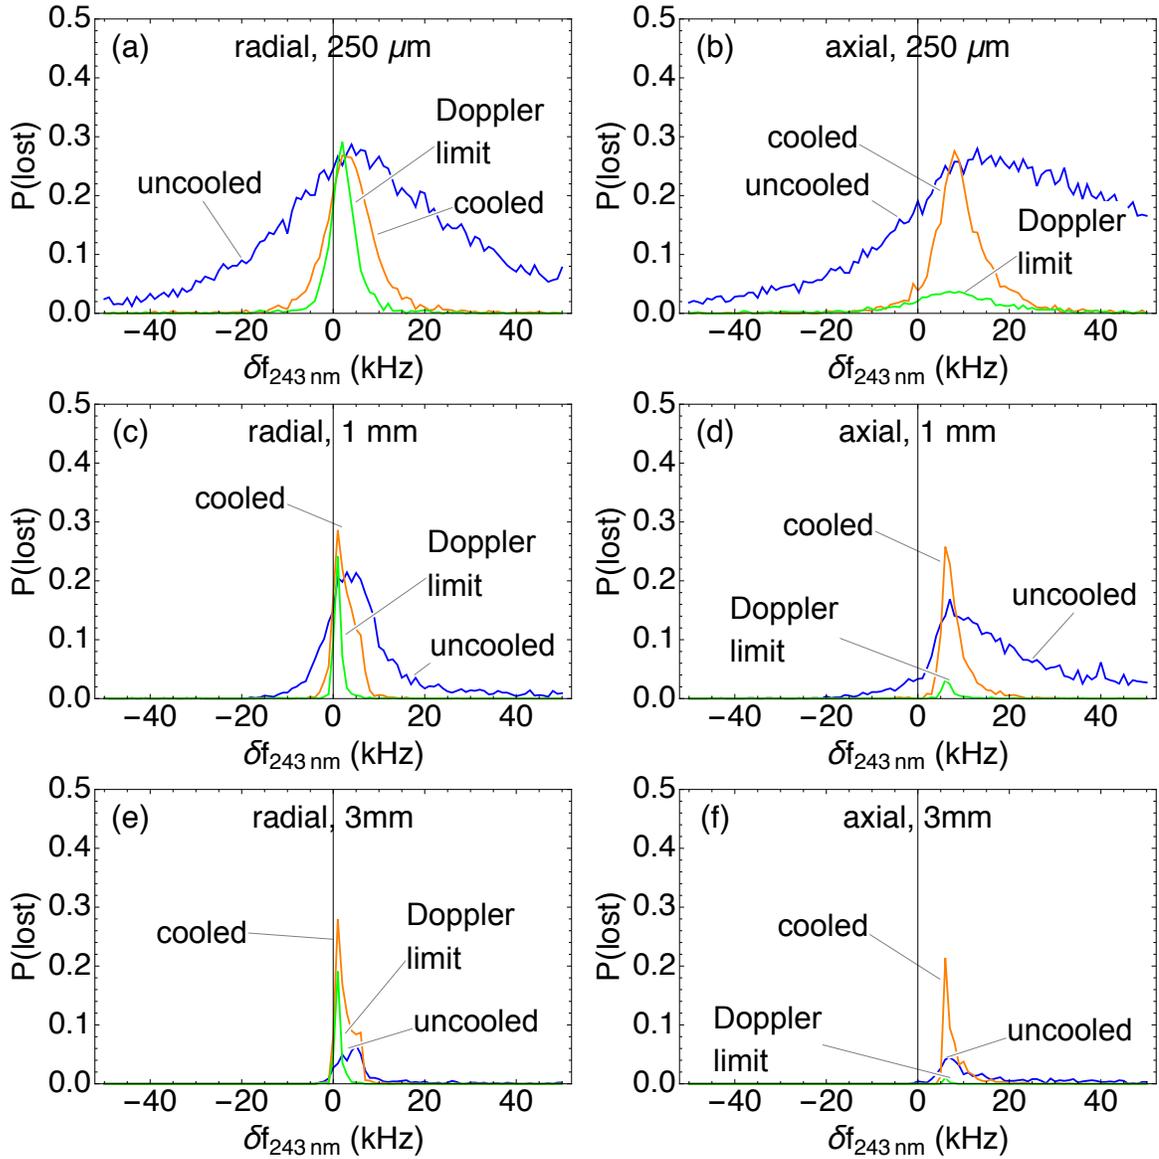


Figure 10.6: 1S-2S spectra in the octupole trap with sideports, with labels giving the laser's path and the beam's waist

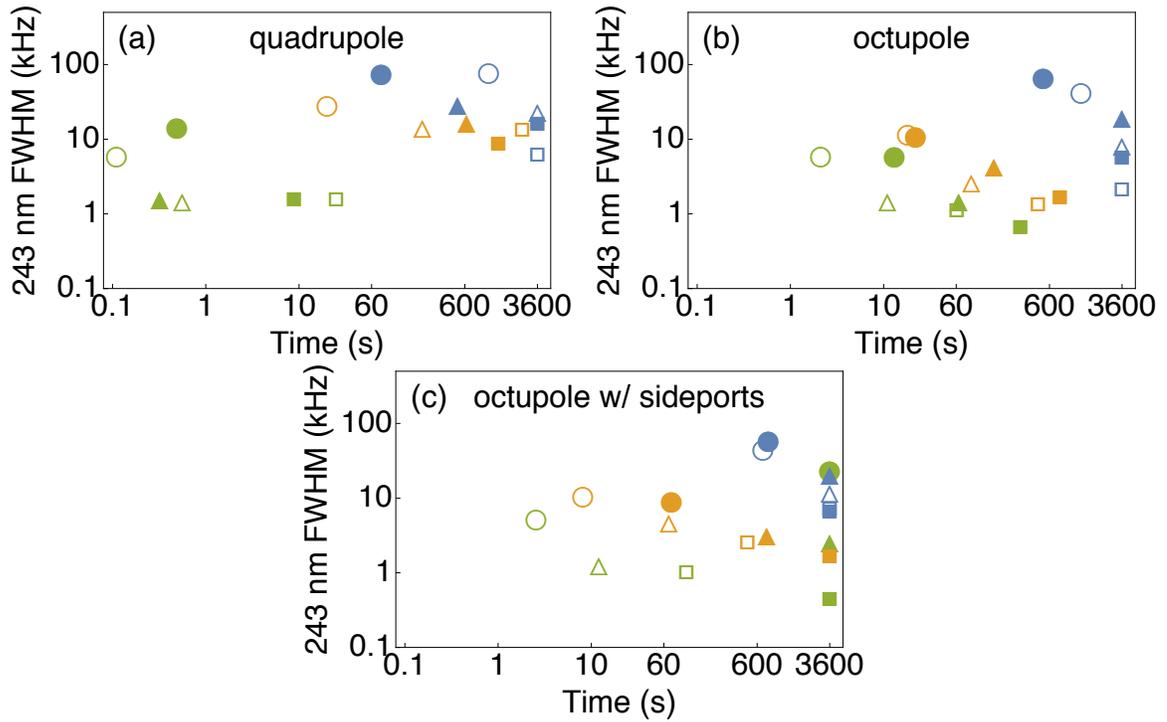


Figure 10.7: For CTRAP's traps, the width of each 1S-2S spectrum is plotted against the time required to collect the data. Blue, orange, and green markers denote uncooled, cooled, and Doppler limit cooled samples. Circular, triangular, and square markers denote beam waists of 250  $\mu\text{m}$ , 1 mm, and 3 mm. Filled markers denote an axial laser while open markers denote a radial beam.

higher than our 0.3 target. While spectroscopy generally proceeds faster in ALPHA's trap, likely due to its extremely flat field, CTRAP is capable of generating spectra from uncooled atoms in less than 15 minutes in any of the trapping configurations under consideration. Of particular interest is the spectrum generated in the quadrupole trap with an axial, 1 mm laser. For an exposure time of less than 10 minutes, a spectrum with about half the FWHM seen by ALPHA may be generated under these conditions.

With cooling, significant narrowing of spectral linewidths is possible. However, it's important to recognize the qualitative difference between cooling in the quadrupole trap and cooling in the octupole traps. From Tab. 10.1, recall that a cooled sample in one of the octupole traps comes much closer to the Doppler limit than one in the quadrupole trap. The consequences of this can be seen in Fig. 10.7 by comparing cooled samples to their Doppler limit cooled counterparts. In the octupole traps, both samples yield narrower lines as the beam waist expands, with the widths due to laser cooled atoms slightly exceeding those due to samples cooled to  $T_D$ . In the quadrupole trap, atoms cooled to the Doppler limit yield widths that plunge to 1.4 kHz while laser cooled spectra only reach 8 kHz.

Another takeaway from this study is the usefulness and sufficiency of the laser cooling schemes we have developed for octupole traps. We ignore spectra that require more than an hour of exposure time. The narrowest linewidths we can get with uncooled atoms, laser cooled atoms, and atoms at  $T_D$ , respectively are 39.4 kHz, 1.2 kHz, and 600 Hz in the octupole trap. For the octupole trap with sideports, these numbers are 41.6 kHz, 2.3 kHz, and 900 Hz. There is a modest reduction

in the time needed to collect data for each spectrum as the atoms get colder, but more importantly, laser cooling the atoms leads to a significant reduction in spectral widths. In fact, the cooled linewidths are close to the 2 kHz FWHM of the line used to determine hydrogen's 1S-2S frequency to 10 Hz [18]. Reducing the sample temperature to  $T_D$  squeezes the linewidths slightly more, but most of the achievable narrowing can be done with the cooling techniques we have already developed.

Finally, we point out the effects of the off-axis wells illustrated in Fig. 3.7. In spectra generated by radial beams with broad waists, we see structure beyond a single peak for both octupole traps. The structure appears to have a width of  $\sim 1$  kHz for the octupole trap and  $\sim 6$  kHz for the octupole trap with sideports. These values correspond to about half the depths of the wells in Fig. 3.7 (half because we're looking at 243 nm spectra). We see the effects of these side wells also in the amount of exposure time needed to construct a spectrum. Fig. 10.7 shows that spectroscopy-induced losses in the octupole traps from cooled and Doppler limit cooled samples proceed slower for an axial laser than they do for a radial beam. This effect is more pronounced in the octupole trap with sideports as the off-axis wells are deeper. This shows up most clearly for the atoms of temperature  $T_D$  in the trap with sideports, as axial beams of the radii we consider induce less than a 5% loss over an hour. In the other two traps, losses above 30% can be generated in a Doppler limit cooled sample in less than 5 minutes.

Finally, it's important to reiterate a point made in 7.1.1. The simulated lineshapes give a good idea of what experimental results should look like and the expected results of changing the laser setup or the addressed atomic distribution. However, it's

unlikely Hertz-level precision is possible by fitting experimentally generated curves to our model. To make such fine comparisons between hydrogen and antihydrogen, both will need to be addressed in the same trapping environment.

### 10.3 Summary

In this chapter, we explored the possibilities for 1S-2S spectroscopy in the CTRAP apparatus. We discovered that even without laser cooling, there exist laser configurations that should allow us to capture spectra much narrower than ALPHA saw in their most recent result. It also became clear how strongly cooling affects the widths of our spectral lines. With the schemes developed in the previous chapter for using a Lyman- $\alpha$  laser to lower the energies of trapped atoms, linewidths similar to those seen in hydrogen beam experiments may be realized. We will certainly have fewer atoms available, but narrowing the observed lineshape promises to improve the precision of our measurement.

These observations reflect both the usefulness of the spectroscopy simulation and the optimism with which we view antihydrogen studies in CTRAP. Although no 1S-2S experiment has yet been done in our apparatus, we have a reasonable expectation of what the results will look like. We also know how to set up the lasers for best results. It's unclear yet how many antihydrogen atoms we will have to participate in each spectroscopic run, as improvements to our antiproton cloud may significantly boost their numbers (Ch. 6). Still, it appears that we are able to significantly reduce the linewidth of the spectra we produce so the precision of the measurement is as high as possible for however many atoms we are able to use.

# Chapter 11

## Conclusions and Discussion

It has been a long, difficult journey, but the new CTRAP apparatus (our nickname for the latest apparatus used by the ATRAP collaboration) has finally begun to fulfill its intended purpose. Although CTRAP's Ioffe magnet promised to be a significant improvement to the one at the heart of the BTRAP apparatus, it took the work described in this thesis to allow the enclosure holding its cryogen supply to survive thermal cycling. I was also a part of construction efforts that made the electrode stack's vacuum enclosure more robust, added laser access to CTRAP, and allowed us to directly image our plasmas. These upgrades will contribute greatly to our success going forward.

We not only built up the CTRAP apparatus, but we also used it to trap antihydrogen atoms. The  $5 \pm 2$  atoms-per-trial trapping rate we observed was somewhat disappointing given the  $5 \pm 1$  atoms per trial seen in BTRAP's 2011 experiments. However, while BTRAP used 1 million antiprotons in their experiments, we used only 360000, while also cutting trial times in half. The 3-fold boost in trapping efficiency is

not so large, but it does indicate that the changes we have made to our antihydrogen synthesis routine in recent years is beneficial. Also, our new plasma imaging system suggests that the anti-atom numbers may be substantially improved by better-shaped antiproton clouds that our recent imaging shows should be possible. More work on this front, helped along by the plasma imaging diagnostic, should yield tremendous gains in antihydrogen production and trapping rates.

In parallel to the experimental work, I also conducted computational studies of laser cooling and spectroscopy. I showed that trapped atoms can be dramatically laser cooled in CTRAP over times comparable to other trap procedures (e.g. antiproton trapping). It also appears that it might be possible to use the same cooling laser to roughly measure the distribution of energies among trapped atoms. Spectroscopy simulations, on the other hand, assure us that the 1S-2S line can be measured over reasonable time intervals without laser cooling. Laser cooling was shown to greatly reduce the widths of the spectra observed, allowing for even more precise measurement. Overall, the simulations give great direction for how best to set up future laser experiments while also assuring us that high-precision 1S-2S spectroscopy is within reach for the apparatus we have.

Despite the progress made and the knowledge gained during this PhD, we have not yet reached our final goal. However, the strides made over the last seven years place the ATRAP collaboration closer than ever to achieving precision 1S-2S antihydrogen spectroscopy. Experimental directions for increasing anti-atom trapping rates and setting up both laser cooling and spectroscopy have become clear. All that is needed is a little more time to pursue these leads.

# Bibliography

- [1] G. Gabrielse *et al.*, Phys. Rev. Lett. **108**, 113002 (2012).
- [2] M. Ahmadi, B. X. R. Alves, C. J. Baker, W. Bertsche, E. Butler, A. Capra, C. Carruth, C. L. Cesar, M. Charlton, S. Cohen, R. Collister, S. Eriksson, A. Evans, N. Evetts, J. Fajans, T. Friesen, M. C. Fujiwara, D. R. Gill, A. Gutierrez, J. S. Hangst, W. N. Hardy, M. E. Hayden, C. A. Isaac, A. Ishida, M. A. Johnson, S. A. Jones, S. Jonsell, L. Kurchaninov, N. Madsen, M. Mathers, D. Maxwell, J. T. K. McKenna, S. Menary, J. M. Michan, T. Momose, J. J. Munich, P. Nolan, K. Olchanski, A. Olin, P. Pusa, C. O. Rasmussen, F. Robicheaux, R. L. Sacramento, M. Sameed, E. Sarid, D. M. Silveira, S. Stracka, G. Stutter, C. So, T. D. Tharp, J. E. Thompson, R. I. Thompson, D. P. van der Werf, and J. S. Wurtele, Nat. Comm. **8**, 681 (2017).
- [3] M. Ahmadi *et al.*, Nature **557**, 71 (2018).
- [4] G. Gabrielse, X. Fei, K. Helmerson, S. L. Rolston, R. L. Tjoelker, T. A. Trainor, H. Kalinowsky, J. Haas, and W. Kells, Phys. Rev. Lett. **57**, 2504 (1986).
- [5] G. Gabrielse, X. Fei, L. A. Orozco, S. L. Rolston, R. L. Tjoelker, T. A. Trainor, J. Haas, H. Kalinowsky, and W. Kells, Phys. Rev. A **40**, 481 (1989).
- [6] G. Gabrielse, X. Fei, L. A. Orozco, R. L. Tjoelker, J. Haas, H. Kalinowsky, T. A. Trainor, and W. Kells, Phys. Rev. Lett. **63**, 1360 (1989).
- [7] G. Gabrielse, D. S. Hall, T. Roach, P. Yesley, A. Khabbaz, J. Estrada, C. Heimann, and H. Kalinowsky, Phys. Lett. B **455**, 311 (1999).
- [8] G. Gabrielse, J. Estrada, J. N. Tan, P. Yesley, N. S. Bowden, P. Oxley, T. Roach, C. H. Storry, M. Wessels, J. Tan, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, W. Breunlich, M. Cargnelli, H. Fuhrmann, R. King, R. Ursin, J. Zmeskal, H. Kalinowsky, C. Wesdorp, J. Walz, K. S. E. Eikema, and T. Haensch, Phys. Lett. B **507**, 1 (2001).
- [9] 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, Phys. Rev. Lett. **89**, 213401 (2002).

- [10] M. Amoretti, C. Amsler, G. Bonomi, A. Bouchta, P. Bowe, C. Carraro, C. L. Cesar, M. Charlton, M. J. T. 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. Holzschneider, L. V. Jorgensen, V. Lagomarsino, R. Landua, D. Lindelof, E. L. Rizzini, M. Macri, N. Madsen, G. Manuzio, M. Marchesotti, P. Montagna, H. Pruyss, C. Regenfus, P. Riedler, J. Rochet, A. Rotondi, G. Rouleau, G. Testera, A. Variola, T. L. Watson, and D. P. van der Werf, *Nature* **419**, 456 (2002).
- [11] C. H. Storry, A. Speck, D. Le Sage, N. Guise, G. Gabrielse, D. Grozonka, W. Oelert, G. Scheppers, T. Sefzick, J. Walz, H. Pittner, M. Herrmann, T. W. Hänsch, D. Comeau, and E. A. Hessels, *Phys. Rev. Lett.* **93**, 263401 (2004).
- [12] G. Gabrielse, P. Larochele, D. Le Sage, B. Levitt, W. S. Kolthammer, R. McConnell, P. Richerme, J. Wrubel, A. Speck, M. C. George, D. Grzonka, W. Oelert, T. Sefzick, Z. Zhang, A. Carew, D. Comeau, E. A. Hessels, C. H. Storry, M. Weel, J. Walz, and A. T. R. A. P. Collaboration, *Phys Rev Lett* **100**, 113001 (2008).
- [13] G. B. Andresen *et al.*, *Nature* **468**, 673 (2010).
- [14] G. B. Andresen *et al.*, *Nature Phys* **7**, 558 (2011).
- [15] Alan Kostelecky, *The Status of CPT*, 1998.
- [16] G. Lüders, *Ann. Phys.* **2**, 1 (1957).
- [17] M. Tanabashi *et al.*, *Phys. Rev. D* **98**, 030001 (2018).
- [18] C. Parthey, A. Matveev, J. Alnis, B. Bernhardt, A. Beyer, R. Holzwarth, A. Maistrou, R. Pohl, K. Predehl, T. Udem, T. Wilken, N. Kolachevsky, M. Abgrall, D. Rovera, C. Salomon, P. Laurent, and T. W. Hänsch, *Phys. Rev. Lett.* **107**, 203001 (2011).
- [19] T. D. Lee and C. N. Yang, *Phys. Rev.* **104**, 254 (1956).
- [20] R. L. Garwin, L. M. Lederman, and M. Weinrich, *Phys. Rev.* **105**, 1415 (1957).
- [21] J. I. Friedman and V. L. Telegdi, *Phys. Rev.* **106**, 1290 (1957).
- [22] L. D. Landau, *Nucl. Phys.* **3**, 127 (1957).
- [23] M. Gell-Mann and A. Pais, *Phys. Rev.* **97**, 1387 (1955).
- [24] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, *Phys. Rev. Lett.* **13**, 138 (1964).
- [25] A. D. Sakharov, *J. of Exp. and Th. Phys.* **5**, 24 (1967).

- [26] O. Bertolami, D. Colladay, V. A. Kostelecky, and R. Potting, *Phys. Lett. B* **395**, 178 (1997).
- [27] A. D. Dolgov and Y. B. Zeldovich, *Rev. Mod. Phys.* **53**, 1 (1981).
- [28] C. L. Cesar, D. G. Fried, T. C. Killian, A. D. Polcyn, J. C. Sandberg, I. A. Yu, T. J. Greytak, D. Kleppner, and J. M. Doyle, *Phys. Rev. Lett.* **77**, 255 (1996).
- [29] J. Wrubel, G. Gabrielse, W.S. Kolthammer, P. Larochele, R. McConnell, P. Richerme, D. Grzonka, W. Oelert, T. Sefzick, M. Zielinski, J.S. Borbely, M.C. George, E.A. Hessels, C.H. Storry, M. Weel, A. Mullers, J. Walz, and A. Speck, *Nuc. Inst. Meth. A* **640**, 232 (2011).
- [30] G. Gabrielse, X. Fei, L. A. Orozco, R. L. Tjoelker, J. Haas, H. Kalinowsky, T. A. Trainor, and W. Kells, *Phys. Rev. Lett.* **65**, 1317 (1990).
- [31] S. Agostinelli *et al.*, *Nuc. Inst. and Meth. A* **506**, 250 (2003).
- [32] D. Zambrano, personal communication.
- [33] G. Gabrielse *et al.*, *Opt. Lett.* **43**, 2905 (2018).
- [34] I. D. Setija, H. G. C. Werij, O. J. Luiten, M. W. Reynolds, T. W. Hijmans, and J. T. M. Walraven, *Phys. Rev. Lett.* **70**, 2257 (1993).
- [35] M. Ahmadi *et al.*, *Nature* **561**, 211 (2018).
- [36] C. M. Surko, M. Leventhal, and A. Passner, *Phys. Rev. Lett.* **62**, 901 (1989).
- [37] D. Comeau, A. Dror, D. W. Fitzakerley, M. C. George, E. A. Hessels, C. H. Storry, M. Weel, D. Grzonka, W. Oelert, G. Gabrielse, R. Kalra, W. S. Kolthammer, R. McConnell, P. Richerme, A. Mullers, and J. Walz, *New Journal of Physics* **14**, 045006 (2012).
- [38] William H. Wing, *Progr. Quant. Electr.* **8**, 181 (1984).
- [39] T. Bergman, G. Erez, and H. Metcalf, *Phys. Rev. A* **35**, 1535 (1987).
- [40] E. Majorana, *Nuovo Cimento* **9**, 43 (1932).
- [41] P. Richerme, Ph.D. thesis, Harvard University, 2012.
- [42] F. M. Esser, J. Schillings, H. Soltner, M. Schmitt, and G. Hansen, *IEEE Trans. on Appl. Superconductivity* **19**, 1278 (2009).
- [43] R. Kalra, Ph.D. thesis, Harvard University, 2014.

- [44] Harold J. Metcalf and Peter van der Straten, *Laser Cooling and Trapping* (Springer, New York, 1999).
- [45] G. Gabrielse *et al.*, Phys. Rev. Lett. **106**, 073002 (2011).
- [46] F. J. Ridgeon, M. J. Raine, D. P. Halliday, M. Lakrimi, Adrian Thomas, and D. P. Hampshire, IEEE Trans. on Appl. Superconductivity **27**, (2016).
- [47] Standard, American Society for Testing and Materials, West Conshohocken, PA (unpublished).
- [48] R.G. Sharma, in *Superconductivity* (Springer, Switzerland, 2015), pp. 49–69.
- [49] Jack Ekin, in *Experimental Techniques for Low-Temperature Measurements: Cryostat Design, Material Properties and Superconductor Critical-Current Testing* (Oxford University Press, New York, 2006).
- [50] L. S. Brown and G. Gabrielse, Rev. Mod. Phys. **58**, 233 (1986).
- [51] S.A. Prasad and T.M. O’Neil, Phys. Fluids **22**, 278 (1979).
- [52] L. D. Landau and E. M. Lifshitz, *Statistical Physics*, 2nd ed. (Pergamon Press, New York, 1969).
- [53] D. H. E. Dubin and T. M. O’Neil, Rev. Mod. Phys. **71**, 87 (1999).
- [54] H. Goldstein, C. Poole, and J. Safko, *Classical Mechanics*, 3rd ed. (Addison Wesley, Reading, MA, 2002).
- [55] Leaf Turner, Phys. Fluids **30**, 3196 (1987).
- [56] J. J. Bollinger and D. J. Wineland, Phys. Rev. Lett. **53**, 348 (1984).
- [57] J. J. Bollinger, D. J. Wineland, and D. H. E. Dubin, Phys. Plas. **1**, 1403 (1994).
- [58] R. L. Spencer, S. N. Rasband, and R. R. Vanfleet, Phys. Fluids B **5**, 4267 (1993).
- [59] A. Speck, G. Gabrielse, P. Laroche, D. Le Sage, B. Levitt, W.S. Kolthammer, R. McConnell, J. Wrubel, D. Grzonka, W. Oelert, T. Seifick, Z. Zhang, D. Comeau, M. C. George, E.A. Hessels, C. Storry, M. Weel, and J. Walz, Phys. Lett. B **650**, 119 (2007).
- [60] J. D. Jackson, *Classical Electrodynamics, 3rd Edition* (John Wiley and Sons, Inc., New York, 1975).
- [61] W Fundamenski and OE Garcia, Technical report, EFDA-JET (unpublished).

- [62] A. S. Richardson, Technical report, U.S. Naval Research Laboratory (unpublished).
- [63] M. Glinsky, T. O'Neil, M. N. Rosenbluth, K. Tsuruta, and S. Ichimaru, *Phys. Fluids B* **4**, 1156 (1992).
- [64] S. L. Rolston and G. Gabrielse, *Hyperfine Interact.* **44**, 233 (1988).
- [65] G. Gabrielse, W. S. Kolthammer, R. McConnell, P. Richerme, J. Wrubel, R. Kalra, E. Novitski, D. Grzonka, W. Oelert, T. Seifick, M. Zielinski, J. S. Borbely, D. Fitzakerley, M. C. George, E. A. Hessels, C. H. Storry, M. Weel, A. Mullers, J. Walz, and A. Speck, *Phys. Rev. Lett.* **105**, 213002 (2010).
- [66] 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, S. Jonsell, N. Madsen, S. Menary, P. Nolan, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Phys. Rev. Lett.* **106**, 145001 (2011).
- [67] S. P. Pucic, in *1993 IEEE Instrumentation and Measurement Technology Conference* (PUBLISHER, ADDRESS, 1993), pp. 114–117.
- [68] D. H. E. Dubin, *Phys. Rev. Lett.* **66**, 2076 (1991).
- [69] M. D. Tinkle, R. G. Greaves, C. M. Surko, R. L. Spencer, and G. W. Mason, *Phys. Rev. Lett.* **40**, 352 (1994).
- [70] Johnny K. Jennings, Ross L. Spencer, and K.C. Hansen, *Phys. Plas.* **2**, 2630 (1995).
- [71] E. Nottet, Ph.D. thesis, Northwestern University, 2020.
- [72] 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, S. Jonsell, N. Madsen, S. Menary, P. Nolan, A. Olin, A. Povilus, P. Pusa, F. Robicheaux, E. Sarid, D. M. Silveira, C. So, J. W. Storey, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, and Y. Yamazaki, *Phys. Rev. Lett.* **106**, 145001 (2011).
- [73] F. Anderegg, E. M. Hollmann, and C. F. Driscoll, *Phys. Rev. Lett.* **81**, 4875 (1998).

- [74] E. M. Hollmann, F. Anderegg, and C. F. Driscoll, *Physics of Plasmas* **7**, 2776 (2000).
- [75] J. R. Danielson and C. M. Surko, *Phys. Rev. Lett.* **94**, 035001 (2005).
- [76] J. R. Danielson and C. M. Surko, *Phys. Plas.* **13**, 055706 (2006).
- [77] R. D. Dixon and L. A. Lott, *Journal of Applied Physics* **40**, 4938 (1969).
- [78] D. Möhl, *Hyperfine Int.* **109**, 33 (1997).
- [79] J. F. Ziegler, M. D. Ziegler, and J. P. Biersack, *Nuc. Inst. Meth. B* **268**, 1818 (2010).
- [80] D. R. Bates and A. Dalgarno, in *Atomic and Molecular Processes*, edited by D. R. Bates (Academic Press, New York, 1962), pp. 245–271.
- [81] E. A. Hessels, D. M. Homan, and M. J. Cavagnero, *Phys. Rev. A* **57**, 1668 (1998).
- [82] G. Gabrielse, S. L. Rolston, L. Haarsma, and W. Kells, *Phys. Lett. A* **129**, 38 (1988).
- [83] M. Amoretti, C. Amsler, G. Bonomi, P. D. Bowe, C. Canali, C. Carraro, C. L. Cesar, M. Charlton, A. M. Ejsing, A. Fontana, M. C. Fujiwara, R. Funakoshi, P. Genova, J. S. Hangst, R. S. Hayano, L. V. Jorgensen, A. Kellerbauer, V. Lagomarsino, E. Lodi Rizzini, M. Macri, N. Madsen, G. Manuzio, D. Mitchard, P. Montagna, L. G. C. Posada, H. Pruys, C. Regenfus, A. Rotondi, H. H. Telle, G. Testera, D. P. Van der Werf, A. Variola, L. Venturelli, Y. Yamazaki, and N. Zurlo ATHENA Collaboration, *Physical Review Letters* **97**, 213401 (2006).
- [84] R. McConnell, Ph.D. thesis, Harvard University, 2011.
- [85] C. Amole, M. D. Ashkezari, M. Baquero-Ruiz, W. Bertsche, E. Butler, A. Capra, C. L. Cesar, 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, C. A. Isaac, S. Jonsell, L. Kurchaninov, A. Little, N. Madsen, J. T. K. McKenna, S. Menary, S. C. Napoli, K. Olchanski, A. Olin, P. Pusa, C. O. Rasmussen, F. Robicheaux, E. Sarid, C. R. Shields, D. M. Silveira, C. So, S. Stracka, R. I. Thompson, D. P. van der Werf, J. S. Wurtele, A. Zhmoginov, and L. Friedland, *Phys. Plas.* **20**, 043510 (2013).
- [86] P. Richerme, G. Gabrielse, S. Ettenuer, R. Kalra, E. Tardiff, D.W. Fitzakerley, M.C. George, E.A. Hessels, C.H. Storry, M. Weel, A. Mullers, and J. Walz, *Phys. Rev. A* **87**, 023422 (2013).
- [87] D. L. Moskovkin *et al.*, *Phys. Rev. A* **70**, 032105 (2004).

- [88] D. L. Moskovkin and V. M. Shabaev, Phys. Rev. A **73**, 052506 (2006).
- [89] C. Parthey *et al.*, Phys. Rev. Lett. **107**, 203001 (2011).
- [90] Bo nian Dai and P. Lambropoulos, Phys. Rev. A **34**, 3954 (1986).
- [91] M. Haas *et al.*, Phys. Rev. A **73**, 052501 (2006).
- [92] S. Klarsfeld, Phys. Lett. A **30**, 382 (1969).
- [93] G. Breit and E. Teller, Astrophys. J. **91**, 215 (1940).
- [94] Keith McDowell, J. Chem. Phys. **65**, 2518 (1976).
- [95] D. J. Wineland and W. M. Itano, Phys. Rev. A **20**, 1521 (1979).
- [96] C. Ø. Rasmussen, N. Madsen, and F. Robicheaux, J. Phys. B **50**, 184002 (2017).
- [97] J. Alnis, A. Matveev, N. Kolachevsky, Th. Udem, and T. W. Hänsch, Phys. Rev. A **77**, 053809 (2008).