Progress towards a sub-ppb measurement of the antiproton magnetic moment

A thesis presented by

Kathryn Marable

 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 July 2019 O2019- Kathryn Marable

All rights reserved.

Gerald Gabrielse

Kathryn Marable

Progress towards a sub-ppb measurement of the antiproton magnetic moment

Abstract

The antiproton magnetic moment was measured for the first time using a single antiproton. This allowed a 680-fold improvement in precision, and a resulting value in Bohr magnetons of $\mu_p/\mu_N = 2.792845(12)$. This first direct comparison of the antiproton and proton magnetic moments, resulted in a measured ratio of -1.000000(5), a stringent test of CPT. An apparatus that separates the particle excitation and detection was designed, constructed and partly demonstrated. The proposed measurement sequence and result is presented.

Contents

	Title Page
	Abstract
	Table of Contents
	List of Figures
	List of Tables
	Publications
	Acknowledgments
1	Introduction
	1.1 The magnetic moment
	1.2 History
	1.3 CPT symmetries
	1.4 Overview
2	The Penning trap
	2.1 Axial motion and frequency
	2.2 Radial motions
	2.3 Typical energies and amplitudes
	2.4 Detecting particle motions
	2.5 Driven signals
	2.6 Feedback and the self-excited oscillator
	2.7 Magnetron cooling
	2.8 Magnetic bottle
	2.9 Driving spin flips
	2.10 Lineshape
	2.11 Apparatus
	2.11.1 Cryogenic system and electronics
	$2.11.2 \text{ Traps} \qquad \qquad$
3	A part-per-million antiproton measurement 3
-	3.1 Loading antiprotons

		3.1.1 Slowing antiprotons	36
		3.1.2 Detection	38
		3.1.3 Electron cooling	39
	3.2	Detecting spin flips	41
		3.2.1 Instability due to cyclotron transitions	43
		3.2.2 Cyclotron noise drive	45
		3.2.3 Minimizing the scatter	48
	3.3	Measurement sequence	49
	3.4	Antiproton measurement	51
4	Sin	gle spin flips	53
	4.1	Introduction	53
	4.2	The single spin flip dataset	54
		4.2.1 Efficiency and fidelity	56
		4.2.2 Verification with correlations	59
		4.2.3 Application to <i>q</i> -factor measurement	61
	4.3	Improving stability	62
		4.3.1 Effect of fidelity on measurement time	63
		4.3.2 Estimating number of trials as a function of E, F and σ_0	65
		4.3.3 Adiabatic Fast Passage	68
	4.4	Conclusion	71
5	Sep	parated Oscillatory Fields Measurement of the g -factor	72
5	Sep 5.1	DescriptionDescriptionHistorical overview	72 73
5	Sep 5.1 5.2	Parated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances	72 73 75
5	Sep 5.1 5.2	Parated Oscillatory Fields Measurement of the g-factor Historical overview	72 73 75 75
5	Sep 5.1 5.2	Parated Oscillatory Fields Measurement of the g-factor Historical overview	72 73 75 75 75 76
5	Sep 5.1 5.2 5.3	Parated Oscillatory Fields Measurement of the g-factor Historical overview	72 73 75 75 76 78
5	Sep 5.1 5.2 5.3 5.4	Parated Oscillatory Fields Measurement of the g-factor Historical overview	72 73 75 75 76 78 81
5	Sep 5.1 5.2 5.3 5.4	Parated Oscillatory Fields Measurement of the g-factor Historical overview	72 73 75 75 75 76 78 81 81
5	Sep 5.1 5.2 5.3 5.4	Parated Oscillatory Fields Measurement of the g-factor Historical overview	72 73 75 75 75 76 78 81 81 83
5	Sep 5.1 5.2 5.3 5.4	Parated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3	72 73 75 75 76 78 81 81 83 84
5	Sep 5.1 5.2 5.3 5.4 5.5	arated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3 Other sources of error Procedure for compensating field drift	72 73 75 75 76 78 81 81 83 84 86
5	Sep 5.1 5.2 5.3 5.4 5.5 5.6	arated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3 Other sources of error Procedure for compensating field drift Demands on experimental performance	72 73 75 75 76 78 81 81 83 84 86 87
5	Sep 5.1 5.2 5.3 5.4 5.5 5.6 5.7	arated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3 Other sources of error Procedure for compensating field drift Demands on experimental performance	72 73 75 75 76 78 81 81 83 84 86 87 88
5	Sep 5.1 5.2 5.3 5.4 5.5 5.6 5.7 Tra	arated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field Experimental errors affecting the cyclotron measurement 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3 Other sources of error Procedure for compensating field drift Predicted measurement time Predicted measurement time Predicted measurement time	72 73 75 75 76 78 81 81 83 84 86 87 88 91
5	Sep 5.1 5.2 5.3 5.4 5.5 5.6 5.7 Tra 6.1	arated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field Experimental errors affecting the cyclotron measurement 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3 Other sources of error Procedure for compensating field drift Demands on experimental performance Predicted measurement time Overview	72 73 75 75 76 78 81 81 83 84 86 87 88 91 91
5	Sep 5.1 5.2 5.3 5.4 5.5 5.6 5.7 Tra 6.1 6.2	arated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field Experimental errors affecting the cyclotron measurement 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3 Other sources of error Procedure for compensating field drift Predicted measurement time Predicted measurement time Predicted measurement time	72 73 75 75 76 78 81 81 83 84 86 87 88 91 91 94
5	Sep 5.1 5.2 5.3 5.4 5.5 5.6 5.7 Tra 6.1 6.2	arated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3 Other sources of error Procedure for compensating field drift Demands on experimental performance Predicted measurement time Predicted measurement time 6.2.1 Overview	72 73 75 75 76 78 81 81 83 84 86 87 88 91 94 94
5	Sep 5.1 5.2 5.3 5.4 5.5 5.6 5.7 Tra 6.1 6.2	arated Oscillatory Fields Measurement of the g-factor Historical overview Derivation of spin and cyclotron resonances 5.2.1 Spin resonance 5.2.2 Cyclotron resonance g-factor lineshape in a drifting magnetic field Experimental errors affecting the cyclotron measurement 5.4.1 Error from the unknown initial state 5.4.2 Error due to cyclotron excitation drifts in the analysis trap 5.4.3 Other sources of error Procedure for compensating field drift Predicted measurement time Predicted measurement time Predicted measurement time 6.2.1 Overview of damping time reductions 6.2.2 Calculating the optimized trap geometry	72 73 75 75 75 76 78 81 81 83 84 86 87 88 91 94 94 94

	6.3	Analysis trap	102
	6.4	Interactions between the cooling and analysis traps	106
	6.5	Conclusion	110
7	\mathbf{Exp}	erimental Apparatus	112
	7.1	Electrode construction	112
		7.1.1 Polishing, brazing, and splitting	113
		7.1.2 Plating	116
		7.1.3 Trap assembly \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots	119
	7.2	Wiring	120
	7.3	Thermal and magnetic shielding	131
	7.4	Thermal shield and magnetic shielding	131
		7.4.1 New bottom radiation shield	135
	7.5	Magnetic shielding and stability	136
	7.6	Alignment	138
8	Det	ecting and driving particles in a new apparatus	143
	8.1	Amplifier theory	144
	8.2	Axial amplifiers	146
		8.2.1 Design and construction of the new axial amplifiers	148
		8.2.2 Results and future steps	152
		8.2.3 Heat sinking	154
	8.3	Cyclotron amplifiers	157
		8.3.1 The cyclotron amp in the cooling trap	157
		8.3.2 Loading cyclotron amplifier	161
		8.3.3 Loading cyclotron amplifier design	163
		8.3.4 Loading cyclotron amplifier performance	165
		8.3.5 Future steps	166
	8.4	Sweeping the cyclotron drive in a magnetic gradient	169
	8.5	Frequency uncertainty due to magnetic gradients	171
		8.5.1 Gradient effects on decay shape	173
		8.5.2 Linewidth broadening and cyclotron sweeps	175
		8.5.3 Next steps for cyclotron detection	180
	8.6	Degrader and loading	181
	8.7	Conclusion	186
9	Con	clusion	187
Bi	Bibliography 19		

List of Figures

1.1	History of measurements of the proton and antiproton magnetic mo- ment.	4
1.2	Fractional precision of four kinds of tests of CPT symmetry: parti-	
	cle/antiparticle comparisons of g-factors, masses m , charges q , and	
	lifetimes τ	7
2.1	Motions of a particle in a penning trap	10
2.2	Section view of the precision trap of the antiproton measurement with	
	labeled dimensions	12
2.3	Schematic of feedback drive	21
2.4	Radiofrequency schematic for applying either a two-drive scheme or	
	feedback for the self-excited oscillator	22
2.5	Energy levels with an axial-magnetron coupling drive	23
2.6	(a) Schematic of the analysis trap, with the iron ring shown in grey	
2.7	(b) Magnetic field along the central axis of the trap	25
	spin flips	27
2.8	Axial frequency shift as a function of the drive applied at the hat.	29
2.9	Spin and cyclotron lineshapes	31
2.10	Solidworks model of the trap stack as used in the antiproton magnetic	
	moment measurement	34
3.1	Diagram of PPAC tube	36
3.2	Antiprotons trapped in the high voltage well as a function of timing .	37
3.3	The 'long well' potential originally used to trap antiprotons	39
3.4	The 'short well' potential that successful caught antiprotons and sym-	
	pathetically cooled them with electrons.	40
3.5	The stability of axial frequency measurements in the analysis trap	42
3.6	Histograms of axial frequency shifts with and without drive	44
3.7	Allan deviation as a function of cyclotron quantum number	45
3.8	Measurement steps versus time for a single data point on the spin line	49

3.9	Data points and fit from the measurement	50
$4.1 \\ 4.2 \\ 4.3$	Stability of axial frequency measurements from single spin dataset Histograms of frequency differences for single spin data	55 56
4.4 4.5	tributions with two different σ_0	57 61
4.6	scatter	65
$4.7 \\ 4.8$	scatter	67 · 67
4.9	and 100 percent spin flip probability	68 69
5.1	Calculated spin transition probability fringe from separated oscillatory fields method	76
$5.2 \\ 5.3$	G-factor fringe with different evolution times	80 85
6.1 6.2 6.3	Picture of the new trap stack	92 96 99
6.5	function of separation between electrode center and analysis trap Calculated κ_x as a function of distance between the analysis and cooling	105
$\begin{array}{c} 6.6 \\ 6.7 \end{array}$	traps Cyclotron frequency as a function of distance from analysis trap z γ as a function of z , corrected for shifting frequency	108 109 110
$7.1 \\ 7.2 \\ 7.3$	Polishing electrodes on a lathe	114 115
$7.4 \\ 7.5$	platers	118 119
$7.6 \\ 7.7$	baffleCAD model of the tripod regionArrangement of cables through the support structure of the experiment	121 123 124
7.8 7.9 7.10	Photo of components and wiring inside the trap stack DC and RF wiring for precision trap	125 128
1.10	$D \cup$ and rep withing for analysis and cooling trap \ldots \ldots \ldots	129

$7.11 \\ 7.12$	DC and RF wiring for loading trap	$130 \\ 132$
7.13	Magnetic field variation due to the AD cycle in the zone and measured from the particle's frequency	132
7.14	Solidworks model of the new thermal shield extension that fits over the	100
7.15	Photograph of the new shield mounted on the PPAC before re-installation	104
7.16	Solidworks model of the magnetometer mounted on the endplate of the trap stack	136 139
8.1	Diagram showing two models of the how the cyclotron motion of the particle across a split comp induces currents in a tuned circuit: (a)	144
8.2	Photo of the new analysis axial amp with the top of the can open	$144 \\ 147$
8.3	SolidWorks model of the new axial amplifier design, showing the coil	
0.4	support mounted on Teflon bottom support. A.Top view. B. Side view	150
8.4	Layout of the new axial amplifier board, top view. Copper is shown in	159
85	Noise spectrum from the new analysis amplifier with a test capacitor	152 152
8.6	Noise signal with a dip due to roughly 50 particles.	153
8.7	Photograph of the loading cyclotron amp mounted inside the trap can	
	next to the loading trap	161
8.8	Wiring diagram of the loading trap resonator and FET	164
8.9	Noise signal from loading cyclotron amp	165
8.10	Proposed revision of the loading trap for better RF performance	168
8.11	Cyclotron peaks from antiprotons.	169
8.12	Estimated cyclotron frequency as a function of position	173
8.13	Projected change of cyclotron frequency with radius due to cooling	
	trap gradients	175
8.14	Projected change of cyclotron frequency with radius due to loading	
	trap gradients.	176
8.15	Steady state response of a particle to a cyclotron drive due to magnetic	
	gradients or relativistic shift	179
8.16	SolidWorks model of the degrader of the new loading trap and degrader	100
0.15	assembly mounted inside the trap can	183
8.17	The number of antiprotons loaded as a function of SF6 fraction in the	104
0 10	energy tuning cell	184
8.18	Antiprotons caught as a function of time to close the trap	199

List of Tables

2.1 Frequencies, amplitudes, and quantum numbers for all three motion These numbers are taken from the analysis trap in 2012, but oth			
	traps vary by at most a factor of two	16	
3.1	Uncertainties for the antiproton magnetic moment measurement	52	
6.1	Heights of the new cooling and analysis trap electrodes $\ldots \ldots \ldots$	93	
8.1	Properties of the three axial amps	149	
8.2	Trap capacitance sources in current and proposed revised trap	168	
8.3	Gradients in the loading and cooling traps	174	
8.4	Linewidths in the loading and cooling traps	177	

Publications

- Resolving an Individual One-Proton Spin Flip to Determine a Proton Spin State
 J. DiSciacca, M. Marshall, K. Marable, and G. Gabrielse, *Phys. Rev. Lett.* **110**, 140406 (2013).
- One-Particle Measurement of the Antiproton Magnetic Moment
 J. DiSciacca, M. Marshall, K. Marable, G. Gabrielse, S. Ettenauer, E. Tardiff, R. Kalra, D. W. Fitzakerley, M. C. George, E. A. Hessels, C. H. Story, M. Weel, D. Grzonka, W. Oelert, and T. Sefzick, *Phys. Rev. Lett.* **110**, 130801 (2013).

Acknowledgments

First, I would like to thank Jerry Gabrielse for his support, advice and encouragement. I am grateful for the opportunity to work on this project thanks to his vision in creating it and the support of the National Science Foundation. I am also indebted to the work of Nicholas Guise, who started this project. Jack DiSciacca provided invaluable mentorship and guidance during my first two years as I was starting out on this project. He was an excellent teacher, scientist and example for a younger student.

I joined the experiment only six months after Mason Marshall, and working directly with him for the first half of my thesis was a pleasure. Even after moving to CERN, we have both benefited from sharing advice, support and a second pair of eyes as necessary, and I am grateful for his help.

I thank Ed Myers for his valuable insights into measurement methods and the year he spent on sabbatical helping us set up the experiment at CERN. His wisdom and advice were greatly appreciated.

I also thank Gunn Khatri for his good humor and contributions to the experiment, particularly the antiproton annhibition detection systems. I thank Yao Chen for his help taking data and constructing the new upgraded apparatus, as well as many insightful conversations.

I thank Geev Nahal for his assistance with construction of the new amplifiers for the experiment, and Andra Ionescu for her efforts designing the new loading trap. I believe the experiment will be in good hands as they take over.

I additionally benefited from the advice, support and wisdom from many other postocs and students in the group. I would like to thank Elise Novitski, Shannon Fogwell, Joshua Dorr, Cris Panda, Cole Meisenhelder, Melissa Wessels, Ronald Alexander and Maryrose Barrios for helpful discussions and advice.

I thank members of the ATRAP collaboration for their advice, support, and willingness to lend a hand to get things running at CERN, including Stefan Ettanauer, Dan Fitzakerley, Dieter Grzonka, Walter Oelert, Eric Tardiff, Rita Kalra, Nate Jones, Chris Hamley, and Tharon Morrison.

I thank my committee members, Professors Masahiro Morii and Isaac Silvera, for their support during my graduate studies.

I have also benefited from the hard work and expertise of many talented specialists at the Harvard SEAS Scientific Instrument shop and the Harvard Instrument design lab, with special thanks to Stan Cotreau, Mike McKenna and Jim McArthur.

I also had the opportunity to work with several bright and hardworking undergraduates. In particular, I would like to thank Makinde Ogunnaike, Jonah Philion, and Meutia Wulansatiti for their work on the experiment.

Finally, I would like to thank my parents Jeanne Wilson and James Marable, my brother Eli Marable, and my fiancee Maarten Van Dijk for their support and encouragement without which I would never have accomplished this.

Chapter 1

Introduction

This thesis reports on comparisons of the proton and antiproton magnetic moments. Such measurements are a precise test of CPT symmetry in the baryon sector. We report on a measurement of the antiproton magnetic moment at the ppm level [1], a 680-fold improvement on the best previous comparison. It reviews initial work towards a sub-ppb measurement in the same apparatus, which has successfully detected spin flips as first reported in [2]. We then present the design and construction and first results of a major upgrade program to allow higher precision. We conclude with a report on the upgrade program, which centers on a completely rebuilt trap apparatus with new and improved analysis and precision traps as well as two new specialized traps for loading antiprotons and cooling their motion.

1.1 The magnetic moment

The spin magnetic moment μ of a particle is a fundamental property defined by

$$\vec{\mu} = -g \frac{e\hbar}{2m} \vec{S} \tag{1.1}$$

where \vec{S} is its spin, m its mass, -e its charge, and g is a dimensionless factor derived from the underlying physics. The spin magnetic moment of a Dirac particle with the charge, mass and spin of the proton has a magnitude equal to the nuclear magneton $\mu_N = e\hbar/(2m_p)$, multiplied by the scaled spin $\vec{\sigma} = \vec{S}/(\hbar/2)$. As the proton is a composite rather than a point particle, its magnetic moment is larger than a Dirac particle by a factor of g/2. The magnitudes of the proton and antiproton magnetic moments, given by $\vec{\mu}_p = \mu_p \vec{\sigma}$ and $\vec{\mu}_{\bar{p}} = \mu_{\bar{p}} \vec{\sigma}$ respectively, are then equal to

$$\mu_p = \frac{g_p}{2} \mu_N \tag{1.2}$$

$$\mu_{\overline{p}} = -\frac{g_{\overline{p}}}{2}\mu_N,\tag{1.3}$$

A measurement of the g-factor is equivalent to measuring the magnitude of the magnetic moment in units of nuclear magnetons.

1.2 History

The magnetic moment of the proton has been subject of increasingly precise measurements since the 1930s. Initial measurements were based on the deflection of molecular beams in a magnetic field [3, 4], which already indicated a g-factor significantly larger than the value of two predicted from the Dirac wave equation for a point particle [5]. From their invention in the late 1930s on, nuclear magnetic resonance methods and cyclotron orbit frequencies were used for increasingly precise measurements of the proton spin magnetic moment [6]. Like our experiment, these experiments measured the *g*-factor by simultaneously determining the cyclotron and spin frequencies in a strong magnetic field, but they needed to use separate probes for the two frequencies. The cyclotron frequency was measured in a cyclotron accelerator, and the spin frequency was measured using the NMR resonance of a separate water or oil-filled probe inserted into the cyclotron field [7, 8]. An early particle trap called an omegatron is more similar to our experiment. However, lacking our non-destructive RF detection methods, the resonance was detected from exciting the particles until they struck an ion collector at the edge of the trap, and a water probe was used to find the proton spin resonance [9]. All of these measurements were limited by theoretical corrections for protons in water or oil to an accuracy of a few ppm.

The development of the hydrogen maser allowed a further improvement in precision by measuring both frequencies in the simpler system of a hydrogen atom. A 1972 measurement combining the hydrogen maser with several other measurements and theoretical corrections increased the precision by two orders of magnitude to the 10 ppb level [10]. This result is limited by the accuracy of the ratio of the proton and electron magnetic moments in hydrogen $\frac{\mu_p(H)}{\mu_e(H)}$, which has only been measured to 10 ppb [11]. A Penning trap measurement of a single proton would not be limited by this uncertainty.

Previous measurements of the antiproton magnetic moment have not been able to take advantage of the methods used for the proton, as they would require macroscopic



Figure 1.1: History of measurements of the proton and antiproton magnetic moment. Data gathered from [6, 7, 8, 9, 12, 13, 1, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24]

quantities of anti-atoms. Instead, the most accurate experiments use antiprotonic atoms, which are formed by colliding an antiproton beam with a target. When these atoms decay, they emit x-rays that show a fine structure splitting which depends on $\mu_{\overline{p}}$ [15]. The latest measurement using antiprotonic helium measured $\mu_{\overline{p}} = 2.7862(83)$ [14]. These experiments faced many challenges due to the short lifetimes of the atoms, theoretical uncertainty in their structure and the accelerator environment. These limit prospects for higher precision measurement using these techniques.

A Penning trap measurement has the advantage that the same apparatus can be used for both protons and antiprotons. Our apparatus was used in a demonstration measurement of the proton at ppm precision [25]. A year later, I and others made the first single-antiproton measurement in the same apparatus. At ppm precision, this represented a 680-fold improvement in precision over the previous measurement [1]. A parallel effort at the university of Mainz has since extended that to 0.3 ppb measurements of the proton [12] and 1.5 ppb of the antiproton [13]. The historical progression of magnetic moments towards better precision is shown in Figure 1.1. Meanwhile, we have been improving our apparatus as described in Chapter 6 with the goal of a 0.1 ppb comparison of the antiproton and proton.

1.3 CPT symmetries

There is a long history of symmetry tests leading to the discovery of new physics. Originally, it was believed that all fundamental interactions would be the same under a parity transformation that swapped the positions of every particle from \vec{x} to $-\vec{x}$. However, in 1957, Wu experimentally discovered parity violations in certain weak interactions in the beta decay of Cobalt-60 [26]. At that point, it was suggested that symmetry under a combination of charge conjugation and parity transformation would form a good symmetry. Yet, less than ten years later, CP symmetry was also found to be violated in the decay of neutral kaons by Cronin and Fitch [27]. The combined symmetry of charge, parity and time reversal, however, has been preserved in all particle experiments thus far. It is believed to be an exact symmetry because any locally Lorentz-invariant quantum field theory, like the standard model, unavoidably has CPT symmetry [28].

Despite the great success of the Standard Model in the last decades, major features of the universe are not explained. An obvious example is the dominance of matter over antimatter in the universe. Limits from astronomical searches for gamma rays from matter/antimatter annihilation, as well as isotope abundances, give confidence that this is not just a feature of our local galaxy but the entire visible universe [29]. CP violation of the standard model develops from a complex phase in the mixing matrix of weak interactions. Based on the best current measurement of this phase, CP violation in the standard model is about 10 orders of magnitude smaller than needed to explain the matter dominance observed in the universe [29].

CPT symmetry predicts that the mass, lifetime, charge magnitude, and magnetic moment magnitude of every antiparticle should be identical to that of the corresponding particles [30]. Thus, precise measurements of these properties are very strict tests of this symmetry. The most stringent of these comparisons is the K_0/\overline{K}_0 mass, because the kaon mass eigenstate mixes with the anti-kaon to form a natural interferometer. The mass of the K_0 and \overline{K}_0 were compared to to better than a part in 10^{-19} in 2011 [31].

However, without strong theoretical guidance about what sorts of forces or interactions would be expected to violate CPT, it is necessary to probe a number of different systems. The current state of these comparisons across a variety of sectors and parameters is shown in Figure 1.2. The electron and positron magnetic moment comparison at Washington is the strongest test in a lepton system with ppt precision [32]. A new experiment in development by our group (now moved to Northwestern University) is seeking to improve this by a factor of 150 [33, 34]. Precise comparisons of baryons are particularly interesting since a matter-antimatter asymmetry is observed in the baryon sector. Asymmetries in other sectors would require an additional coupling mechanism outside the standard model to explain how that asymmetry is transferred into baryons.



Figure 1.2: Fractional precision of four kinds of tests of CPT symmetry: particle/antiparticle comparisons of g-factors, masses m, charges q, and lifetimes τ . Data compiled from [11, 12, 13]

1.4 Overview

This thesis begins with work done early in my Ph.D. on measuring the antiproton magnetic moment to ppm precision and detection of individual spin flips. The rest of the thesis presents a redesign and rebuilding of the apparatus with the goal of a new measurement at sub-ppb precision.

- Chapter 2 reviews the theory of Penning traps and the basic techniques we use for detecting and driving particles in our apparatus. The apparatus was already constructed and ready for use in a part-per-million measurement when I joined and more information is available in the thesis of Nick Guise [35] and Jack DiSciacca [36].
- Chapter 3 presents a ppm measurement of the magnetic moment of the antipro-

ton, which I contributed to. This is also presented in Jack DiSciacca's thesis [36] as he was the senior graduate student. The chapter includes issues encountered that will be addressed by the new design presented in the later chapters, particularly loading of antiprotons and stability of the analysis axial frequency.

- Chapter 4 describes detection of single spin flips in the old apparatus, an important step towards reaching higher precision. It also analyzes two strategies for future improvements to detection efficiency.
- Chapter 5 describes a new measurement strategy based on the method of separated oscillatory fields.
- Chapter 6 describes a new trap electrode design. The new electrodes include a dedicated loading trap and cooling trap, as well an improved analysis trap design. The new cooling and analysis trap are predicted to increase the rate of cyclotron cooling by more than an order of magnitude.
- Chapter 6 and 7 describe the construction and early commissioning of the new trap stack and apparatus. Chapter 6 discusses the construction of the new electrodes, new wiring, and a new alignment system. Chapter 7 focuses on the construction of the new cryogenic amplifiers and the new loading trap, and investigates the challenges that both systems faced due to magnetic gradients.

Chapter 2

The Penning trap

An ideal Penning trap is a combination of a spatially homogeneous magnetic field and a quadrupole electric potential that traps a charged particle, producing exactly defined motions. The magnetic field causes the particle to circle around its field lines in a cyclotron motion, trapping it radially. The electrostatic field confines the particle along the axis of the magnetic field resulting in a harmonic axial oscillation. The electrostatic potential also has a component that attracts the particle radially outward towards the walls of the trap. The interaction of this force with the magnetic field causes a slow drift in the center of the cyclotron motion called the magnetron motion. The trapping fields and the resulting motions are depicted in Figure 2.1. The theory of a cylindrical Penning trap can be solved analytically, allowing for a theoretically simple interpretation of our measurement [37].

We use a series of coaxial Penning traps. The new apparatus described in Chapter 6, our apparatus included 4 traps, for capturing, cooling, exciting, and detecting particle motions. Most of the basic principles described in this section apply for all



Figure 2.1: Motions of a particle in a penning trap [36]

traps.

2.1 Axial motion and frequency

The frequency of the oscillation along the axis of the magnetic field depends on the electrostatic potential applied. An ideal quadrupole potential is

$$V_{ideal}(z,\rho) = V_q \frac{z^2 - \rho^2/2}{2d^2}$$
(2.1)

where ρ is the trap radius, V_q gives the strength of the trapping potential and d is a trap size parameter that is defined in terms of the trap radius ρ_0 and distance from trap center to the endcap z_0 by $d^2 = \frac{1}{2} \left(z_0^2 + \rho_0^2/2 \right)$. For a particle moving along the trap axis with mass m and charge q, the angular axial frequency would be given by

$$\omega_z^2 = qV_q/\left(md^2\right) \tag{2.2}$$

The actual potential inside the trap electrodes $V(z, \rho)$ is determined by the shape of the trap electrodes, and by the applied potentials and imperfections. While the original Penning traps used hyperboloid electrodes to produce a quadrupole field [38], in our trap we use cylindrical electrodes with open bore endcaps [39] that allow particle transfers between traps and loading of antiprotons from an incoming beam. Our traps consist of five segments like the precision trap shown in Figure 2.2¹: a ring, two endcaps, and two compensation electrodes. We approximate a quadrupole potential near trap center with only two precision voltage sources by applying a potential V_0 to the ring, V_{comps} to the compensation electrodes, and grounding the endcaps. The calculation of the potential at trap center is more straightforward if we shift the ground potential by $V_0/2$ and calculate using the equivalent potentials of $V_0/2$ on the ring, $-V_0/2$ on the endcaps and $V_c = V_{comps} - V_0/2$ on the compensation electrodes. The potential difference between the ring and the endcaps keeps the particles trapped along the axis of the electrode stack, which we align with the magnetic field axis. The compensation electrodes allow better approximation of a harmonic potential.

The potential near trap center can be written as a sum of Legendre polynomials due to its cylindrical and reflection symmetries:

$$V\left(\vec{r}\right) = \frac{1}{2} \sum_{\substack{k=0\\even}} \left(V_0 C_k^{(0)} + V_c D_k \right) \left(\frac{r}{d}\right)^k P_k\left(\cos\theta\right)$$
(2.3)

The coefficients $C_k^{(0)}$ and D_k can be calculated by solving Laplace's equation subject to the boundary conditions. The small $C_k^{(0)}$ are typically strongly affected by trap imperfections. We can choose the ratio of the trap voltages V_0/V_c so that $C_4^0 = -\frac{V_c}{V_0}D_4$, making the fourth order terms vanish. Fixing this ratio, we can simplify the solution

¹In the new apparatus, a simplified three-electrode design that lacks compensation electrodes suffices for the cooling trap. See Section 6.2.3



Figure 2.2: Section view of the precision trap of the antiproton measurement with labeled dimensions [35]

by defining $C_k = C_k^{(0)} + \frac{V_c}{V_0} D_k$ so that:

$$V\left(\vec{r}\right) = \frac{V_0}{2} \sum_{\substack{k=0\\even}} C_k \left(\frac{r}{d}\right)^k P_k\left(\cos\theta\right)$$
(2.4)

By careful choice of the trap dimensions z_c and z_0 , we can arrange that $C_6 = 0$ as well. Higher order terms can be neglected for a well-cooled particle because its entire motion is close to trap center where $r/d \ll 1$. Then the only non-vanishing term in the series is

$$V(z) = \frac{C_2 V_0 z^2}{2d^2}$$
(2.5)

which is the quadrupole potential of Equation 2.1 when $\rho = 0$ and $V_q = C_2 V_0$, yielding

$$\nu_z = \frac{1}{2\pi} \sqrt{\frac{qV_0 C_2}{md^2}}$$
(2.6)

for the axial frequency in our trap.

2.2 Radial motions

The particle also moves in the plane perpendicular to the magnetic field, which we separate into the cyclotron and magnetron motions. In the absence of the electric field, the particle would oscillate around the magnetic field lines with a frequency

$$\omega_c = 2\pi\nu_c = \frac{eB}{m} \tag{2.7}$$

This is referred to as the "free space" cyclotron frequency. In our trap, this motion is complicated by the radial component of the electric field. The full motion of the particle can be found by solving the radial equation of motion:

$$\ddot{\vec{\rho}} - \omega_c \hat{z} \times \dot{\vec{\rho}} - \frac{1}{2} \omega_z^2 \vec{\rho} = 0$$
(2.8)

The solution is a superposition of two circular motions at ω_{\pm} [37], with frequencies related by

$$\omega_+ = \omega_c - \omega_- \tag{2.9}$$

$$\omega_{-} = \omega_z^2 / 2\omega_{+} \tag{2.10}$$

We always choose the magnetic and electric fields such that $\omega_c >> \omega_z$ and thus $\omega_c >> \omega_-$. That means that ω_+ only differs from the free-space cyclotron frequency by a small correction, so it is called the trap-modified cyclotron frequency ω'_c . The second orbit at oscillation frequency $\omega_- = \omega_m$ is the magnetron motion. In all of our traps, $\nu_m = \omega_m / (2\pi)$ is typically on the order of a few kHz, $\nu_z = \omega_z / (2\pi)$ is around 1 MHz, and $\nu'_c = \omega'_c / (2\pi)$ is around 85 MHz. There is a resulting frequency hierarchy

$$\omega_m << \omega_z << \omega'_c \tag{2.11}$$

The g-factor in a Penning trap

The particle also has an intrinsic spin that will align or anti-align with the magnetic field. The frequency of the spin transition is given for a particle of charge -e

$$\omega_s = 2\pi\nu_s = -2\vec{\mu}\cdot\vec{B} = g\frac{eB}{m} \tag{2.12}$$

. The ratio of the free space cyclotron frequency and the spin frequency is proportional to the g-factor

$$\frac{g}{2} = \frac{\omega_s}{\omega_c} \tag{2.13}$$

We rely on the Brown-Gabrielse Invariance theorem to find the free space cyclotron frequency from the motional frequencies of the particle [37]:

$$\omega_c^2 = \left(\omega_c'\right)^2 + \omega_z^2 + \omega_m^2 \tag{2.14}$$

This relationship is valid even in the presence of the leading quadratic imperfections of the harmonic potential or the alignment of the electric and magnetic fields. As a result of the hierarchy of our trap (Equation 2.11), to measure ω_c with ppb precision we only need to measure (ω'_c) to such a high level of precision. For a similar contribution to the error, the axial frequency only needs to be measured to 10 ppm, and the magnetron frequency only needs to be measured to about 25 percent.

2.3 Typical energies and amplitudes

The motional energies of the particle in our trap are large enough that they can be treated classically. A classical particle oscillating with axial amplitude A and cyclotron radius ρ_+ has axial energy E_z and (trap-modified) cyclotron energy E_c :

$$E_z = \frac{1}{2}m\omega_z^2 A^2 \tag{2.15}$$

$$E_{c} = \frac{1}{2}m\left(\omega_{c}'\right)^{2}\rho_{+}^{2}$$
(2.16)

The axial motion and cyclotron motion are typically coupled to amplifiers that act as thermal reservoirs with average temperatures around 8 K, described in Section 2.4. Given these temperatures, we can calculate typical amplitudes and quantum numbers from the equipartition theorem (shown for the analysis trap in Table 2.1). A cold particle's amplitude is measured in μ m as shown in the table, which validates the assumption of the previous section that cold particles stay close to trap center for their entire orbit. For the axial motion,

$$\langle z_{rms} \rangle = \sqrt{\frac{k_B T_z}{m\omega_z^2}} \tag{2.17}$$

Because the magnetron motion comes from a negative electrostatic potential, its energy is unstable and given by:

$$V_m(\rho) = -V_0 \frac{\rho^2}{4d^2}$$
(2.18)

$$E_m = -\frac{1}{4}m\omega_z^2 \rho_m^2 \tag{2.19}$$

Despite the generally classical behavior, a quantum formulation is necessary when these motions interact with the spin state. In that case, each motion can be described as a quantum harmonic oscillator. Its state can be identified with the quantum numbers for each motion: cyclotron n, axial k, and magnetron l. The energies are

Motion	Frequency	Amplitude (μ m)	Quantum number
Cyclotron	$79 \mathrm{~MHz}$	0.5	2000
Axial	920 kHz	40	90000
Magnetron	$5~\mathrm{kHz}$	5	90000

Table 2.1: Frequencies, amplitudes, and quantum numbers for all three motions. These numbers are taken from the analysis trap in 2012, but other traps vary by at most a factor of two.

then

$$E_c = \hbar \omega_c' \left(n + \frac{1}{2} \right) \tag{2.20}$$

$$E_z = \hbar\omega_z \left(k + \frac{1}{2}\right) \tag{2.21}$$

$$E_l = -\hbar\omega_m \left(l + \frac{1}{2}\right) \tag{2.22}$$

The average magnetron temperature is determined by the procedure used to cool the magnetron motion (described in Section 2.7), which reaches a limit when the magnetron quantum number is equal to axial quantum number. By equating the magnetron energy at this quantum number with the classical expression for the radius, we find that a well-cooled particle should have a magnetron radius of about 5 μ m.

2.4 Detecting particle motions

The motion of a charged particle in our trap induces currents between the electrodes of the trap. The induced current I is proportional to the charge e of the particle, its axial velocity \dot{z} and its distance to the electrodes z_0 . We can write this relationship as

$$I = \frac{e\kappa}{2z_0} \dot{z} \tag{2.23}$$

where κ is a geometrical constant close to 1, calculated in more detail in Section 6.2.2.

These currents are on the order of only 1-100 fA. To maximize the signal power, we maximize the impedance in series with this current. We tune out the intrinsic capacitive coupling between the electrodes with an added inductor so that the resonance is close to the particle's frequency. The tuned circuit is also connected through a capacitive divider to a cryogenic field effect transistor (FET), which provides power gain by buffering the voltage into a 50 Ω cable leading to the rest of the detection chain. The capacitive divider only partially isolates the FET from the performance of the "front end" tuned circuit, so they are often referred to and analyzed together as "the amplifier". The performance of these amplifiers is critical for the measurement, Chapter 8 describes extensive upgrades to the amplifiers on the experiment.

While so far we have discussed the axial motion, the cyclotron motion also induces counter-rotating currents $I = e\kappa_x \dot{\rho}_c / (2\rho_0)$ in the inner surface of the electrodes. We split an electrode into half-cylinders so that a tuned circuit can be connected in series with these currents so that we can detect a voltage and damp the cyclotron energy. Figure 8.1 depicts a tuned circuit in parallel with the cyclotron motion. The general principles of this section apply to both the cyclotron and axial motion.

The tuned circuit also damps the motion of the particle. First, the currents induced by the particle in the tuned circuit dissipate power at a rate of I^2R . Since the induced current is proportional to the velocity (Equation 2.23), an excited particle will lose energy from its axial motion at a rate proportional to \dot{z}^2 . Since the kinetic energy of the axial motion is proportional to \dot{z}^2 as well, the particle's energy will decay exponentially with a time constant

$$\tau = \frac{1}{2\gamma} = \frac{1}{2} \left(\frac{2z_0}{e\kappa}\right)^2 \frac{m}{R} \tag{2.24}$$

A short time constant τ is often desirable because it means tighter coupling to the amplifier and more detected signal power.

Additionally, thermal fluctuations in the voltage across the tuned circuit can drive the particle's motion. The particle's equation of motion when resonant with the amplifier is that of a damped, driven harmonic oscillator

$$\ddot{z} + \gamma \dot{z} + \omega_0^2 z = \frac{F_d(t)}{m} \tag{2.25}$$

where the driving force $F_d(t)$ is due to the thermal voltage noise across the resistor v(t) [37]. Instead of damping to zero amplitude, the particle will damp until it is in equilibrium with the noise temperature of the amplifier.

We can also treat the particle as another element in parallel with the tuned circuit. The voltage due to a noise current I_N is $v(t) = I_N R$. If we insert this into Equation 2.25 and substitute in Equation 2.24, we get an equation for the induced current:

$$\frac{dI}{dt} + \gamma I + \omega_0^2 \int I dt = \gamma I_N \tag{2.26}$$

If we take the Fourier transform of both sides,

$$\tilde{I}(\omega) = \frac{\gamma}{\left(i\left(\omega - \frac{\omega_0^2}{\omega}\right) + \gamma\right)}\tilde{I}_N(\omega)$$
(2.27)

The induced currents from the particle peak when $\omega = \omega_0$, at which point the particle carries the entire noise current. The noise signal from the tuned circuit is thus shorted

out close to the resonant frequency of the particle, forming a dip in the noise signal with a width equal to the coupling constant γ . This result can also be obtained by modeling the particle as a series LC circuit in parallel with the tuned circuit and matching the terms in Equation 2.26 with the terms of an equation for the current in an LC oscillator [38]. A similar model can be made for multiple particles. For small numbers of particles, the dip is Lorentzian with a width proportional to the number of particles [40].

2.5 Driven signals

We can also detect a peak in the amplifier spectrum from an excited particle damping energy into the tuned circuit. The minimum width of the peak is equal to the coupling constant γ . We can excite the particle directly with a drive near the resonant frequency ν_z . However, the drive itself would also be picked up by the amplifier and overwhelm the particle signal. We use a two-drive scheme to avoid this problem.

One drive is at an intermediate frequency ν_i and the other at $\nu_z - \nu_i$. The first drive modulates the trapping potential, so that the axial resonance develops sidebands at $\nu_z \pm \nu_i$, and the second drive excites the particle [37]. The particle responds at ν_z , which is the only frequency picked up by our narrow bandwidth detection system. We chose $\nu_i = 455$ kHz where commercial crystal filters are easily available. For very strong drives, the nonlinear response of the FET will produce spurious signals at at the axial frequency even if there are no particles. However, this usually occurs at least 10 dB above even the high powers typically used. Axial drives are an important tool for finding the axial frequency and optimal tuning after a cooldown. The approximately 200 Hz bandwidth of our detection system is much narrower than the possible change in the particle frequency between cooldowns. We search a larger range by keeping the drive frequency constant at the center of the amplifier, and ramping the trap voltage to sweep the antiproton frequency. We must wait several seconds after each voltage step because of long time constant filters on the trap DC biases. The shape of the swept response also gives information about the sign of anharmonic terms of the trap as discussed in [37], which we can then reduce by optimizing the voltage ratio V_0/V_c . This must be done iteratively by lowering the drive strength each time to narrow the response. Once the response looks harmonic and is narrower than 1 mV, we can either switch to the faster procedure of sweeping the drive frequency, or we can try to look directly for a dip signal.

2.6 Feedback and the self-excited oscillator

The signal to noise for dips is typically close to one and it takes several minutes of averaging to resolve a dip. Dips are also limited in their frequency resolution by the coupling constant γ . For higher signal to noise together with an extremely narrow frequency response, we turn to feedback techniques [41].

The voltage detected from the particle is proportional to the current and thus \dot{z} . If we feed back the signal as a drive to the particle with a gain G chosen so that $\frac{F_d}{m} = G\gamma_z \dot{z}$, the equation of motion for the particle (Equation 2.25) is modified to

$$\ddot{z} + (1 - G)\gamma_z \dot{z} + \omega_z^2 z = 0$$
(2.28)



Figure 2.3: A feedback drive can be applied through a gain G to adjust the coupling of the particle to the resistance R_{eff} [36]

By adjusting G, we can adjust the strength of the coupling between the amplifier and the particle. For 0 < G < 1, the coupling to the amplifier is decreased by a factor of 1 - G. The effective temperature of the particle is also reduced by the same factor [41]. This effect was used to reduce the axial temperature of the particle for the antiproton measurement described in Chapter 3.

In principle, by setting G=1 we could completely cancel out the effect of damping. This would allow an excited particle to maintain a large constant amplitude giving a an extremely narrow and strong signal on the amplifier. This is the basic idea of the the self-excited oscillator technique [42].

In practice, we must use adaptive feedback control rather than simply finding the "right" gain and phase. This is evident if we consider that the amplitude of the particle will exponentially increase or decrease if G is different from 1 by any amount, no matter how tiny. We avoid this by continuously adjusting G to counteract fluctuations in the amplitude of the axial peak. The axial peak amplitude is measured by a dedicated digital signal processor (DSP) which performs a fast fourier transform of



Figure 2.4: Radiofrequency schematic for applying either a two-drive scheme or feedback for the self-excited oscillator

the axial signal. It outputs an analog signal that controls a voltage variable attenuator. With the active control of the DSP, we are able to maintain a large SEO signal with a width an order of magnitude smaller than γ [35].

Figure 2.4 shows the radiofrequency components used to detect and drive the axial motion of the particle. A switch on the bottom endcap lets us change modes between driving the particle from a frequency generator as described in the last section and driving it with feedback controlled by the DSP.



Figure 2.5: Energy levels with an axial-magnetron coupling drive. (a) The two transitions from a cooling drive (left) and a heating drive (right). (b) Transitions involved in calculation of the steady state limit of sideband cooling. [35]

2.7 Magnetron cooling

The magnetron motion cannot be cooled with resistive damping techniques. Since the motion is unstable, any radial damping force on the particle would cause an exponential increase in the radius until the particle leaves the trap. We are careful to make sure that any electrodes split for the application of radial forces at the cyclotron or axial frequency are shorted at the kHz scale magnetron frequency through an inductor (resulting in a time constant of magnetron expansion on the order of years). Despite this, it is still possible for the magnetron radius to increase due to noise or transfers. In addition it is not unusual for particles to be initially loaded with a large magnetron radius. In those situations, we use a sideband cooling technique to reduce the magnetron radius [37].

To couple the magnetron and axial motions, we apply a drive at $\nu_z + \nu_m$ on one half of a split compensation electrode. This creates an oscillating field near trap center
proportional to $xz \cos(\omega_z + \omega_m) t$. The drive causes a splitting of the axial energy levels, shown in Figure 2.5, and transitions from $|kl\rangle \rightarrow |k+1, l-1\rangle$ and $|kl\rangle \rightarrow$ $|k-1, l+1\rangle$. For l > k transitions to lower magnetron quantum number happen at a faster rate than transitions in the other direction. The magnetron quantum number decreases until l = k at which point the two processes are in equilibrium. The energy transferred into the axial motion is damped by the axial amplifier at temperature T_z , so the minimum magnetron temperature achievable by the process is can be found from:

$$T_m = \frac{\omega_m}{\omega_z} T_z \tag{2.29}$$

For a particle starting from a large magnetron radius, the axial frequency will increase as the particle approaches trap center. We increase the trapping potential in several steps keep the drive resonant when cooling particles that start in a large magnetron orbit.

2.8 Magnetic bottle

To detect the spin state, we deliberately introduce a strong magnetic gradient that couples the axial motion to the angular momentum. The basic spin flip detection principle is to measure the axial frequency, apply a spin flip drive, and determine if the drive shifted the frequency by the characteristic amount due to the spin flipping.

A quadratic magnetic gradient or "bottle" shifts the magnetic field by [37]

$$\Delta \vec{B} = B_2 \left[\left(z^2 - \rho^2 / 2 \right) \hat{z} - \left(\hat{B} \cdot \vec{z} \right) \vec{\rho} \right]$$
(2.30)

The interaction of the particle's magnetic moment μ with the bottle adds a term to



Figure 2.6: (a) Schematic of the analysis trap, with the iron ring shown in grey (b) Magnetic field along the central axis of the trap [36]

the Hamiltonian

$$\Delta H = -\vec{\mu} \cdot \Delta \vec{B} \tag{2.31}$$

For a well-cooled particle, we assume $\rho \approx 0$ and this simplifies to

$$\Delta H = \mu B_2 z^2 \tag{2.32}$$

Since this has the same z^2 dependence as the electrostatic potential, it effectively shifts the frequency of the axial motion ω_z to

$$(\omega_z')^2 = \omega_z^2 + \frac{2\mu B_2}{m}$$
(2.33)

The magnetic moment of a particle in our trap includes the moment due to the cyclotron and magnetron motions as well as the intrinsic magnetic moment of the proton's spin, $\vec{\mu}_S = g\mu_B \vec{S}/\hbar$. When the spin state flips from up to down, the frequency will shift by

$$\delta\omega_z \approx \frac{2\mu_p B_2}{m\omega_z} \tag{2.34}$$

or about 130 mHz in our analysis trap.

We generate a strong bottle gradient in our analysis trap by using a ring machined out of iron. In the strong magnetic field of the solenoid, the ferromagnetic material saturates saturates and produces a field with a profile shown in Figure 2.6. The field due to the iron ring can be calculated from the solution to Laplace's equation for the magnetic potential inside the trap [35, 37]. The leading term in the expansion of the magnetic field near trap center has a similar form to Equation 2.33. For the analysis trap, $B_2 = 290,000 \text{ T/m}^2$. Out of the 920 kHz analysis trap axial frequency, a single spin flip shifts the axial frequency by 130 mHz.

The axial frequency in the bottle can also shift due to changes in the magnetic moment from the cyclotron and magnetron motions. The shift due to a spin state s, a cyclotron state n, and a magnetron state l is [43]:

$$\Delta\omega_z = \Delta\tilde{\omega}_z \left(\frac{gs}{4} + n + \frac{1}{2} + \frac{\omega_m}{\omega'_c} \left(l + \frac{1}{2}\right)\right) \tag{2.35}$$

where

$$\frac{\Delta\tilde{\omega}_z}{\omega_z} = \frac{\hbar}{2m\omega_m} \frac{B_2}{|B_0|} \frac{\omega_c}{\omega'_c - \omega_m}$$
(2.36)

Each cyclotron transition $|n\rangle \rightarrow |n \pm 1\rangle$ causes a shift that only differs from the spin flip transition by a factor of 2/g. Three cyclotron transitions together (50 mHz each) can mimic a spin flip (130 mHz). This means the cyclotron energy must be stable to better than 0.1 percent, a problem discussed in Section 3.2.1. The shift due to the magnetron motion is a factor of $\frac{\omega_m}{\omega_c}$ smaller, or 3 μ Hz per quanta. It still must be taken into account for a precise measurement, and is discussed in Section 2.10.



Figure 2.7: Schematic of currents through the compensation electrode for driving spin flips

2.9 Driving spin flips

We must be able to drive spin state transitions in order to measure the spin flip frequency. We accomplish this by generating an oscillating magnetic field perpendicular to the Penning trap z-axis. We run the drive current through both halves of the split compensation electrode as shown in Figure 2.7. The net magnetic field from this current at trap center $\vec{B_1}$ is perpendicular to the solenoid field. This field can be written as the sum of co-rotating and counter-rotating terms:

$$\vec{B}_{1} = \frac{|B_{1}|}{2} \left(\hat{x} \cos\left(\omega_{s}t\right) - \hat{y} \sin\left(\omega_{s}t\right) \right)$$
(2.37)

The co-rotating term will drive spin flips with a Rabi frequency

$$\Omega_R = \frac{|B_1|\,\mu_p}{\hbar} \tag{2.38}$$

For the current trap geometry, we can calculate B_1 using the Biot-Savart law and the assumption that most current is carried on the inner surface of the electrode.

$$B_1 = \frac{\mu_0 \rho_0 I}{2z_c} \left(\frac{1}{\sqrt{z_0^2 + \rho_0^2}} + \frac{1}{\sqrt{(z_0 - z_c)^2 + \rho_0^2}} \right)$$
(2.39)

This yields 1.4 Gauss/A in the analysis trap, or a Rabi frequency of 18 kHz per Ampere.

As long as the 25 kHz bottle linewidth ($\Delta \omega_s / (2\pi)$) is large compared to the Rabi frequency, we are not able to coherently drive spin flips. Instead, the probability of driving a transition after a time Δt is given by [36]

$$P_{sf}\left(\Delta t\right) = \frac{1}{2} \left(1 - \exp\left(\frac{-\pi \Omega_R^2 \Delta t}{\Delta \omega_s}\right)\right)$$
(2.40)

With a strong enough drive, this will saturate at 50 percent probability.

The Rabi frequency is limited by the current we can apply to the electrodes. At 220 MHz, it is difficult to prevent a large fraction of the power applied from dissipating in other parts of the experiment. At high drive power, there is a noticeable shift in the axial frequency of the particle, which is attributed to heating of the experiment from this lost RF power. Weaker drives were used during a measurement of the proton g-factor in 2011 to avoid this systematic effect, but this limited the speed of data accumulation [25].

Shortly before the measurement in 2012, a coaxial matching line was installed inside the trap can. The coaxial line minimizes the impedance of the path to the electrodes, increasing the current flowing through them for a given drive power. For the antiproton measurement, we were able to increase the drive strength while the frequency shift remained at the ppb level or below (shown in Figure 2.8).



Figure 2.8: Axial frequency shift as a function of the drive applied at the hat. The dashed line is a fit to data taken before the new transmission line transformer was installed [36]

2.10 Lineshape

The bottle also has the undesirable effect of broadening the cyclotron and spin flip transitions. In the bottle gradient, the magnetic field reaches a minimum at trap center. Particles with a larger amplitude will thus experience a stronger average field over the course of each oscillation. For both the spin and cyclotron frequencies, the bottle shifts the gradient-free frequency ω_0 to

$$\omega\left(z\right) = \omega_0 \left(1 + \frac{B_2}{B_0} z^2\right) \tag{2.41}$$

Thermal fluctuations in the axial motion caused by its coupling to the amplifier (Section 2.4) thus cause fluctuations in the spin and cyclotron frequencies. Since the standard deviation of a thermal distribution is equal to the mean, we naively expect

the standard deviation of the transition frequency to be

$$\Delta\omega = \omega_0 \left(\frac{B_2}{B_0}\right) \left\langle z_{rms}^2 \right\rangle \tag{2.42}$$

This width $\Delta \omega$ is referred to as the linewidth parameter. We can find z_{rms}^2 using Equation 2.17. In the analysis trap, $\Delta \omega$ is typically about $2\pi \times 25$ kHz for the spin transition at 220 MHz, and $2\pi \times 8$ kHz for the cyclotron transition at 79 MHz. Since the rate of damping to the axial amplifier $\gamma \approx 1$ second is much less than $\Delta \omega$, the axial amplitude and magnetic field are essentially constant over the time needed for a single spin flip trial. During the much longer time in between trials the particle reequilibrates with the amplifier, and the average amplitude and thus field may be very different for the next trial. If we make many trials at different possible frequencies, the lineshape mirrors the exponential distribution of axial energies with a sharp rise at ω_0 corresponding to a zero-amplitude particle:

$$\chi(\omega) = \frac{\Theta(\omega - \omega_0)}{\Delta\omega} \exp\left[-\frac{\omega - \omega_0}{\Delta\omega}\right]$$
(2.43)

where $\Theta(\omega)$ is the Heaviside step function. This expression for the lineshape when $\gamma \ll \Delta \omega$ and for other cases are derived in [37]. To measure g we need to compare ω_c and ω_s at the same amplitude. We do this by carefully measuring below and above the sharp rise to determine the zero-amplitude frequencies with high precision.

A similar analysis applies to the magnetron motion [36]. We typically side-band cool before each spin flip trial to minimize the magnetron radius. Each application of the cooling drive equalizes the axial and magnetron quantum numbers as described in Section 2.7. The distribution of axial states from trial to trial will be reflected in the resulting distribution of magnetron states. A nonzero magnetron radius shifts the



Figure 2.9: (a) The spin or cyclotron lineshape due to axial thermal fluctuations in the bottle (b) The lineshape due to magnetron fluctuations is similar but reversed. [36]

field $\Delta \vec{B} = B_2 \left(-\rho_m^2/2\right) \hat{z}$, which shifts the frequency by

$$\omega = \omega_0 \left(1 - \frac{B_2}{B_0} \rho^2 \right) \tag{2.44}$$

The resulting linewidth is calculated in the same way as for the axial motion, but scaled by $T_m = \frac{\omega_m}{\omega_z} T_z$. Because of the negative sign of the radial bottle potential, the frequency increases up to the maximum given by a zero-energy particle. The final lineshape is the convolution of the magnetron lineshape and the axial lineshape, illustrated in Figure 2.9. The magnetron lineshape is much narrower, but it limits the sharpness of the initial rise and thus contributes directly to the uncertainty in determining ω_0 . The measured lineshape and its uncertainty is presented in Section 3.4.

2.11 Apparatus

This section provides a brief overview of the apparatus as used in the ppm measurement of Chapter 3. That version of the apparatus is documented in more detail in [35] and [36]. A substantial program to rewire and upgrade the apparatus is described in Chapters 7 and 8.

2.11.1 Cryogenic system and electronics

This section gives an overview of the parts of the apparatus common to the old measurement and the rebuilt experiment. The magnetic field of the Penning trap is produced by a 5.6 T superconducting persistent solenoid cooled by liquid helium. The 4" diameter bore of the magnet is a separate vacuum space so that the magnet can be left on if the trap is removed for repairs and upgrades. When the trap is inserted into the bore, an O-ring seal is made between the hat and the magnet top plate and the bore space is pumped out. Then liquid nitrogen is filled into a 50 L auxiliary magnet dewar which is heat sunk to the copper inner wall of the bore.

The experimental insert holding the trap has five main sections: the hat, the baffles, the dewar, the tripod, and the trap can. Everything is supported vertically from a vacuum flange called the hat that mates to the magnet top. The hat has a liquid helium fill port on top and eight KF vacuum flanges around its diameter to allow electrical and vacuum connections. The thermal baffles hang from G10 rods bolted to the bottom of the hat. These consist of three copper baffles at different temperatures which provide thermal radiation shielding to the colder parts of the experiment. The top baffle has a cylindrical copper shield hanging from it that surrounds the lower baffles. This thermal shield has beryllium copper fingers at the bottom and extends far enough down that the fingers press against the liquid nitrogen cooled bore.

The copper liquid helium dewar holds 3.6 L [35]. The region below it is called the

tripod after the three OFHC support posts that connect the dewar to the pinbase. These posts also provide heat sinking for the electronics mounted to them. Space in the tripod region is limited as we need to fit large electrical filters as well as the amplifiers for several traps. At the bottom of the tripod is an OFHC vacuum flange with 3 larger high voltage feedthrough pins and 51 standard vacuum feedthrough pins.

The trap can is a separate vacuum space containing the electrodes of the trap. The top flange makes an indium seal to the bottom of the pinbase. After this seal is made, the titanium trap can is pumped out to roughly 10^{-7} Torr through a copper flange at the bottom, then a pinch off tool is used to cold weld the pump out port shut. On cooldown, cryopumping of the closed volume creates an extremely high vacuum that has allowed us to store a single antiproton for more than a month. A pressure better than 5×10^{-17} Torr was achieved in a similar system [44]. Such high vacuum will be essential for the long antiproton lifetimes we will need for a precision measurement.

2.11.2 Traps

The electrode stack arrangement used for the antiproton measurement (described in Chapter 3) is shown in Figure 2.10. In addition to the analysis trap used for the measurement, there is a second trap called the precision trap at the bottom of the stack. The larger, low-gradient precision trap was used to cool antiprotons immediately after loading. We also hope to use the precision trap for a future measurement at higher precision as described in Section 4.1. Antiprotons are loaded from the bot-



Figure 2.10: Solidworks model of the trap stack as used in the antiproton magnetic moment measurement [36]

tom of the stack and cooled in the precision trap before transfer to the analysis trap. The field emission point at the top of the stack fires electrons. These can ionize adsorbed gas from the degrader, allowing us to load electrons, ions and protons in this apparatus as well.

Chapter 3

A part-per-million antiproton measurement

We measured the antiproton magnetic moment to 4.4 ppm [1], a 680-fold improvement in precision over the previous best measurement [14]. On a very short time scale, the measurement required overcoming several experimental challenges to load antiprotons into our proton apparatus and achieve sufficient axial frequency stability to measure spin flips. This chapter describes the measurement procedure and result.

3.1 Loading antiprotons

Antiprotons are provided by the Antiproton Decelerator (AD) facility at CERN in Geneva. Every 120 seconds, about 35 million antiprotons are ejected to our experimental beamline at 5 MeV. Antiproton loading was first demonstrated, in 1986, by



Figure 3.1: Diagram of PPAC tube showing the beam passing through the beam position detector and the energy tuning cell on its way to the trap.

our group [45]. The techniques are well-developed as reviewed in [46]. This section provides a brief outline of the trapping procedure, focusing on the particular challenges faced by this experiment. These challenges informed the design of an upgraded apparatus described in Section 8.6.

3.1.1 Slowing antiprotons

The energy of the antiprotons need to be reduced by three orders of magnitude to the keV range before it is feasible to catch them in an electrostatic trap. Before entering our trap, the beam passes through a beam position detector and a gas cell.



Figure 3.2: Antiprotons trapped in the high voltage well as a function of delay in closing the trap relative to the CERN warning trigger[36]

The beam position detector is a parallel plate avalanche counter (PPAC), described in detail in [47] and shown in Figure 3.1. Just after the PPAC, the beam passes through a gas cell containing an adjustable mixture of SF_6 and helium at 1 atm. The purpose the gas cell is to modify the amount of energy absorbed from the beam. When the cell is filled with helium, the cell absorbs nearly no beam energy, but as the fraction of SF_6 flowing through the cell is increased the absorbed energy increases up to a maximum of roughly 0.5 MeV.

Most of the beam's energy is absorbed in the degrader, which is a 125 μ m thick sheet of beryllium mounted at the entrance of the trap stack. That thickness is carefully chosen so that the average beam energy will be reduced to zero with a 50:50 He/SF₆ ratio in the gas cell [48] in order to maximize the energy range of the cell.

The low end of this distribution has energies of only a few hundred eV. The degrader bias voltage is kept low when the beam arrives to allow these low energy

particles to exit into the trap stack. We place a blocking potential of -300 V partway up the electrode stack, usually on T5 (see Figure 3.4). Antiprotons with less than 300 eV of energy will be reflected when they reach this part of the stack. The bias applied to the degrader is controlled by a fast, high voltage switch that can lower its voltage from +100 V to -300 V before the reflected antiprotons have hit the degrader again, leaving them trapped between two strong negative potentials referred to as the "HV well". Figure 3.2 shows the sharp rise in the number of antiprotons caught when the delay from the AD shot warning signal is long enough that the antiprotons have time to reach the degrader before the trap is switched closed. The number caught starts declining when the delay is extended by more than 600 ns as faster antiprotons then have time to leave the stack after reflection before the trap is closed.

3.1.2 Detection

We can check for successful trapping by removing the trapping potentials and looking for annihilation signals from antiprotons leaving the trap and hitting the electrodes. This is an essential technique for quick optimization of the multiple parameters necessary for successful catching. When antiprotons annihilate with protons in the trap walls, they generate several charged pions that pass through the rest of the material in the trap stack and magnet. The pions generate light in the scintillating paddles arranged around the magnet which is detected by PMTs.

The total efficiency of the paddles is about 50 percent [49], and there is a background signal from cosmic rays that typically causes on average 2 counts during a 10 ms ramp. As little as 10 or more antiprotons in a short well-timed ramp can provide



Figure 3.3: The 'long well' potential originally used to trap antiprotons

a clear signal above background, making this a sensitive test for successful trapping. As we typically only caught 10-100 particles at a time, these signals are less useful as a passive monitor of particle loss as slow losses are not visible above the background.

3.1.3 Electron cooling

The energy of the antiprotons must be further reduced after trapping from hundreds of electron-volts (> 11,500 K) to 4 K. This is done by co-trapping a cloud of electrons which sympathetically cool the antiprotons [50]. The electrons damp into thermal equilibrium with the walls of the trap through synchrotron radiation within minutes.

Initially we put the blocking electrode on the far side of the precision trap. After finding the right parameters to catch antiprotons in the HV well between the blocking potential and the degrader, we added a harmonic well on the precision trap as shown in Figure 3.3. When we loaded electrons into the harmonic well before the shot, we did not detect any antiprotons leaving when we grounded the degrader at the end of the cooling period. However, we also were not able to detect any evidence of antiprotons in the precision trap. Presumably the antiprotons were leaving the long



Figure 3.4: The 'short well' potential that successful caught antiprotons and sympathetically cooled them with electrons.

well as they cooled, but slowly enough that in any given time bin the signal was not distinguishable from background.

One possibility that is that cold particles in a long flat well are more susceptible to magnetron heating. Such behavior has been observed by [51], which suggests that in regions of flat potentials, torques due to residual electrical fields from patch potentials and other imperfections will overwhelm the weak trapping fields. Hot particles would be less sensitive since they would still spend much of their oscillation time near the trap edges the electric field is large. However, as they cool to the center of a flat well the effect of patch potentials would increase. Dramatically shortening the well by moving the blocking electrode to T5 as shown in Figure 3.4 solved the problem.

After the particles are cooled, antiprotons and electrons are transferred together to the precision trap using an "inchworm" technique. A deeper potential is placed on the next electrode up the stack, then the electrode currently holding the particles is zeroed to push them into it. The original voltage is subsequently restored and the process is repeated on the next electrode.

Upon arrival in the precision trap, we immediately eject the electrons from the trap. Otherwise, they can interfere with signals from the particles. First, a ramp of

increasing potentials is applied to T5, EXT, and the degrader to guide the particles away from the trap. We send a 10 V pulse from a saturated switch driven by a pulse generator down the axial drive line to the bottom endcap. This opens one side of the trap for the 100 ns duration of the pulse. The electrons leave the trap, but the heavier, slower antiprotons do not have time to leave before the pulse ends. After a few initial pulses, a magnetron-axial cooling drive re-centers the electrons before ten more ejection pulses are applied.

Before detecting radiofrequency signals in the precision trap, the antiproton magnetron radius needs to be reduced with two sideband ramps. Then we excite the cyclotron motion to count the number of particles, and repeatedly dip the trap until we have a single antiproton as described in [36]. Once the particle is well cooled in all three motions in the precision trap, it can be transferred to the analysis trap where we can detect spin flips.

3.2 Detecting spin flips

We detect spin flips by measuring the 130 mHz shift in the 920 kHz axial frequency due to the bottle, as described in Section 2.8. We use the self-excited oscillator described in Section 2.6 to achieve the required frequency precision, but the small size of a spin flip also presents a stringent requirement on the stability of the axial frequency. Before describing the process for achieving good stability, we will detail how to characterize and measure it.

First, we typically look at the stability between consecutive measurements only. While the electron g-factor measurement can infer the spin state directly from the



Figure 3.5: (a) Axial stability over several hours in the analysis trap (b) Axial stability as a function of averaging time [36]

axial frequency [52], in our analysis trap the axial frequency trap drifts by much more than a spin flip over long periods of time. Figure 3.5a shows example data from repeated measurements of the axial frequency with the self-excited oscillator where the frequency drifts by nearly a Hz over several hours. However, the frequency differences between consecutive measurements are much smaller. We can isolate the effect of the spin flip drive by measuring the frequency before and after driving a spin transition and taking the difference.

We can characterize the point to point scatter with the standard deviation of these differences σ . We also often refer to the Allan deviation, a standard parameterization of the frequency noise of an oscillator that happens to be equal to σ aside from a factor of $\sqrt{2}$. The Allan deviation of a series of N frequency difference measurements with differences $\Delta_i = f_{i+1} - f_i$ is given by

$$\sigma_A = \sqrt{\frac{1}{2N} \sum_{i=1}^{N} \left(\Delta_i\right)^2} \tag{3.1}$$

When we don't apply a drive, we have a natural Allan deviation σ_0 due to the drifts

and averaging time. If we apply a drive, σ_A increases because on top of the normal scatter there will be jumps in frequency due to spin flips. In the presence of a spin flip drive that flips the spin with a probability P, the total Allan deviation of a set of measurements σ_f is [36]

$$\sigma_f = \sqrt{\sigma_0^2 + \frac{P}{2}\Delta_s^2} \tag{3.2}$$

We can solve this expression to find the fraction of spins flipped in a dataset based on the measured Allan deviation:

$$P = \frac{2\left(\sigma_f^2 - \sigma_0^2\right)}{\Delta_s^2} \tag{3.3}$$

We isolate the effect of the drive by cycling between measuring the frequency before and after applying a near-resonant drive and an off-resonant drive. The second set of differences acts as a control, which we monitor for any changes or increased scatter not due to spin flips. We bin both sets of data into histograms like the ones shown in Figure 3.6, and fit a Gaussian to confirm the normal distribution of the scatter.

As long as σ_0 is Gaussian, the statistical error in the measurement of the Allan deviations is [36]

$$\delta\sigma_A = \frac{\sigma_A}{\sqrt{2N-2}} \tag{3.4}$$

By driving at several frequencies along the resonance, we can measure the spin flip probability at each point and thus the lineshape to a desired statistical precision.

3.2.1 Instability due to cyclotron transitions

The dominant source of axial frequency instability in the analysis trap are fluctuations in the cyclotron state, which also couple to the axial frequency through the



Figure 3.6: (a) Histograms of frequency differences for near resonant drive in outline $(f_1 - f_2)$ with histogram of the control drive $(f_3 - f_1)$ shown in grey. The dashed line is a Gaussian fit to the control drive, and the solid line is the Gaussian fit to the near resonant drive. (b)The outline is a histogram on non-resonant control drive data $f_2 - f_3$ and the grey is the control data $(f_3 - f_1)$. The fits for the two datasets overlap so only the solid line fit to the non-resonant drive is visible.[36]

magnetic bottle. We can deduce this from the dependence of the Allan deviation on the averaging time and cyclotron quantum number.

First, the axial frequency is drifting on timescales similar to our typical averaging time. We can see this from Figure 3.5b. It shows that as the averaging time T increases, the scatter first decreases as electronic noise is reduced, as would be expected from the 1/T resolution of a fast Fourier transform. However, at longer averaging times the Allan deviation starts increasing. This is consistent with a slower drift that increases as \sqrt{T} . We typically find the optimal averaging time where these two effects cross over before taking a dataset.



Figure 3.7: Allan deviation as a function of cyclotron quantum number. Dashed line indicates a fit to \sqrt{n} [36]

3.2.2 Cyclotron noise drive

An important feature of this drift is that its rate depends on the cyclotron quantum number. Figure 3.7 shows the Allan deviation of the axial frequency for a constant 60 second averaging time as a function of cyclotron quantum number. To make this plot, we excite the cyclotron motion in the precision trap then move the particle to the analysis trap before it has a chance to damp its energy into the precision cyclotron amplifier. The cyclotron state can be measured on arrival in the analysis trap because changes of thousands of quanta of cyclotron energy will cause shifts of 1-100 Hz. The drift rate is proportional to \sqrt{n} , which is shown on the plot with a dashed line.

The \sqrt{n} dependence is consistent with cyclotron state transitions driven by small voltage fluctuations on the electrodes. We derive the axial frequency fluctuations

expected from random electric field noise, and find the observed stability could be explained by voltage fluctuations not much larger than known sources of Johnson noise.

The cyclotron transition rate due to an electric field can be calculated using perturbation theory. The Hamiltonian due to the homogeneous magnetic field is the standard ladder of states of a harmonic oscillator, $H_0 = \omega_c \left(a^{\dagger}a + \frac{1}{2}\right)$. Here a^{\dagger} and aare the creation and annihilation operators for the cyclotron quantum number. We assume that the particle radius is constant over the time needed to measure the axial frequency. For a general calculation, we account for the effects of the bottle by allowing the cyclotron frequency $\omega_c(t)$ to randomly vary in time as the axial amplitude fluctuates. A time-varying electric field $\epsilon(t)$ adds a term $H_1(t) = qx\epsilon(t)$ to the Hamiltonian. In the notation of ladder operators [37], $x = \sqrt{\frac{\hbar}{2m\omega_c(t)}} \left(a^{\dagger} - a\right)$. Then if the particle starts in state n, to a first order approximation the rate of transitioning up to state n + 1 is

$$\frac{d\left|c_{n+1}\left(t\right)\right|^{2}}{dt} = \int_{0}^{t} \frac{q^{2}\left(n+1\right)}{2m\omega_{c}\left(t\right)\hbar} \left\{\epsilon\left(t\right)^{*}\epsilon\left(t'\right)\exp\left[i\int_{t'}^{t}\omega_{c}\left(\tau\right)d\tau\right] + c.c.\right\}dt' \qquad (3.5)$$

and the rate of transitioning down to state n-1 is

$$\frac{d\left|c_{n-1}\left(t\right)\right|^{2}}{dt} = \int_{0}^{t} \frac{q^{2}n}{2m\omega_{c}\left(t\right)\hbar} \left\{\epsilon\left(t\right)^{*}\epsilon\left(t'\right)\exp\left[-i\int_{t'}^{t}\omega_{c}\left(\tau\right)d\tau\right] + c.c.\right\}dt' \quad (3.6)$$

If the cyclotron frequency is constant in time, and we only consider transitions from state 0 to 1, this reduces to a common formula for heating in ion trapping experiments [53]:

$$\left\langle \frac{d \left| c_{0}\left(t\right) \right|^{2}}{dt} \right\rangle = \frac{q^{2}}{4m\omega_{c}\hbar} S_{E}\left(\omega\right)$$
(3.7)

where $S_E = 2 \int_{-\infty}^{\infty} d\tau e^{i\omega_c \tau} \langle \epsilon(t) \epsilon(t+\tau) \rangle.$

For easy comparison to known noise sources, we can approximate the effect of a noise source that has a flat frequency spectrum with a root mean squared amplitude of $v_{rms} V/\sqrt{\text{Hz}}$. Such a source would produce a field $\langle \epsilon(t') \epsilon(t)^* \rangle = \left(\frac{e\kappa}{2\rho_0}\right)^2 v_{rms}^2 \delta(t-t')$ at trap center. The Dirac delta function greatly simplifies evaluation of the integral, which yields¹

$$\frac{d\left|c_{n+1}\left(t\right)\right|^{2}}{dt} \approx \frac{q^{2}\left(n+1\right)}{2m\omega_{c}\hbar} \left(\frac{e\kappa}{2\rho_{0}}\right)^{2} v_{rms}^{2}$$

$$(3.8)$$

The rate of transitioning down $d |c_{n-1}|^2 / dt$ is almost identical, except proportional to n rather than n+1. This means transitions up occur at a faster rate than transitions down, agreeing with the intuition that noise should heat the particle. However, since $n \approx 2000$ for a thermal particle, these two rates are almost identical in our experiment. It would take months for a thermal (8 K) particle to double in temperature.

Instead, we are concerned with the total rate of transitions out of the initial state, given by the sum of those two transition rate rates. We can approximate this as

$$\frac{dn}{dt} = \frac{n}{\hbar\omega_c} \left(\frac{e\kappa}{2\rho_0}\right)^2 \frac{v_{rms}^2}{m} \tag{3.9}$$

. The particle will like transition up and down several times over the time to average the axial signal, effectively causing a small drift in the axial frequency over the averaging time. For large n, this is essentially a one-dimensional random walk over the ladder of cyclotron states. The average distance is proportional to the square root of the number of steps taken over the averaging time T. The variance of the frequency change over an averaging time is then

$$\sigma^2 = n\Gamma T \Delta_c^2 \tag{3.10}$$

¹The contribution from the exponential term goes to 1 due to the delta function, but we retain $\omega_c(t)$ in the denominator. We approximate this as a constant as the cyclotron frequency varies by at most parts per million due to the bottle

where the transition rate $\frac{dn}{dt}$ has been divided into n and $\Gamma = \frac{dn}{dt} \frac{1}{n}$ and $\Delta_c = 50$ mHz is the frequency shift per cyclotron quantum. If we substitute Equation 3.9 into this, we find that the standard deviation is

$$\sigma = \sqrt{\frac{T}{\hbar m \omega}} \left(\frac{e\kappa}{2\rho_0}\right) \times v_{rms} \Delta_c \sqrt{n} \tag{3.11}$$

Since the Allan deviation is proportional to $\sigma/\sqrt{2}$, we have found the \sqrt{n} dependence shown in Equation 3.7.

We can also use Equation 3.11 to solve for the equivalent noise amplitude v. We can extract $\Gamma = .000021$ transitions per second per quanta from the fit in Figure 3.7. This corresponds to $v = 2.5 \text{ pV}/\sqrt{\text{Hz}}$. For comparison, the Johnson noise from a 50 cable at room temperature is roughly 1 nV/ $\sqrt{\text{Hz}}$. This motivated the addition of a Chebyshev high pass filter to the analysis spin flip line, which successfully improved our stability by a factor of 3. The new filter attenuates the cyclotron frequency by more than 80 dB [36]. Our observed voltage noise of approximately 2.5 pV/ $\sqrt{\text{Hz}}$ is also equivalent to a 0.1 resistor at 4 K without any filtering. This illustrates the importance of minimizing loss in the cold electronics and filtering electrical lines as close to the trap as feasible.

3.2.3 Minimizing the scatter

After minimizing the source of the noise, we can minimize the scatter of a given particle by selecting a low cyclotron quantum number. We estimate its cyclotron number based on the frequency in the analysis trap. If the shift is larger than desired, we send the particle back to the precision trap to thermal with the cyclotron amplifier there for several damping times. Then we send it back to the analysis trap, isolating



Figure 3.8: Measurement steps vs time for a single data point on the spin line [1]

it from the amplifier and randomly selecting a new cyclotron state from the thermal distribution of the amplifier. We repeated this process until the quantum number was less than 2000 for the antiproton measurement.

3.3 Measurement sequence

Both the spin and cyclotron frequency can be measured through repeated measurements of the axial frequency. The data taking procedure is shown in Figure 3.8. It is also described in detail in [36] and [1]. For each point, a cycle of three axial frequency measurements using the SEO is performed. A spin flip drive is applied between the first two measurements f_1 and f_2 , and a drive far off resonance with the same power is applied between the second and third measurements f_2 and f_3 . Between f_3 and f_1 of the next cycle we apply sideband cooling to control the magnetron radius. The frequency differences $f_2 - f_3$ are collected into a "control" histogram, which tells us the size of background fluctuations not due to spin flips. We use a



Figure 3.9: Data points and fit from the measurement of (a) The spin flip line and (b) the cyclotron line [36]

detuned drive so we can account for any extra scatter due to the strong power of the drive. This is compared to the distribution of $f_2 - f_1$, which are collected into the "drive" histogram. When the drive successfully causes spin flips, the drive histogram will be broader than the control histogram and the histogram for the off-resonant drive.

The cyclotron frequency is measured in an analogous manner. A weak resonant drive was used that caused the cyclotron state to fluctuate on average by 8 quanta, or about 400 mHz. The larger frequency difference allowed us to track a dip signal in place of the SEO, and achieve statistical significance with fewer data points at each frequency. See Figure 3.9b for the cyclotron data used in the measurement. The frequency at which the scatter increases abruptly is the trap-modified cyclotron frequency. We can extract the free-space cyclotron frequency from this using the Brown-Gabrielse Invariance theorem given in Equation 2.14. We took data for 2-4 hours at each point on the cyclotron line, for 1.5 total days of data taking. In contrast, each point for the spin flip line took 24-48 hours of data taking and three weeks to take the entire dataset.

3.4 Antiproton measurement

Using the procedure described in the previous section, we can measure the spin frequency f_s and the cyclotron frequency f_c . We then compute the g-factor from their ratio according to

$$\frac{g_{\overline{p}}}{2} = \frac{f_s}{f_c} \tag{3.12}$$

For direct comparison to the proton g-factor, we would like to compute the ratio of the antiproton magnetic moment to the nuclear magneton, which depends on the proton charge and mass:

$$\frac{\mu_{\overline{p}}}{\mu_N} \equiv \frac{g_{\overline{p}} \, q_{\overline{p}}/m_{\overline{p}}}{2 \, q_p/m_p} \approx -\frac{g_{\overline{p}}}{2} \tag{3.13}$$

This is equal to the g-factor within the precision of this measurement, thanks to measurements of the charge to mass ratio at the 0.1 ppb level before this measurement[54], and even more accurate measurements since [55].

Our result for the \overline{p} magnetic moment [1] is

$$-\frac{g_{\overline{p}}}{2} = -2.792845(12) \,[4.4\text{ppm}] \tag{3.14}$$

From which we can calculate the ratio of the proton to antiproton magnetic moments:

$$\frac{\mu_{\overline{p}}}{\mu_p} = -1.000000 \pm 0.000005 \tag{3.15}$$

consistent with CPT symmetry.

The uncertainty in these measurements is dominated by the spacing between frequencies chosen to map out the lineshape. The measured spin and cyclotron line shapes are shown in Figure 3.9. The uncertainties for the spin and cyclotron frequency extracted from those lines is shown in Table 3.1. The uncertainties in the axial and magnetron frequency contribute much less than a ppm uncertainty to the calculation of the free space cyclotron frequency from the Brown-Gabrielse Invariance Theorem and are not included. The resonance frequency uncertainty is determined by the half-width between the points measured at the top and bottom of the lead-ing edge. In addition, magnetron broadening blurs the line at the sharp leading edge, adding an uncertainty equal to the linewidth calculated in Section 2.10. The magnetron uncertainty was lower by a factor of 2 for the cyclotron measurement, because feedback cooling was used during the sideband cooling step to lower the axial temperature and narrow the distribution of magnetron states.

Table 3.1: Uncertainties for the antiproton magnetic moment measurem	ent.
--	------

Resonance	Source	ppm
spin	resonance frequency	2.7
spin	magnetron broadening	1.3
cyclotron	resonance frequency	3.2
cyclotron	magnetron broadening	0.7
total		4.4

Chapter 4

Single spin flips

4.1 Introduction

The measurement of last chapter was limited by the large linewidth in the analysis trap due to the strong magnetic bottle. The precision of the measurement can be improved by several orders of magnitude if we apply the spin and cyclotron drives in a separate, low-gradient precision trap, and then move the particle to the analysis trap to measure the change in the spin state. This is known as the double-trap technique.

There are four parts to this measurement:

- 1. Preparing the particle in a known spin state
- 2. Sending the particle to the precision trap and applying drives to excite spin and cyclotron transitions
- 3. Cool the particle until the standard deviation σ_0 is low enough.
- 4. Measuring the spin state again in the analysis trap

For step 1 and 4, we need to detect individual spin flips as they occur rather than the average spin flip rate over a large dataset. To do this, the fluctuations in the axial frequency must be smaller than the shift due to a single spin flip Δ_s . This chapter presents a dataset with sufficient stability that we were able to identify individual spin flips [2]. The rest the chapter focuses on the two paths for further improving the accuracy and speed of spin flip identification by reducing the variance of the frequency fluctuations σ_0^2 and increasing the spin flip probability in the analysis trap P_a . The next chapter will go through the rest of the measurement sequence, and a proposed improved new trap design is presented in Chapter 6.

4.2 The single spin flip dataset

A key step towards a higher precision measurement was taking a dataset in the analysis trap with sufficient stability that we could distinguish individual spin flips. We will briefly review the data-taking procedure and the data, then the following sections present and then validate the spin flip analysis.

The dataset was taken using a cycle of three axial frequencies similar to the one used for the measurement of Chapter 3 (shown in Figure 3.8). The main difference is that the higher cyclotron stability of this dataset allowed the use of a longer averaging time (32 seconds instead of 24 seconds). Between the f_1 and f_2 we apply a drive near the spin flip frequency, and the differences $f_1 - f_2$ is collected into a "drive" dataset. Between f_2 and f_3 we apply a non-resonant drive, which is used as a control.

The dataset comprises 15 hours of data and 450 full measurement cycles. The control data has a standard deviation σ_0 of 63 mHz, about half a spin flip. As shown



Figure 4.1: (a) Axial stability over several hours in the analysis trap (b) Axial stability as a function of average time [2]

in Figure 4.1, over the course of the 15 hours of data the frequency drifted by many multiples of a spin flip shift. However, almost all the differences between consecutive measurements that are larger than a spin flip only occurred when a drive was resonant.

The frequency is sufficiently stable that we can identify these large jumps as spin flips as they occur, rather than detecting an average spin flip rate over many trials as in the part-per-million measurement. We identify spin flips with a simple analysis: we set thresholds at a frequency differency $\pm \Delta_t$ equal to $\pm \Delta_s$, and declare any measured frequency difference larger than the thresholds to be the result of a successful spin flip. We assess the performance using the efficiency and fidelity, which are defined



Figure 4.2: Histogram of frequency shifts for control data (gray) and driven data (outline), plotted with a calculated control distribution (dashed) and distribution with a drive probability of 50 percent (solid) [36]

and calculated in the next section.

4.2.1 Efficiency and fidelity

Our threshold-based analysis misses some spin flips that did not result in large enough shifts due to random noise. It also mis-identifies some shifts that were not due to a spin flip but rather an especially large fluctuation due to background noise. The likelihood of these unwanted outcomes can be calculated from the expected distribution of frequency shifts due to the drive. We can calculate the efficiency and fidelity using these probabilities.

The fluctuations measured in the absence of a spin flip drive are well described by a random variable Δ selected from a Gaussian distribution, written $G(\Delta, \sigma_0)$. The shaded histogram in Figure 4.2 shows the control fluctuations from the single spin flip dataset with a Gaussian fit overlaid on it. After a spin flip drive, there are four possible outcomes: a $P_a/2$ probability that the spin was up and flips to down, a $P_a/2$



Figure 4.3: Calculated distribution of axial frequency shifts due to underlying distributions with two different σ_0

probability that the spin was down and flips to up, or a $(1-P_a)$ probability that the spin does not flip which is equally divided between particle staying up and staying down. The same random fluctuations are present in each of the four scenarios, so the full distribution is described by the sum of three scaled Gaussians with means offset by $\Delta_s = 130$ mHz:

$$s(\Delta) = \frac{P_a}{2}G(\Delta - \Delta_s, \sigma_0) + \frac{P_a}{2}G(\Delta + \Delta_s, \sigma_0) + (1 - P_a)G(\Delta, \sigma_0)$$
(4.1)

When $\sigma_0 \ge \Delta_s$, the overall distribution still looks Gaussian but with a wider standard deviation. This was the case for the data taken for the part per million measurement where we used the larger standard deviation to measure spin flip probability (Chapter 3). Below that limit, as the peaks get narrow the overall shape of the distribution becomes less Gaussian and the individual peaks start to appear. This is visible in the hollow histogram showing the drive data in Figure 4.3.

In order to evaluate our success rate at inferring spin states from our identified spin flips, we define two quantities: the efficiency E and the fidelity F. The definition is based on the initial spin state rather than the presence of spin flips, as the ultimate goal of the measurement is to infer spin flips in the precision trap based on the state changing between measurements in the analysis trap. The efficiency is defined as the fraction of all trials where the spin state is correctly identified. The fraction of trials where we observe a fluctuation above threshold due to noise and incorrectly assign a spin state is called the inefficiency, I. The fidelity F is the fraction of correctly identified spin states over all spin assignments, which is equal to

$$F = \frac{E}{E+I} \tag{4.2}$$

All three quantities depend on Δ_t and σ_0 .

The total probability of observing an event above the threshold can be calculated by integrating the full probability distribution in Equation 4.1. We can also calculate probabilities for observing a fluctuation above the threshold for each of the four possible drive outcomes individually. The largest probability should come from a spin that has flipped from down to up, called $P_{\downarrow\uparrow}(\Delta_t)$. The next most likely is that the particle particles spin did not flip but a large fluctuation happen to cause an above-threshold shift. We distinguish between the probability of this happening when the spin is up $P_{\uparrow\uparrow}(\Delta_t)$ or down $P_{\downarrow\downarrow}(\Delta_t)$ because in the second case the incorrect determination of a spin flip up still correctly identifies the initial spin state. Finally, it is also possible that an extremely large positive fluctuation causes a mis-identification of the sign of a spin flip from up to down, referred to as $(P_{\uparrow\downarrow})$. These probabilities can be calculated from integrating over the scaled Gaussian distributions:

$$P_{\downarrow\uparrow}(\Delta_t) = P_a \int_{\Delta_t}^{\infty} G(\Delta - \Delta_s, \sigma_0) \,\mathrm{d}\Delta, \qquad (4.3)$$

$$P_{\uparrow\uparrow}(\Delta_t) = P_{\downarrow\downarrow}(\Delta_t) = (1 - P_a) \int_{\Delta_t}^{\infty} G(\Delta, \sigma_0) \,\mathrm{d}\Delta, \qquad (4.4)$$

$$P_{\uparrow\downarrow}(\Delta_t) = P_a \int_{\Delta_t}^{\infty} G(\Delta + \Delta_s, \sigma_0) \,\mathrm{d}\Delta.$$
(4.5)

The efficiency is equal to the probability of correctly inferring a down state from an above threshold event:

$$E = P_{\downarrow\uparrow} \left(\Delta_t \right) + P_{\downarrow\downarrow} \left(\Delta_t \right) \tag{4.6}$$

The inefficiency I is the fraction of events where an initial state up caused an abovethreshold event due to positive random fluctuations:

$$I = P_{\uparrow\uparrow}(\Delta_t) + P_{\uparrow\downarrow}(\Delta_t) \tag{4.7}$$

The fidelity F is the fraction of assigned spin state that have been correctly identified. It can be calculated from E and I as in Equation 4.2.

The control data has a standard deviation of 63 mHz. Setting the threshold equal to the spin flip size at 130 mHz and using this σ_0 , we find that E = .26 and F = .96. The efficiency roughly corresponds to assigning a spin state in one out of four attempts. However, of the assigned spin states, greater than 95 percent are correct. This is sufficient for a double trap measurement, as explored in more detail in Section 4.2.3.

4.2.2 Verification with correlations

To validate that events above the threshold are truly due to spin flips, we conducted an analysis of the correlation between frequency differences. Out of the 450
points in our dataset, roughly 30 of them were identified as two spin flips in a row, in agreement with the calculated $(E/F)^2 \approx 0.07$ probability of identifying two sequential flips. For each pair of above threshold events, the event has a shift Δ_1 , and the following event has a shift Δ_2 . If Δ_1 was truly due to a spin flip, the particle is in the up state at the start of the next event. In that case is not possible for the next event to be a true flip from down to up, which alters the distribution of Δ_2 . If we are correctly identifying most spin flips, it should be much more likely that Δ_2 negative due to a spin flip down than that it is positive due to random noise. As a result, the difference $\Delta_1 - \Delta_2$ should be correlated, with a distribution dominated by two peaks around $\pm 2\Delta_s$. A third peak at zero corresponds to events where one or both of the spins were misidentified, and should be much smaller than the the other peaks if our estimated 96 percent fidelity is accurate.

Our data shows the predicted double peak structure. The left side of Figure 4.4 shows a histogram of these events for our data, with the expected absence of events around $\Delta_1 - \Delta_2 = 0$. The right side shows the results of a simulation using a random number generator, a fifty percent spin flip probability, and an additional frequency shift randomly selected from a normal distribution with standard deviation equal to $\sigma_0 = 63$ mHz. From repeated simulations, we find 30 ± 7 events in the peaks around $\pm \Delta_s$, and 2 ± 2 events in the central peak due to misidentified random fluctuations of the background. This agrees with the 25 events at the peaks and 3 central events of the experimental data. It also agrees with the calculated rate of two correlations in a row with one in the wrong direction: $P_{\uparrow\uparrow} \times (2E + P_{\uparrow\uparrow})$. Given our parameters, we would expect 0.5 percent of all 450 events to lie in the central peak, or an average



Figure 4.4: (a) Experimental data showing distribution of $\Delta_1 - \Delta_2$ when the consecutive differences Δ_1 and Δ_2 are both above the threshold (b) Simulated data based on theoretical distribution for double spin flip events [36]

of 2.3 events. The strong agreement between simulation, calculation and data give confidence that the identified events are true spin flips.

4.2.3 Application to *g*-factor measurement

The demonstrated performance would be adequate to perform a g-factor measurement. With an efficiency of 26 percent and a 96 percent fidelity, we could prepare the initial spin state in less than three minutes. Detecting the final spin state after the precision trap drive would take comparable time, though the low efficiency would complicate the measurement.

The spin state could be prepared in the analysis trap in an average of four attempts. First, we would measure the axial frequency for 32 seconds, apply a saturated spin flip drive, then measure the axial frequency again. If the measured axial frequency shift is below the threshold, we would repeat the procedure until we see a shift above the threshold. Then we would assign a final spin state based on the sign of the measured shift. The average number of attempts needed for a successful spin state determination is 1/(E+I), or about 4 attempts. The fidelity of the final step for this procedure is equal to the fidelity of a single attempt. The total time to achieve a known spin state with 96 percent fidelity would take less than three minutes.

We can use a similar procedure to detect the spin state upon returning to the analysis trap. Some cooling of the magnetron and axial motion would likely be necessary after transfer. This has not yet been optimized so we use current times of a few minutes as a worst case estimate. After that, we could measure the axial frequency, drive the spin transition, and measure the shift in frequency. If the shift is above the threshold, we can identify a spin flip and from its sign we can infer the state of the antiproton when it arrived in the analysis trap. However, if the shift is not above the threshold there is a 33 percent chance that the spin has still flipped even though it was not detected. In that case, a subsequent spin flip detection would likely infer the wrong initial spin state. The safest approach is to restart the entire measurement cycle when the spin state cannot be determined.

This would mean the total number of attempted trials needed for a measurement is inflated by a factor of four. Improving the efficiency of spin flip detection could reduce the data-taking time needed by up to this factor.

4.3 Improving stability

An obvious route to improve the efficiency of spin flip detection would be to reduce σ_0 so that there is less overlap between peaks and we can choose a threshold that accepts all of the true spin flips. Chapter 6 will discuss a proposed new experimental

design that includes several features aimed at reducing σ_0 . This section estimates the effects of changing the threshold Δ_t and reducing the scatter σ_0 on the efficiency, fidelity, and total number of trials needed for a measurement. We will show that for $\sigma_0 \leq 30$ mHz efficiency and fidelity are both essentially 1. The last part of this section presents an alternate approach to increase efficiency by increasing the fraction of spins flipped by the drive.

4.3.1 Effect of fidelity on measurement time

The measurement time is approximately inversely proportional to the efficiency, as we will need to repeat the entire measurement cycle whenever a spin state cannot be determined. To some extent, efficiency can be gained at the expense of fidelity by lowering the threshold, Δ_t so we accept a larger fraction of events. However, a low fidelity will require taking more points to average out the effect of mis-identifications. To find the optimal threshold for a given dataset, we will need to find the exact relationship between fidelity and total measurement time.

We start by defining the precision trap fidelity F_p , as is the probability of correctly identifying the effect of the drive in the precision trap. This fidelity depends on both the fidelity of preparing the spin state F_1 and of correctly detecting the spin state after returning from the precision trap F_2 . In fact, it is exactly equal to the probability of measuring both spin flips correctly, plus the probability of measuring them both incorrectly so that the effect of the error cancels out:

$$F_p = F_1 F_2 + (1 - F_1) (1 - F_2)$$
(4.8)

When the precision drive and cyclotron cooling steps of the measurement are relatively

time-consuming, it may be advantageous to lower the thresholds for spin flip detection during the readout step to enhance E_2 at the cost of reduced F_2 . Since a low E_1 is much less costly in time than a low E_2 , we can almost certainly maximize overall fidelity vs time by setting the state preparation threshold relatively high so that $F_1 \approx 1$. In the limit $F_1 = 1$, we have $F_p \approx F_2$, and we assume we are in that limit for the rest of the calculation.

The effect of $F_p < 1$ is to reduce the contrast between off-resonance and onresonant drives. An obvious reason is that even for non-resonant drives, there will be some background of observed spin flips at a rate $1 - F_p$. However, it will also mean the average probability of detecting a spin flip for a resonant drive \hat{P}_{max} differs from the true probability of flipping the spin P_{max} . They are related by

$$\hat{P}_{max} = P_{max}F_p + (1 - F_p)(1 - P_{max})$$
(4.9)

The difference in probability between the peak and the background becomes

$$\hat{P}_{max} - (1 - F_p) = 2P_{max} (F_p - .5)$$
 (4.10)

This is maximized when $F_p = 1$ and is zero when $F_p = .5$ as expected.

The reduced contrast from imperfect fidelity forces us to take extra data points in order to reduce the statistical error on the measurement of the spin flip rate. In practice we will determine the number of trials necessary by choosing an appropriate confidence interval estimator. However, since we are most interested in a case where $P_{max} \approx 0.5$ a normal distribution is a reasonable approximation as long as we have enough trials [56]. We use the mathematically simple requirement that the contrast between peak and background is equal to at least five standard deviations. The



Figure 4.5: For $\sigma_0 = .01$ (blue), $\sigma_0 = .03$ (red), $\sigma_0 = .06$ (black), and $\sigma_0 = .13$ (green), (a) Efficiency (solid) and fidelity (dashed) as a function of threshold Δ_t (b) Attempted trials required N_a as a function of threshold Δ_t

average proportion of spin flips measured over N trials will randomly vary according to a binomial distribution, with a standard deviation given by

$$\sigma = \sqrt{\frac{\hat{P}_{max}\left(1 - \hat{P}_{max}\right)}{N}} \tag{4.11}$$

By setting the contrast equal to 5σ , we can estimate the number of trials needed as a function of \hat{P}_{max} and F:

$$N = \frac{25\hat{P}_{max}\left(1 - \hat{P}_{max}\right)}{4\hat{P}_{max}^2\left(F_p - 0.5\right)^2}$$
(4.12)

As expected, the number of trials needed increases as the fidelity decreases.

4.3.2 Estimating number of trials as a function of E, F and σ_0

In total, we will need N completed trials to get the desired statistical precision, and we will need on average $N_a = F/E$ attempted trials for every completed trial. We use these relationships to find the optimal choice of threshold for $\sigma_0 < \Delta_s$. Then using this optimized threshold we can compute the estimated number of trials needed as a function of σ_0 . This will show that reducing σ_0 by a factor 2 is sufficient to achieve the desired factor of four improvement in measurement time. Further reductions in stability will only marginally improve the measurement time.

First, we further simplify Equation 4.12 by assuming that we use a saturated drive with $P_{max} = 0.5$. A saturated drive would minimize the number of points needed by maximizing the contrast between resonant and non-resonant frequencies. In this case, $\hat{P}_{max} = P_{max}$ since on average there are equal numbers of misidentified flips and non flips. The number of attempted trials needed for a measurement N_a becomes

$$N_a = \frac{25}{4\left(F_p - .5\right)^2} \frac{F_p}{E_p} \tag{4.13}$$

Since the efficiency and fidelity are proportional to error functions, it is not generally possible to analytically calculate the optimal thresholds to minimize N_a . Instead, we numerically compute the efficiency, fidelity, and N_a for a number of interesting combinations of σ_0 and Δ_t given $P_a = 0.5$. The results are shown in Figure 4.5. We find that while a poorly chosen threshold can require many additional trials, N_a is relatively insensitive to the choice of threshold over a wide range roughly centered on $\Delta_t = \Delta_s/2$ for the values of σ_0 studied.

In that case, we can calculate the relationship between σ_0 and N_a by fixing the threshold $\Delta_t = \Delta_s/2$. The result, shown in Figure 4.6, fits well to a quadratic curve for larger values of σ_0 but flattens out below a critical stability $\sigma_0 \approx 30$ mHz. This indicates improving stability will reduce the measurement time until this threshold is reached.

Likely we can reduce the number of attempted trials even further at that stability by changing our procedure. The three peaks of the distribution would all be distinct,



Figure 4.6: Number of trials needed to complete a measurement as a function of scatter. Blue dots are numerically calculated points, dashed line is a quadratic polynomial fit



Figure 4.7: Distribution of axial frequency shifts from $P_a = 0.5$ and $\sigma_0 = .03$ Hz. Vertical lines show thresholds $\Delta_t = \Delta_s/2$

as shown in Figure 4.7. In that case, if $\Delta < \Delta_t$ there is a 98.4 percent probability that the spin did not flip. We could attempt to flip the spin multiple times while remaining confident that the first above-threshold observed will be the first spin flip after returning from the precision trap. A simulated procedure that keeps flipping the spin until we detect a shift over threshold and then infers the initial spin state had an overall fidelity of 97 percent for the initial spin state. With this modified procedure



Figure 4.8: Distribution of axial frequency shift with 63 mHz background scatter and 100 percent spin flip probability.

we could identify the spin state in almost 100 percent of trials, further reducing the total number of trials needed to the minimum possible of 25.

4.3.3 Adiabatic Fast Passage

The other possible approach to improve the efficiency and fidelity of single spin flip detection is to increase the spin flip probability in the analysis trap [2]. If we could increase P_a from 0.5 to 1, we could eliminate the central peak due to spins that don't transfer in Figure 4.3. The separation between the remaining peaks at $\pm \Delta_s$ would be effectively doubled. We could set a single threshold at $\Delta_t = 0$ and assign 100 percent of measurements as transitions up or down with a fidelity of 98 percent. This is illustrated in Figure 4.8.

Achieving adiabatic fast passage might also allow us to robustly distinguish spin flips even in the case of scatter higher than 63 mHz. This might allow us to spend less time averaging the axial signal or cooling the cyclotron motion to reduce scatter. Figure 4.9 show the fidelity as a function of σ_0 with the threshold fixed at $\Delta_t = 0$ so



Figure 4.9: Calculated fidelity in the limit that $P_a \rightarrow 1$, assuming $\Delta_t = 0$

that the efficiency is always 100 percent. As in the last section drive there is a critical stability above which the fidelity starts to decline. It happens that is very close to the 63 mHz stability achieved in the single spin flip dataset.

We could experimentally realize $P_a \approx 1$ with a coherent drive technique known as adiabatic fast passage (AFP). This technique was originally developed in nuclear magnetic resonance to completely invert the state of a population of spins [57]. By adiabatically sweeping the frequency of the spin flip drive, the probability of a transition accumulates until the entire population is transferred.

It is possible that the strong magnetic gradients in the analysis trap will prevent us maintaining the coherence necessary for adiabatic fast passage. In a typical AFP experiment, the drive mixes the spin states and effectively tilts the magnetic field seen by the particle in its rotating frame through a 180 degree rotation. However, if the drive is swept too quickly, this process is not adiabatic and the particle state does not follow the drive. The "adiabatic condition" for a sweep slow enough that the particle's spin can follow it is

$$\left|\frac{d\omega}{dt}\right| << \Omega_R^2 \tag{4.14}$$

where Ω_R is the Rabi frequency. However, in our case, the 'static' magnetic field seen by the particle is also changing in time due to the fluctuating axial amplitude of the particle in the magnetic gradient of the analysis trap. The average magnetic field seen by the particle is

$$\langle B \rangle = B_0 + B_2 \left\langle z^2 \right\rangle = \frac{k_B T}{m \omega_z^2} \tag{4.15}$$

The random fluctuations in $\langle z^2 \rangle$ are exponentially damped with a time constant γ_z , so that field is changing at approximately $B_2 \langle z^2 \rangle \gamma$. Then the adiabatic condition becomes:

$$\left\langle \frac{d\omega}{dt} \right\rangle = \left\langle \frac{gq}{m} \frac{dB_{eff}}{dt} \right\rangle = \gamma_z \frac{gqB_2}{m} \left\langle z^2 \right\rangle$$
(4.16)

We can use the definition of the linewidth from Equation 2.10 and write the condition as simply

$$\Omega_R^2 >> \gamma_z \Delta \omega \tag{4.17}$$

The Rabi frequency of the drive can only be increased so much before the dissipated drive power causes systematic shifts in the frequency. We tested the AFP technique briefly after completing the measurement of Chapter 3, but at that time the antiproton measurement the coupling constant γ was 2.2 Hz, the linewidth was 24 kHz, and the Rabi frequency was only 410 rad/s which does not satisfy the inequality.

However, since then we have implemented a series of improvements to the drive frequency. We should now be able to safely apply a 10 dB stronger drive, and further improvements have been made to the next generation [58].

We are still likely at the edge of feasibility for this technique, and if we find we cannot apply sufficient drive power we could instead try to reduce γ_z . One option

would be to deliberately detune the particle frequency from the axial amplifier by adding a bias to the endcaps. The second would be to use axial feedback cooling (described in Section 2.6) which would reduce $\langle z^2 \rangle$ in addition to γ_z .

4.4 Conclusion

We have detected individual spin flips, a necessary precondition for a double trap measurement at the ppb level. We have studied the best routes to improve the performance of spin flip detection. The number of attempted spin flip trials needed at each point on the spin flip line depends on both the efficiency and fidelity. Those depend on the threshold Δ_t chosen, as well as the underlying distribution determined by σ_0 and P_a . Numerical calculation indicates that $\Delta_t = \pm \Delta_s/2$ is close to the optimal threshold for a wide range of distributions when $P_a = 0.5$. The total number of trials needed can be reduced by a factor of four in two ways: Improving the stability until $\sigma_0 \leq 30$ mHz or increasing the transition rate P_a to close to 1 using AFP and maintaining $\sigma_0 \leq 63$ mHz. Reducing the stability below the identified critical level does not further reduce the number of trials.

Chapter 5

Separated Oscillatory Fields Measurement of the *g*-factor

Obtaining the narrowest possible linewidth is a critical goal for making the most accurate measurement. The linewidth in the precision trap was originally estimated to be 2.7 ppb due to the estimate 8 T/m² magnetic gradients from the trap spacers and electrodes. A measurement of the proton g-factor by the BASE collaboration that obtained 1.5 ppb precision with a full width half max (FWHM) of 32 ppb was ultimately limited by random drifts of the magnetic field between the interlaced spin and cyclotron measurements [13]. A more recent measurement simultaneously measured the spin and cyclotron frequencies, but using a cyclotron-axial coupling that broadens the cyclotron linewidth to 3 ppb [12]. To achieve simultaneous spin flip measurements with the narrowest possible resonances, we propose a measurement based on a separated oscillatory fields technique¹.

 $^{^1 \}rm We$ thank Professor Ed Myers of Florida State University for useful discussions of using separated oscillatory fields measurements

We outline the full measurement procedure and predict that a resonance with a line-width narrower than 0.1 ppb can be measured with using less than half an hour per data point, leading to a measurement in as little as 100 hours (based on the estimate of the number of data point from Chapter 4).

5.1 Historical overview

The method of separated oscillatory fields has a long history in ion experiments. In essence, SOF measurements drive a resonance with multiple phase-coherent pulses separated by a long free oscillation time, T. The resulting interference pattern can be narrowed by extending T. This general technique has been extended to a wide variety of contexts [59, 60, 61, 62], but we focus here on historical applications directly related to measurement of spin and cyclotron transitions.

The SOF technique was originally developed by Norman Ramsey to measure the nuclear magnetic moments of molecules in a beam [63]. Previous experiments were limited by experimental difficulties with long RF pulses within a homogeneous field, which determine the width of the response to a single coherent response. Instead, Ramsey split the RF drive into two short pulses at the entrance and exit of a region of constant field. The total probability of transition is dominated by an interference pattern between the phase accumulated by the particle's rotating spin in the constant field region, $\int_{t_0}^{t_0+T} \omega(t) dt$, and the phase accumulated by the drive during the same period, $\omega_d T$. The width of the central peak of the interference pattern is determined by the free oscillation time.

A variation on this technique was used to compare the cyclotron frequencies of

ions with differing masses in a Penning trap [64]. Both ions were driven with an initial drive pulse, allowed to freely oscillate in the same trap, then driven with a second identical pulse. The final cyclotron amplitude for each particle depended on the phase difference accumulated between the cyclotron drives and the particle during the free oscillation period. The trap voltage was adjusted to bring first one ion then the other into resonance with the axial amplifier so that the cyclotron amplitude could be read out using a cyclotron-axial coupling technique. In this experiment, the SOF technique allows the ions to be driven simultaneously even though the results needed to be measured one ion at a time [65]. In our experiment, this allows us to separate driving the cyclotron state from reading it out from the resulting amplitude. We can further take advantage of our existing analysis trap to directly measure the cyclotron amplitude from the axial frequency shift by transferring particles to that trap after application of the pulses.

Since both the nuclear magnetic resonance frequency and the cyclotron frequency must be measured for a g-factor measurement, we have the possibility of simultaneously applying pulses for both measurements. Then the interference patterns in the resulting cyclotron amplitude and spin transition probability are determined by the average magnetic field over the exact same oscillation time. This avoids potential experimental errors due to the magnetic field drifting between the spin and cyclotron measurements.

5.2 Derivation of spin and cyclotron resonances

The key advantage of SOF for our experiment is that we can use the exact same pulse time τ and free oscillation time T to drive spin and cyclotron resonances of similar linewidths, so that we are measuring both frequencies in the exact same magnetic field. We will show that the drive parameters can be chosen such that both resonances are proportional to $\cos^2 \Delta T$, where Δ is the average detuning of the drive from the resonance over the free oscillation period. The resulting fringe patterns and lineshapes can be narrowed by extending T.

5.2.1 Spin resonance

The proton spin transition is a special case of the nuclear magnetic moment resonances calculated by Ramsey. The probability of a spin transition after both pulses is [66]

$$P_{sf} = \frac{\Omega^2}{\mu^2} \sin^2(\mu \tau/2) \left[\cos(\Delta_s T/2) \cos(\mu \tau/2) - \frac{\Delta_s}{\mu} \sin(\Delta_s T/2) \sin(\mu \tau/2) \right]^2 \quad (5.1)$$

where $\mu = \sqrt{\Delta_s^2 + \Omega^2}$, and We drive the spin transition in our precision trap with an oscillating magnetic field at ω_d with average detuning Δ_s from the spin frequency $\omega_s(t)$. The drive strength is characterized by the Rabi frequency Ω , which can be increased to as much as 60 Hz as described in Section 2.9.

This expression can be greatly simplified if we assume that we are only measuring over a small frequency range close to the resonant frequency (i.e. $\Delta_s \tau \ll 1$) with a strong drive. In this case, we can ignore the second term because $\Delta_s / \Omega \approx 0$, yielding

$$P_{sf} \approx \sin^2\left(\Omega\tau\right) \cos^2\left(\Delta_s T/2\right) \tag{5.2}$$



Figure 5.1: Blue line: Calculated fringe for the spin transition probability with $\Omega = 10$ Hz, $\tau = \pi/(2\Omega)$ and T =1 s. Yellow line: Approximate formula $P = \cos^2 \Delta T/2$

For a resonant drive ($\Delta = 0$), the maximum transition probability is achieved when both pulses are $\pi/2$ -pulses with duration $\tau = \pi/(2\Omega)$. If we assume we chose our pulse length to be *pi*-pulses, we get

$$P_{sf} \approx \cos^2\left(\Delta_s T/2\right) \tag{5.3}$$

This approximate formula describes the shape of the response close to the central fringe, as shown in Figure 5.1.

5.2.2 Cyclotron resonance

For the cyclotron motion, a similar fringe pattern arises from the classical equations of motion [65]. A general calculation starts from the equation of motion for the radial motion with the magnetron motion neglected and the cyclotron position expressed as a complex amplitude $\rho_c = x + iy^2$. The resulting general equation of

² Here the x axis is perpendicular to the split in the electrode used to apply a drive f(t)

motion is:

$$\frac{d\rho_c}{dt} + i\omega_c\rho_c = f\left(t\right) \tag{5.4}$$

where f(t) is a general drive, in our case a series of two pulses of length τ , frequency ω_d and amplitude V separated by an evolution time T. We can find the frequency response of the particle from the Fourier transform:

$$\tilde{\rho}(\omega) = \frac{\hat{f}(\omega)}{i(\omega + \omega_c)}$$
(5.5)

We can invert the Fourier transform to solve for the radial position at the end of the pulse:

$$\rho_c(\tau) = V \int_0^\tau dt' \sin \omega_d \left(\tau + t'\right) \int_{-\infty}^\infty dw \frac{1}{(\omega + \omega_c)} e^{-i\omega t'} = V \int_0^\tau dt' \sin \omega_d \left(\tau + t'\right) e^{i\omega_c t'}$$
(5.6)

After approximating to zero the effect of fast-rotating components at $\omega_c + \omega_d$, we get

$$\rho_c(\tau) = V e^{i\omega_d \tau} \frac{e^{i(\omega_c - \omega_d)\tau}}{\omega_c - \omega_d}$$
(5.7)

The full solution for the particle's motion includes the transient solution which accounts for the initial position, which we assume are minimal for now and explore further in Section 5.4. We also can assume $(\omega_d - \omega_c) \tau \ll 1$ since we are choosing a detuning on the central fringe and an evolution time much longer than the pulse time. In this limit, the response to the pulse approaches

$$\rho_c\left(\tau\right) = \frac{A}{2} e^{i\omega_d \tau} \tag{5.8}$$

where we have defined the final amplitude $A = 2V\tau$.

Over the period of free oscillation, the particle acquires an additional phase $\overline{\omega_c}T = \int_0^T \omega_c(t)$ and advances to $\rho_c = A e^{i\omega_d \tau} e^{i\omega_c T}$. The second pulse has effect identical to

the first pulse, except the drive has also acquired an extra phase $\omega_d T$. The final position of the particle is the vector sum of its initial position at the end of the evolution time and the change due to the pulse:

$$\rho_c \left(T + 2\tau\right) = \frac{A}{2} e^{i\omega_d \tau} \left(e^{i\omega_d T} + e^{i\omega_c T}\right)$$
(5.9)

The final energy of the particle is then proportional to the radius squared, or

$$|\rho_c (T+2\tau)|^2 = A^2 \cos^2 \left[(\omega_c - \omega_d) T/2 \right]$$
(5.10)

which is a very similar pattern with a dependence on T to the spin case scaled by the respective frequencies.

5.3 g-factor lineshape in a drifting magnetic field

While drifts over each individual measurement cycle are accounted for with this technique, they can still impact the overall measurement. We will need multiple measurements at each spin flip frequency in order to compute an average spin flip probability. Drift in the magnetic field in between these measurements complicate this process and must be accounted for.

Previous experiments provide some guidance about plausible targets for controlling the drift of magnetic fields. The electron experiment in an identical solenoid at Harvard University saw drifts rates on the order of a few tenths of a ppb per hour [67]. However, That may be an unrealistic goal in the less well-controlled particle accelerator environment at CERN. The proton/antiproton charge-to-mass ratio comparison in our exact magnet at the LEAR facility saw average drifts of 1 ppb/hour [68]. We use this as a worst-case estimate for the rest of this chapter. In that case, the 25 or so points we would need at each frequency (see Chapter 4) could be measured in fields that differ by several ppb.

To compensate for drift, we must completely determine the cyclotron frequency for each spin flip attempt and use it to adjust the spin drive frequency f_{sd} . We can invert Equation 5.10 in order to find the antiproton's cyclotron frequency from our measurement of the final cyclotron radius. Using it to scale f_{sd} gives us a resonance that depends on the q factor:

$$P_{sf} = \cos^2 \left[\pi \frac{\delta g}{g} g f_c T \right] \tag{5.11}$$

where $\delta g = g - f_{sd}/f_c$.

Additionally, the periodic nature of the response means we can only determine the frequencies up to an integer multiple of π . We resolve this by measuring in several steps of increasing evolution time. Each measurement reduces the error on the spin frequency sufficiently that it is known to be somewhere on the central fringe of the subsequent measurement. This will increase the total measurement time necessary, but only by a factor of three. We can measure the *g*-factor down to 1 ppb by using a single pulse for each frequency in place of the SOF sequence. Then the resonance is Lorentzian rather than periodic and we can uniquely determine the frequency by simply finding the peak. However, around 1 ppb the pulse would would need to be on the same order as the axial damping time, which means the resonance is broadened by the residual gradient in the precision trap. At that point it is advantageous to switch to the full SOF method as then the resonance will only depend on the average field during the long oscillation time.



Figure 5.2: G-factor fringe for T=1 second in blue, T=3 seconds in yellow, and T=9 seconds in green.

We start with an free oscillation time of 1 second, resulting in a half-width of 1 ppb. If we aim for a point halfway down the central fringe, the true detuning could be anywhere within our starting uncertainty of \pm 1 ppb and still be between the maximum at $\Delta_c = 0$ and the minimum at $\Delta_c = \pi/4$. Each time we measure the line, we determine the cyclotron frequency well enough that we can increase T by a factor of 3 while the fringe half-width (distance from a maximum to a minimum) is still greater than the uncertainty. We show that we can measure the cyclotron frequency with sufficient precision to do this in the next section. Figure 5.2 shows three g-factor resonances with evolution times increased by this factor, going from a half-width of 1 ppb to 0.1 ppb.

5.4 Experimental errors affecting the cyclotron measurement

We must demonstrate that we can measure the cyclotron amplitude such that a single measurement cycle determines the cyclotron frequency with our desired subppb precision. The cyclotron amplitude is measured by transferring the particle to the analysis trap and measuring the shift in the axial frequecy from the excited cyclotron radius in the large magnetic gradient there (see Section 2.8). The greatest contribution to the uncertainty in almost all circumstances is the random initial phase of the particle's motion. The second source of error that must be considered are drifts in the cyclotron radius during the measurement of the cyclotron state. We also consider error due to drifts in the precision trap and a few other effects, and show that all are at least an order of magnitude smaller. Errors of less than 100 ppt seem achievable.

5.4.1 Error from the unknown initial state

The particle's final position as given by Equation 5.9 is modified by the assumption that the particle has a nonzero initial position r_0 and phase ϕ_0 . This radius and phase are randomly selected from a thermal distribution when the particle is decoupled from a thermal reservoir by leaving the cooling trap, and are unknown at the start of the pulse sequence. The full expression for the position ρ_c is

$$\rho_c \left(T + 2\tau \right) = r_0 e^{i\omega_c T} e^{i\phi_0} + A e^{i\omega_d \tau} e^{i(\omega_c + \omega_d)T/2} \cos\left(\Delta T/2\right)$$
(5.12)

and the energy is proportional to

$$|\rho_c (T+2\tau)|^2 = r_0^2 + r_0 A \cos(\Delta T/2) \cos(\phi_0 - \omega_d \tau) + A^2 \cos^2(\Delta T/2)$$
(5.13)

The first term is a constant offset that can generally be neglected if $A > r_0$. However, the second term will vary randomly with ϕ_0 . The distribution of resulting errors in $n = \frac{m\omega_c}{2\hbar} |\rho_c (T + 2\tau)|^2$ is highly non-Gaussian, but there is a maximum and a minimum possible value of $\pm r_0 \cos(\Delta T/2)$. We convert this to a range of quantum numbers for easy comparison to the measured $f_z = n * 0.05$ mHz:

$$dn = \sqrt{n_0 n_A} \cos\left(\Delta T/2\right) \tag{5.14}$$

where n_0 is the initial quantum number of the particle and n_A is the quantum number of a particle excited to maximum $\rho_f^2 = A^2$. Figure 5.3 shows how the fringe will vary with five different choices of $\delta \phi$ if we increase the energy of the particle by an average factor of (a) 50 and (b) 800. From these, we can see that stronger pulses that drive the particle to larger amplitudes result in a lower fractional error on the measurement of the quantum number.

Even a relatively large dn can correspond to a small range of cyclotron frequencies df_c due to the steep slope of the fringe and our knowledge of the frequency to within a cycle from prior measurements. To propagate the error, we first solve Equation 5.10 for the detuning Δ_c and substitute in the quantum number n:

$$f_c - f_{cd} = \frac{1}{\pi T} \operatorname{Arccos}\left[\sqrt{\frac{n}{n_A}}\right]$$
(5.15)

where we are assuming we that we are driving on the positive side of the central fringe and thus $0 < f_c - f_{cd} < \frac{\pi}{T}$. We assume we can measure T and n_A with high enough precision that the dominant error comes from dn^{3} . Then we can propagate the error dn to an error in the measured cyclotron frequency df_{c} with a linear approximation:

$$df_c = dn \times \frac{df_c}{dn} = \frac{1}{\pi T} \frac{dn}{n_A} = \frac{1}{\pi T} \frac{1}{2\sqrt{n}\sqrt{n_a - n}} dn$$
(5.16)

If we assume we measure around $n = n_A/2$, this is simply

$$\frac{1}{2\pi T}\sqrt{\frac{n_0}{n_A}}\tag{5.17}$$

Narrowing the fringe by increasing T directly reduces the error on df, as does increasing the final amplitude n_A . For example, if we increase the particle's energy by a factor of 50 using an evolution time of 10 seconds, we have a 14 percent error in measuring the quantum number dn. That is sufficient to constrain the the cyclotron frequency to 50 ppt.

5.4.2 Error due to cyclotron excitation drifts in the analysis trap

The cyclotron state is drifting during the time spent searching for a response in the analysis trap. We estimate the time needed in the analysis trap and the resulting drift, and show it does not contribute significantly to dn in most circumstances.

We first estimate the time needed in the analysis trap and show that it is proportional to n_A . The fastest way to measure the axial frequency is to appy an axial heating drive, which will cause a distinctive peak (see Section 2.7) as long as the true frequency is within the bandwidth of the amplifier. If not, we can step the trap voltage to shift the particle's frequency by 200 Hz and try again. We estimate that

 $^{^{3}}$ if in practice these errors cannot be neglected they should be added in quadrature

the long time constant filters on the trap voltage plus the necessary time to average the axial signal will take about 45 seconds for each step. The final quantum number lies between $n_A/2 - n_A/4$ and $n_A/2 + n_A/4$, and we need to cover the resulting range of frequencies, stopping when we see a response. Thus, on average we need to cover a range $\frac{n_A}{4} \times 0.05$ Hz. The average time needed to see a response is $T_{an} = \frac{n_A}{4} \times \frac{.05}{200} \times 45$ seconds.

The resulting drift is much smaller than the error due to the initial position in almost all cases. We make our estimate using the analysis trap transition rate reported in Chapter 3 as a worst case scenario, though in the following chapter we present a redesigned analysis trap that should have an order of magnitude slower transition rate. We measured a transition rate per quanta of $\Gamma = 0.000021$ transitions/sec/quanta, which we substitute into Equation 3.2.2 to get the drift error $dn_d = \sqrt{n * T_{an} * .000021}$. We can substitute in the averaging time need T_{an} to find the analysis trap drift error:

$$dn_d = n_A \times .00024 \tag{5.18}$$

This is orders of magnitude smaller than the initial phase error as long as $n_A/n_0 << 8.7 * 10^6$. We can assume we are in this limit because otherwise we would have an extremely large final amplitude that would take an impractically long time to measure in the analysis trap.

5.4.3 Other sources of error

The particle's state will also drift due to noise during its time in the precision trap, but generally not by much as the minimum time needed in the precision trap is smaller.



Figure 5.3: Calculated fringes for $\delta \phi = 0$, $\frac{\pi}{4}$, $\frac{\pi}{2}$, $\frac{3\pi}{4}$, and π for a pulse strength that increases the particle energy by a factor of (a) Five (b) Twenty. The horizontal line indicates the half-max height, where without prior knowledge of $\delta \phi$ such an amplitude could have been cause by a range of frequencies.

The minimum time in the precision trap is determined by our need to measure the axial frequency before and after each SOF measurement, which takes 60 seconds each time. However, the cyclotron state transition rates may be higher in the precision trap. All the analysis trap lines are heavily filtered at the cyclotron frequency, but we cannot filter the precision trap cyclotron drive line as heavily because we need to be able to apply a strong coherent cyclotron drive. With no filtering at all, noise from the 50 Ω drive cable at the hat would cause 0.2 transitions per quanta per second in the precision trap. An n = 200000 particle waiting for 60s in the precision trap would drift by 1500 quanta, or 0.8 percent. In practice, we attenuate the cyclotron drive line in the precision trap by roughly -30 dB so the actual average drift should be a factor of 30 smaller and can be ignored.

There are several more possible sources of error due to voltage fluctuations from transferring or the distribution of final states after magnetron cooling, but since they are all are on the order of tens of quanta at most we neglect them here. They are discussed in more detail in [58].

5.5 Procedure for compensating field drift

A sub-ppb measurement of the cyclotron frequency must use repeated measurements of higher precision to track the central fringe in the precision of magnetic field drift. Even under optimistic assumptions (see Section 5.7) we will require roughly 30 minutes for each measurement which translates to an average drift between points of 0.5 ppb. To compensate for this, we perform an initial measurement with a lineshape broadened to several ppb so that the central fringe covers the entire initial uncertainty in the cyclotron frequency due to drift since the last measurement. With a strong enough drive, we can determine the cyclotron frequency with greater precision than our desired final half-width of 0.1 ppb. We can then return to the precision trap and apply simultaneous spin and cyclotron drives with the evolution time and drive detuning adjusted to produce a spin response on a fringe of our desired width. The cyclotron drive strength can be lowered for this second measurement to reduce the readout time, as the narrower fringe will still achieve high precision.

We illustrate the procedure with some example evolution times. An evolution time of 1 second produces a cyclotron fringe with a half-width of 3 ppb to cover the full possible range of drifts. We can measure the cyclotron frequency to 6 mHz, or roughly 70 ppt With a drive strength chosen to increase the final energy by a factor of 1300, . If the particle started at n=300, it would end somewhere between n=100,000 and n=300,000. Searching the resulting 20 kHz range would take more than twenty minutes, by which point the field has drifted again.

This motivates selecting a colder initial particle so that less time is needed to find the initial response. If we start with n=50, we can measure the initial frequency

to 70 ppt in less than four minutes. A 10 second evolution time would suffice to produce a final fringe with a half-width of 0.1 ppb. The error on the cyclotron frequency df_c is highly constrained by the long evolution time, allowing a weaker drive. The main consideration is that we must keep the amplitude error dn less than the amplitude range for a meaningful measurement, which can be accomplished with a 10-fold increase. From this point, we will take on average another 30 minutes finding the final analysis trap response, but as the measurement has been already locked in further magnetic field drifts are not as significant.

5.6 Demands on experimental performance

As shown in the last section, it is essential to be able to quickly select ultracold $(n \leq 50)$ particles and rapidly cool the particle after cyclotron excitation. This is unrealistic due to the long cyclotron damping time in the precision trap, which motivated a trap designed for fast cooling times (presented in Chapter 6).

Selecting a particle with $n \leq 50$ is extremely time consuming. The procedure begins with the antiproton in the cooling trap, where it damps radial energy into a resistance provided by a tuned circuit resonant at the cyclotron frequency (as described in Section 2.4). The electronic temperature of the tuned circuit is set by the electronic noise level which is generally higher than its physical temperature. This is extremely time-consuming to measure for the cyclotron amplifier but our estimate is based on the 8 K measured for our axial amplifiers. The time constant of the damping τ_c was measured to be 200 s on the old trap. The particle must damp for six time constants to reach 8 K from its initial amplitude Then, we transfer the antiproton to the analysis trap where it is isolated from the tuned circuit. This "locks in" a single oscillation amplitude from the Boltzmann distribution. We can measure n by detecting the shift in the analysis axial frequency $\Delta f_z = n \times 50$ mHz. The fastest way to do this is to sweep a weak axial drive over a 15 Hz range and see if the particle responds. If not, the resonant frequency is too high because n > 300, so we transfer the particle back to the cooling trap and repeat until we see a response in our target range. On average we will need 40 attempts to select an n=300 particle from an 8 K distribution, resulting in a total of 46 damping times or nearly two hours to reselect a cold particle before each measurement attempt.

To speed things up, we propose to reduce the cooling time by a factor of 72. This reduces the time needed at three steps. First, it reduces the time needed to select an ultracold particle by a factor of six. We do not realize the full factor of 72 because the selection procedure is now limited by the 30 seconds necessary to measure the cyclotron amplitude in the analysis trap. Second, as presented in Chapter 4 ultracold particles will increase the efficiency of spin flip detection, reducing the predicted measurement time by a further factor of four. Finally, we can recool the particle with a brief trip to the cooling trap after the initial cyclotron measurement, narrowing the expected final range and thus search time by a factor of 30. The resulting measurement time is presented in the following section.

5.7 Predicted measurement time

Finally, we consider the time needed for the full separated oscillatory fields measurement with a lineshape narrower than 100 ppt. We list here the seven expected steps of the measurement, as well as a predicted measurement time. We start the procedure with a particle with n = 50.

- 1. Measure the spin state in the analysis trap. This is estimated to take only thirty seconds with the ultracold particle as calculated in Chapter 4.
- 2. Transfer the antiproton to the precision trap and probe the cyclotron frequency in the precision trap with T=1 second (fringe half-width of 3 ppb
- 3. Send the particle to the analysis trap and find the axial frequency. This will take at a maximum 7 minutes and on average three minutes.
- 4. Send the particle to the cooling trap and allow it to cool for six damping times (30 seconds) until $n \approx 2000$.
- 5. Go back to the precision trap and perform the measurement sequence: Measure the axial frequency, drive both the spin and cyclotron frequencies with T = 10seconds, and measure the axial frequency again. The total time is two minutes, mostly due to measuring the axial frequency.
- 6. Transfer the antiproton to the analysis trap and find the axial frequency response. We must cover a range of 500 Hz, which can be covered in either 45 or 90 seconds.
- 7. Return to the cooling trap and allow the particle to damp back to equilibrium. Repeatedly transfer the antiproton between the analysis and cooling trap to select an n=50 particle. The total time needed is 21 minutes.
- 8. Measure the spin state in the analysis trap, which will take roughly one minute.

The total time predicted per point is thus estimated to be 29 minutes, of which about 21 minutes is spent cooling the particle, to achieve a final linewidth of 0.1 ppb.

Chapter 6

Trap design

Trap designs are presented which reduce the cyclotron cooling time to allow the selection of sub-thermal particles that are colder by a factor of six than previously used. This reduction will allow the use of higher-precision techniques presented in the last section as well as a faster measurement time. The new traps have been constructed. The construction of the new trap electrodes and the initial results from commissioning will be reported in Chapters 7 and 8.

6.1 Overview

The new trap "stack" has four traps that each have a dedicated role in the final measurement: the precision trap, the analysis trap, the cooling trap, and the loading trap. A section view of the new trap stack with all four traps is shown in Figure 6.1. Table 6.1 shows the heights of crucial electrodes.

The precision trap will provide extremely well-controlled electric and magnetic



Figure 6.1: Picture of the new trap stack

fields, allowing for very narrow spin and cyclotron resonances. In the previous trap, the residual gradient of 8 T/m² from the spacers would limit the linewidth to a few ppb. We replaced the spacers with low-susceptibility quartz and took advantage of a new design that uses symmetry to cancel out any remaining contribution from the spacers, presented in detail in [58]. The new precision trap gradient should be much less than the residual gradient from the analysis trap of 0.8 T/m². We predict the

Electrode	Length(mm)
ARING+ACOMPS	2.93
ABEC	1.5
Cooling Inner EC	2.5
CRING	2.50
Cooling Outer EC	2.75
Trap center separation	6.71
SB electrode separation	4.21

Table 6.1: Heights of the new cooling and analysis trap electrodes

linewidth from the unwanted bottle parameter $\Delta \omega$ will be less than 0.2 ppb.

The analysis trap allows measurement of the spin state by coupling it to the axial frequency though a strong magnetic gradient (see Section 2.8). The analysis trap differs from our earlier model in that it reduces the electronic noise reaching the center of the trap (see Section 6.3).

A dedicated cooling trap, just below the analysis trap, is designed to speed up cooling of the cyclotron motion by a factor of 72. The design of the new trap is discussed in Section 6.2.

A large loading trap at the bottom of the stack is dedicated to loading and storing antiprotons. The larger radius is for more efficient loading than realized earlier. We can also use this trap to store a reservoir of antiprotons for use during periods when the AD is not providing antiprotons. The idea and design considerations of this trap are presented together with its early performance in Section 8.6.

6.2 Cooling trap

The new cooling trap is optimized to minimize the cyclotron damping time. As discussed in Section 5.6, the selection of cold particles is performed by repeatedly bring the particle into contact with a tuned circuit, allowing them to equilibrate, then moving the particle to an isolated trap and measuring its cyclotron energy. The cyclotron energy in the new trap is randomly selected from a thermal distribution determined by the tuned circuit (typically 8 K, with average quantum number n =2000). Reducing the time constant of the coupling between the tuned circuit and the particle τ_c reduces the time necessary for this procedure.

We first review the factors that determine τ_c and the adjustments we have made in the new trap. The largest contribution is a reduced geometrical constant, which is calculated in detail in Section 6.2.2. Optimizing the design for fast cyclotron cooling will not introduce new difficulties due to frequency shifts, as calculated in Section 6.2.3.

6.2.1 Overview of damping time reductions

Dedicating the cooling trap to damping the cyclotron motion allows us to minimize the damping time constant τ_c . The factors affecting the damping time can be derived following similar logic as the derivation of the axial damping time from Section 2.4:

$$\tau_c = \frac{1}{2} \left(\frac{2\rho_0}{e\kappa}\right)^2 \frac{m}{R} \tag{6.1}$$

where m is the mass of the particle. A dedicated cooling trap allows us to maximize the resistance R, reduce the trap diameter ρ_0 , and maximize the geometrical factor к.

First, we made two changes that increased R and reduced the damping time by a factor of two. The resistance R at the 100 MHz cyclotron frequency depends on the radiofrequency tuned circuit formed by the trap capacitance and an inductor, as discussed in Section 2.4. We built a new inductor and amplifier circuit which is described in Section 8.3. We additionally reduced losses in the new trap by replacing the Macor spacers between electrodes with quartz, which has a lower dielectric loss.

Second, we reduced the radius of the new trap by a factor of two relative to the old precision trap, which decreases the damping time by a factor of four.

Third, the majority of the reduction in damping time is due to a new trap geometry that maximizes the geometrical constant κ . Intuitively, the coupling to the particle should increase if more surface area closer to the center of the trap is attached to the tuned circuit. The cost is that the sensitive RF tuned circuits of our amplifiers are easily disrupted by other wiring, particularly other tuned circuits. If we attach more electrodes of the trap to the cyclotron amplifier, it becomes extremely challenging to wire up the axial amplifier, spin drive, and cyclotron drive that we also need in the precision trap without unwanted RF couplings that ruin their performance. By using a single-purpose cooling trap, we can dedicate as much area of the trap as we need to the cyclotron amplifier. Details of the calculation of κ are in Section 6.2.2. Our proposed design increases κ by roughly a factor of three, for a nine-fold reduction in the damping time.

All these factors together should reduce the damping time from 200 seconds to 3 seconds. The total time to select a ultracold particle with n = 50 particle is


Figure 6.2: Diagram of the idealized potential on a split electrode.

thereby reduced from nearly two hours to roughly 20 minutes. The remaining cooling time would be dominated by the time required to measure the cyclotron energy in the analysis trap. This could be further optimized with a lower tuned circuit temperature that would reduce the number of selection attempts needed.

6.2.2 Calculating the optimized trap geometry

We calculate the geometrical constant κ by finding the scaling between a voltage on the electrodes and the resulting potential at trap center. We first demonstrate this by revisiting the damping of the particle's motion explained in Section 2.4. Then we solve the Laplace equation for the potential near trap center, take the first order approximation and extract κ , and find the electrode length that maximizes κ .

The geometrical factor κ was originally defined in terms of the induced current in Equation 2.23, but because of conservation of energy it is also proportional to a scaling between the voltage on the electrodes and a voltage at trap center. To see this, we start from defining the voltage induced by the particle's motion $V_A = IR$. This voltage produces a force opposing the motion equal to $e\frac{dV}{dx}$. Here V is the voltage at trap center due to V_A , which is

$$V\left(x\right) \approx V_A K x \tag{6.2}$$

to first order. This force slows the motion with a power equal to $-ef\dot{x}$, or $eKIR\dot{x}$. Since this power must be equal to the power dissipated in the resistor I^2R , we know that $I = eK\dot{x}$. By analogy to Equation 2.23, we can match the voltage scaling factor K to $\kappa/(2\rho_0)$.

The voltage near trap center can be derived using Laplace's equation [39, 68]. The potentials of our trap are shown in Figure 6.2. We define the x-axis perpendicular to the split where $\phi = 0$, and close the boundary by approximating V = 0 at $\pm L$ for a large distance L. We can expand the potential at the center of the trap in terms of solutions to the Laplace equation as follows:

$$V(\rho, z, \phi) = \sum_{j=0}^{\infty} \sum_{\substack{m=1\\m \ odd}}^{\infty} A_{jm} \cos\left(k_j z\right) \cos\left(m\phi\right) I_m\left(k_j \rho\right)$$
(6.3)

Where $k_j = (j + \frac{1}{2}) \frac{\pi}{L}$. To find the coefficients A_{jm} , we match this expression for $V(\rho, z, \phi)$ to the given boundary conditions:

$$V(\rho_0, z, \phi) = \sum_{j=0}^{\infty} \sum_{\substack{m=1 \\ \text{m odd}}}^{\infty} A_{jm} \cos\left(k_j z\right) \cos\left(m\phi\right) I_m\left(k_n \rho_0\right)$$
(6.4)

We can solve for the coefficients using the orthogonality of cosines:

$$A_{jm} = \frac{V_A}{2\pi L I_m \left(k_j \rho_0\right)} \left(\int_{-\pi/2}^{\pi/2} \cos\left(m\phi\right) - \int_{\pi/2}^{-\pi/2} \cos\left(m\phi\right) \right) d\phi \int_{-z_c}^{z_c} \cos\left(k_j z\right) dz \quad (6.5)$$

which simplifies to

$$A_{jm} = \frac{4V_A}{mk_j\pi L} \frac{\sin\left(k_j z_c\right)\sin\left(\frac{\pi}{2}m\right)}{I_m\left(k_j \rho_0\right)}$$
(6.6)

This gives us an expression for V in terms of the trap dimensions and the voltage on the damping electrodes.

To find κ , we start from the potential along the x-axis, take the first order approximation, and match it to Equation 6.2. The x-axis is defined by $\phi = 0$, z = 0, and $\rho = x$, so the voltage along it is simply

$$V(x,0,0) = \sum_{j=0}^{\infty} \sum_{\substack{m=1\\m \ odd}}^{\infty} \frac{4V_A}{mk_j\pi L} \frac{\sin\left(k_j z_c\right)\sin\left(\frac{\pi}{2}m\right)}{I_m\left(k_j \rho_0\right)} I_m\left(k_j x\right)$$
(6.7)

We can use the first order approximation for the Bessel function, $I_m(k_j x) \approx \frac{1}{2}(k_j x) \delta_{1m}$. To rearrange the potential into the form of 6.2, we scale x by $2\rho_0$. The result is

$$V(x,0,0) = \left(V_A \frac{x}{2\rho_0}\right) \sum_{j=0}^{\infty} \frac{4\rho_0}{\pi L} \frac{\sin\left(k_j z_c\right)}{I_1\left(k_j \rho_0\right)}$$
(6.8)

We see that the scaling between the electrode voltage is indeed proportional to $x/2\rho_0$ times a unitless constant, which we identify as κ_x .

$$\kappa_x = \frac{4\rho_0}{\pi L} \sum_{n=0}^{\infty} \frac{\sin\left(k_j z_c\right)}{I_1\left(k_j \rho_0\right)} \tag{6.9}$$

The maximum value of κ decreases the damping time by a factor of nine and occurs when $z_c \geq \rho_0$. Figure 6.3 shows the calculated value of κ as a function of z_c . The vertical lines indicate electrode lengths corresponding to connecting the amplifier across the ring of the precision trap ($\kappa = .335$), and across an electrode with a height equal to the trap radius. The calculated κ_x starts leveling off around 1.2 when $z_c = \rho_0$. The actual previous value of κ was .45 because the amplifier had been moved from a ring electrode to a compensation electrode before the ppm measurement. This value is not shown in the figure because it follows a different formula for an off-axis electrode that is calculated in Section 6.4. Nevertheless, a larger damping electrode increases κ by nearly a factor of three, and thus the damping time by a factor of 9.



Figure 6.3: Calculated κ_x as a function of z_c/ρ_0 . A solid vertical line indicates the height of the current ring electrode and a dashed line indicates $z_c = \rho_0$, the proposed height of the new trap

6.2.3 Impact of amplitude-dependent frequency shifts

We simplified the cooling trap design from the usual compensated five electrode trap to a three electrode trap consisting of only a ring and two endcaps. Additionally, the iron ring of the analysis trap is only 6.8 mm away (see Section ?? for discussion of the motivations behind this position0). However, both the proximity of the analysis trap and the simplified trap design should not cause amplitude dependent frequency shifts.

The average field seen by a particle in the cooling trap decreases as its radius increases, which shifts its cyclotron frequency. Additionally, our three-electrode trap approximates a harmonic potential less precisely than a five-electrode trap. Since the goal of the trap is damping the energy rather than precision measurement, energy dependent frequency shifts are only a problem if they are large enough to shift the particle frequencies out of resonance with the damping circuits. We show that the shifts in the axial frequency due to anharmonic terms are not large enough to do this, shifts in the cyclotron frequency due to trap anharmonicity are negligible, and finally that shifts in the cyclotron frequency due to magnetic gradients are only significant for the most excited particles.

Although the main purpose of the cooling trap is damping the cyclotron motion, we want to make sure the axial motion can be damped as well. Otherwise, any small noise drive at the axial frequency could cause an uncontrolled increase in the amplitude. To estimate the effect of anharmonic shifts, we first calculate the size of anharmonic terms in the potential before calculating the size of the resulting shifts for reasonable amplitude fluctuations for a particle with zero cyclotron energy. Finally we include the effect of a large cyclotron radius.

Both the fourth and sixth order anharmonic terms might contribute and must be calculated. We can minimize the lowest order anharmonic term C_4 through choice of z_c , but without compensation electrodes we have no free parameters left to adjust C_6 as we did for the five-electrode trap (see Section 2.1). We choose the ring length that produces the lowest value of C_4 , $z_c = \rho_0/1.203$ [69], because it also satisfies our goal of $\kappa \geq 1.2$. For that z_c , the second anharmonic term C_6 is about 0.1 and in practice, our machining tolerances can only guarantee $C_4 \leq 0.001$.

The size of the resulting shifts depends on the derivative of the potential and the expected axial amplitude. The amplitude dependent axial frequency for an on-axis particle is given by Equation ??. Since the anharmonic terms are a small correction to the frequency, we can take a Taylor expansion of Equation ?? and find the change

in frequency as a function of amplitude A and radius ρ :

$$\Delta f_z = \frac{1}{4d^4C_2} \left(C_4 \left(3A^2 - 6\rho^2 \right) + \frac{C_6}{d^2} \left(\frac{15}{4}A^4 - \frac{45}{2}A^2\rho^2 + \frac{45}{4}\rho^4 \right) \right)$$
(6.10)

When the radius is assumed to be zero, this reduces to the anharmonic shift calculated in [70]. For a particle with a thermal cyclotron radius, we find the axial amplitude can increase to 0.25 mm before the axial frequency is outside the 200 Hz bandwidth of the amplifier. Therefore, the axial frequency can be assumed to be stable against small fluctuations in energy.

We must also consider anharmonic shifts for a particle with a non-zero radius since the cooling trap's purpose is to damp particles with large cyclotron orbits. We find that for a particle where the axial amplitude has been fully damped, the axial frequency is resonant as long as the cyclotron radius is less than 190 μ m, or $n = 2 \times 10^8$. While we should bear this limitation in mind, this is a much larger excitation than we expect is necessary for measuring the cyclotron frequency (discussed further in Chapter 5).

Shifts in the cyclotron frequency due to anharmonic effects are essentially zero. The cyclotron motion of the particle is damped by an amplifier as intended. Since the cyclotron frequency only indirectly depends on the shape of the electrostatic potential through the Brown-Gabrielse Invariance theorem (see Section 2.2), the corresponding shift in the cyclotron frequency is reduced by a factor of 100. The cyclotron amplifier has a bandwidth that is almost 1000 times wider, so the odds of shifting the frequency outside the amplifier bandwidth are negligible.

However, the cyclotron frequency does change significantly with the particle's energy due to magnetic gradients. Growth in the amplitude and radius changes the average magnetic field sampled and thus directly shifts the cyclotron frequency $f_c = q |B| / m$ (for particle of mass m and charge q). This is discussed in more detail in Section 8.5, where we define the first and second order gradients as

$$\vec{B} = \left[B_0 + B_1 z + B_2 \left(z^2 - \rho^2/2\right)\right] \hat{z} + \left(-B_1 \rho - \frac{3}{2} B_2 z \rho\right) \hat{\rho}$$
(6.11)

The cyclotron frequency for a particle with amplitude A and radius ρ shifts in this gradient from f_c to

$$\tilde{f}_c = f_c \sqrt{\left(1 + \frac{B_2}{B_0} \left(A^2 - \rho^2/2\right)\right)^2 + \left(\frac{B_1}{B_0} + \frac{3B_2}{2B_0}A^2\right)^2 \rho^2 + \mathcal{O}\left(\rho^3\right)}$$
(6.12)

Even for very large excitations, the strong magnetic field we use, B_0 , is much greater than any shift due to the gradient $B_1 z$ or $B_2 \rho$. To leading order, we have

$$\tilde{f}_c - f_c = \frac{B_2}{2B_0} \left(A^2 - \rho^2 / 2 \right)$$
(6.13)

For our trap, $B_2 \approx 1400 \text{ T/m}^2$, which is comparable to the bottle used for electron g-factor measurements [71]. For a radius of 1 mm, the cyclotron frequency is shifted by only 5 kHz out of the 100 kHz bandwidth of the cyclotron amplifier. The shift for the same axial amplitude is larger by only a factor of two. Even particles that have been excited more than halfway to the wall of the trap are still resonant with the cyclotron amplifier.

6.3 Analysis trap

The new analysis trap is designed is designed to minimize the cyclotron transition rate in order to increase the axial stability of particles over a wider range of cyclotron energies. Though not conclusively established, we present reasons to suspect that these transitions are driven by electrical noise reaching the particle from outside the trap. We propose a trap design that uses the electrodes themselves to shield the trap center, and calculate the necessary trap dimensions.

In Chapter 3.2.1, we showed that scatter in the axial frequency can be explained by electric field noise driving cyclotron transitions, each of which shifts the axial frequency by 50 mHz. It has also been suggested that the electric field noise might be caused by patch potentials on the surface of the electrodes [36], often referred to as 'anomalous noise' in the ion trapping community [72]. However, most studies of surface noise in ion traps have been done at room temperature, and scaling from 300 K to 4 K may not be reliable. Many proposed mechanisms for the anomalous noise would freeze out above 4 K [72]. Furthermore, the size of the drive we infer from our observed transition rates is in line with either Johnson noise from a small 4 K resistance that is not filtered, or a 50 ohm room temperature Johnson noise that has somehow partially bypassed our tripod filters, as calculated in Section 3.2.1. Such noise could be filtered before it reaches the center of the trap.

The redesigned analysis trap attempts to do this by electrically shielding the trap center using the trap electrodes themselves. In the new analysis trap, we use only cylindrically symmetric ring and compensation electrodes. In order to apply a magnetron cooling drive as described in Section 2.7, we split one of the endcaps. Since the endcap is further out, the effect of the magnetron coupling drive at trap center is significantly reduced so we will need to increase the power used. However, any noise that reaches the split electrode by any method is similarly reduced, effectively shielding the particle from sources of cyclotron noise. This eliminates the possibility

of noise bypassing filters or radiatively coupling to the magnetron drive line.

It is still possible that noise on the spin flip drive line can by picked up by the particle, but this is thought to be less significant than the sideband drive line. The spin flip line could cause a radially asymmetric drive because it must attach at two points to a compensation electrode (see Section 2.9). However, the spin flip line is heavily filtered at the cyclotron frequency as described in [36], and inside the trap can it is carried by a coaxial cable which should provide better protection against noise radiating from any other source in the trap can. It was thus judged more likely that the sideband drive line across the split electrode was coupling noise into the trap, and by moving it away from the trap center we can reduce the noise seen by the particle.

To calculate the effect of a drive electrode, we follow a similar calculation to Equation 6.9, except now we need to expand the potential at the center in both x and z because we need to create an xz gradient for magnetron cooling. The free parameters determining the coupling are the height of the drive electrode and its distance from the trap center. The result is

$$\kappa_{xz} = \frac{4\rho_0^2}{\pi L} \sum_{j=0}^{\infty} \frac{k_j \sin(k_j z') \sin(k_j z_c)}{I_1(k_j \rho_0)}$$
(6.14)

To achieve the optimal reduction in coupling, we chose to apply the drive on a second endcap which is separated from the trap by an additional isolation endcap. The dimensions of the ring and compensation electrodes are fixed by the compensated trap design for a harmonic trap described in Section 2.1. Early calculations showed that a drive electrode directly adjacent to the compensation electrode would reduce the coupling of the magnetron-axial drive κ_{xz} by at most a factor of 7. This would improve the stability of an n = 300 particle by more than the desired amount, but



Figure 6.4: Coupling of the new drive electrode, relative to the old electrode, as a function of separation between electrode center and analysis trap. The dashed line indicates the center position of an electrode placed just below the compensation electrode, and the solid line indicates the actual center of the drive electrode

ideally if we are rebuilding the trap we would reduce the noise enough that we could accept a larger fraction of the thermal particles and reduce the time spent on cooling. Thus, we decided to split the endcap into two electrodes, a thin isolating endcap just outside the comps and a split drive electrode just beyond that.

The position of the drive electrode was determined by the requirement that the height of isolating endcap be no smaller than 1.5 mm. We were concerned that a very short inner surface would be difficult to machine out of the soft copper used for electrodes and easy to damage. A height of 1.5 mm was chosen because features of that size had been successfully fabricated for the previous analysis trap. That requirement places the center of the drive electrode 4.2 mm away from the analysis trap, where the drive is already suppressed by a factor of 300.

We still expect to effectively cool the magnetron motion despite the large reduction in drive power. We typically used drives of up to 600 nW to magnetron cool particles heated after transfer into the analysis trap. If the drive voltage is reduced by a factor of 300, we would need to put 20 mW in at the hat. The total heat load of the experiment is roughly 110 mW based on the rate of helium boiloff, so this would have a noticeable effect on the hold time. There be may other undesirable effects if the heat is not able to dissipate directly to the dewar but instead locally heats the trap stack or tripod electronics. For commissioning, we wired up the analysis magnetron drive line without an RC filter and a capacitive divider that had been used previously for this drive line. Together, they attenuated the SB drive by 30 dB. Thus, the net reduction in drive voltage is only a factor of 10, and the hat drive power only needs to increase to 60 uW.

6.4 Interactions between the cooling and analysis traps

The relative position of the cooling and analysis trap is important to their functions. We wish to place them as close as possible to simplify transfers between them, which otherwise will be the limiting factor on the reduction of the total cooling time. However, if they are too close together the large resistance of the cooling amplifier may broadcast thermal noise close to the cyclotron frequency into the analysis trap. Ultimately, the cooling trap was positioned directly adjacent to the analysis trap endcaps with 6.8 mm of separation between trap centers. This is close enough that transfers should be significantly faster than from the precision trap, but far enough that noise from the amplifier is shielded by orders of magnitude. The total time to select a cold particle is limited by the time spent transferring and reducing the magnetron radius subsequent to transfer due to the minimal τ_c of the cooling trap. We also sought to make transfers from the cooling trap faster than from the precision trap, though the factors affecting this are less well understood. Our first thought was to reduce the total distance is reduced by a factor of 10, which should prevent increase of the magnetron radius during transfers due to small misalignments. It also made feasible the reduction of the length-to width ratios of the intervening electrodes by up to a factor of three. A long length-to-width ratio is believed to cause to heating of the magnetron motion during transfer [51]. This is attributed to to torques from patch potentials becoming significant in the low-field regions at electrode center. We also reduced the number of electrodes in between the traps and thus the number of steps that must be taken during transfers by a factor of two. This final factor is expected to reduce the necessary transfer time by roughly forty percent, and overall gains may be even greater due to reduced magnetron heating.

However, we were concerned that at this position electrical noise from the cooling trap amplifier could reach the analysis trap, were could drive cyclotron transitions. The effective resistance of the amplifier produces Johnson noise from a frequencydependent resistance (see Section 2.4), which can be treated as white noise since the bandwidth of the tuned circuit is much larger than the linewidth of the cyclotron resonance in the bottle, we can treat this as white noise driving the particle. If we substitute $v^2 = k_b T R$ into Equation 3.9, we can get an expression for the transition rate as a function of the coupling constant γ :

$$\frac{dn}{dt} = n \frac{k_b T}{\hbar \omega} \left(\frac{e\kappa_x}{2\rho_0}\right)^2 \frac{R}{m} = n \frac{k_b T}{\hbar \omega} \gamma \tag{6.15}$$



Figure 6.5: Calculated κ_x as a function of distance between the analysis and cooling traps. After roughly a radius separation, κ_x declines exponentially with further distance.

We derive a minimum acceptable value for gamma by requiring that the transition rate is at least an order of magnitude smaller than measured in the old analysis trap in Section 3.2.1, so that it would not limit us if we can reduce other sources of noise enough to improve the stability as proposed in the next section. Then we can invert this equation to find that the maximum acceptable coupling constant is $\gamma = 3 \times 10^{-9}$.

To find the minimum acceptable separation, we need to calculate how κ_x and thus γ change as we move the amplifier a distance z' away from the center of the analysis trap. This can be calculated from Equation 6.9 if we expand around z = z' rather than z = 0:

$$\kappa_x = \frac{4\rho_0}{\pi L} \sum_{j=0}^{\infty} \frac{\cos\left(k_j z'\right) \sin\left(k_j z_c\right)}{I_1\left(k_j \rho_0\right)} \tag{6.16}$$

The variation of κ_x with z_c is shown in Figure 6.5. The relationship is complicated close to the center of the analysis trap, but after $z' \approx \rho_0$ between the two traps, the decline of κ_x with further distance fits well to an exponential.



Figure 6.6: Cyclotron frequency as a function of distance from analysis trap z. The black line indicates the center of the cooling trap.

In practice, γ declines even faster than κ_x^2 does because the effective resistance also changes with separation between the traps. The cyclotron frequency shifts quite rapidly as the amplifier is moved away from the strong magnetic gradients of the analysis trap, as shown in Figure 6.6. A circuit tuned to the cyclotron frequency in the cooling trap will not be resonant at the analysis trap cyclotron frequency. The exact resistance at the analysis frequency is determined by $R(\omega)$ of the tuned circuit, which has a Lorentzian shape [68]. The tuned circuit is currently about 200 kHz wide, so even just a few millimeters away the analysis trap cyclotron frequency is already separated from the tuned circuit peak by 30 line-widths. This can be seen in Figure 6.7, which shows γ corrected for this effect. The horizontal line shows the maximum allowable γ calculated earlier. It occurs when the cooling trap is centered around 3 mm away from the analysis trap. The chosen position of the cooling trap is highlighted on the plot with a vertical line, showing γ is nine orders of magnitude below the maximum. Thus, the cooling trap amplifier will not drive appreciable



Figure 6.7: γ as a function of z, corrected for shifting frequency. The gold line represents the previous amplifier and the blue line represents our new amplifier which has a 2x higher resistance and a narrower response

transitions in the analysis trap.

6.5 Conclusion

This chapter presents our trap design and explains in detail the design principles of the cooling and analysis trap. The new design improves spin flip detection in the analysis trap in two main ways: by making it easier to cool the cyclotron energy of the particle to the low temperatures necessary for single spin flip stability, and by shielding the analysis trap from electronic noise that drives instability. The second half of the chapter presents a scheme for measuring both the cyclotron and spin frequencies in the precision trap using a separated oscillatory fields technique. Such a method would allow both frequencies to be probed in the same field, and also has the advantage of relatively little heating of the cyclotron motion, and thus could easily be extended to sub-ppb precision. The construction of the new trap and its early commissioning are presented in the following two chapters.

Chapter 7

Experimental Apparatus

Implementing the new apparatus presented in the last chapter required fabricating new electrodes and several other construction projects. Chapter 8 discusses new amplifiers that were constructed for the new traps and the loading trap, where interactions between the amplifiers and the magnetic gradients are causing difficulties detecting particles. This chapter presents the construction of the new electrodes, electrical connections and filters for the new traps, new thermal shields, and a new alignment system. Most topics discussed here are adapted from the old design to accommodate a twofold increase in the number of electrodes, and all incorporate new ideas for better performance.

7.1 Electrode construction

New electrodes were constructed following the design of the last chapter. The most important concerns in constructing the trap electrodes are the accuracy of their dimensions, their surface quality, and their durability. As described in Section 2.1, machining tolerances set the ultimate limit on the level of anharmonicity achievable in the trap, which can broaden frequency measurements or cause a less stable response. A rough surface may cause imperfections in the cylindrical symmetry of the potential that would heat the particle, as explained in [35, 36] and [58]. The electrodes must be very carefully handed to avoid deforming them as leads are attached and drive electrodes are divided into halves. The gold plating method was re-conceived for greater durability as the previous plating had suffered from wear. Finally, a positivepressure cube was constructed to protect the electrodes during the extended assembly process.

7.1.1 Polishing, brazing, and splitting

For a high-quality surface and the highest dimensional accuracy, the inner surface of the electrode must be hand-polished after machining. The electrodes are all made from OFHC copper, with the precision trap made from 5N purity copper to minimize its susceptibility and magnetic gradients. A new polishing procedure was developed [58] for better surface quality and to quickly process the larger volume of electrodes in the new stack.

The electrodes were polished with a cotton swab dipped in diamond lapping paste, while spinning in a lathe¹. A custom-made Teflon sleeve protected the cylindrical symmetry of the electrode from deformation due to the lathe collet. We polish with series of lapping pastes of decreasing grits down to a 1 μ m grit size. After polishing

 $^{^{1}}$ Typical speeds started at 200 rpm and increased up to 1000 rpm for finer grit



Figure 7.1: Polishing electrodes on a lathe, with Teflon collet and diamond lapping paste

with each grit for several minutes, the electrode finish is checked under a microscope to verify that the finish is even and any deeper scratches have been removed. Before proceeding to smaller grit, we sonicate the electrode for 5 minutes in de-ionized water to remove any trapped particles of larger grit size. After 8 μ m we carefully measure the inner diameter using a gauge rod sized 0.002" larger than the target inner diameter. This process can scratch the surface, so after the rod fits we finish polishing to a mirror finish without checking again. Based on test pieces, this should get us within ± 0.0002 inches of our desired final radius.

After polishing, high purity wire leads were inserted into the small side-holes of the electrodes and brazed with a Ag-Cu eutectic at 800 C in a hydrogen brazing oven. This also anneals the copper, which must be done before splitting the drive



Figure 7.2: Electrodes in a fixture prepared for brazing in the hydrogen oven.

electrodes.

Annealing the electrodes also changes the surface structure by promoting grain formation. After polishing to a homogenous surface before brazing, subsequent to annealing grains were clearly visible under a microscope. Annealing can cause grains to grow up to several hundred μ m in size [73]. The final size of grains is generally larger if the piece was worked more before heating or if the heat was maintained for a longer time period. This correlation was observed on test pieces that were brazed with a slower ramp to temperature and longer total time above the annealing point of copper, resulting in larger grains as predicted.

We restored the homogeneity of the surface by lightly polishing after brazing. Annealing also softens the metal so that holding the electrode in a lathe collet for further polishing would risk distorting its cylindrical shape. Instead, we gently held the electrode in a gloved hand and swabbed it with 1 μ m paste until the grains were no longer visible under a microscope. Several of the electrodes need to be split into two halves so that we can apply noncylindrically symmetric drives to excite the radial motions of the particle. Annealing relieves the strain from machining while the part is still cylindrically symmetric so that the electrode will not relax out of round when divided [74], so we perform this step last. We use wire electrical discharge machining (EDM), which removes material with repeated sparks between a wire electrode and the work piece. The wire electrode's motion is computer controlled, so we cut circular recesses for the sapphire ball spacers between the two halves at the same time. We machined tight-fitting Teflon plugs which protected the inner surface from pitting due to the EDM.

7.1.2 Plating

The electrodes are plated with gold to prevent oxidation for long-term maintenance of the surface quality. The electrodes of the old trap had developed small bare patches in the plating by the time we opened the trap stack in 2016 as part of troubleshooting persistent problems detecting particle signals. If these patches were oxidized they would charge up when particles were dumped from the trap. The resulting potential could severely distort the harmonic shape of the well, possibly explaining inconsistent signals. It is unclear if these potentials were the source of the problem as older traps had operated successfully without plating the electrodes at all, but any surface inhomogeneities are highly undesirable. A new method was developed for an extremely durable plating.

The bare patches were attributed to the thinness of the original gold plating. The old precision trap was coated using a thermal evaporation process. The thin plating was a deliberate choice intended to minimize the size of grains in the deposited gold [35]. However, gold plated directly onto copper will diffuse into it at a slow rate. For the 100 nm thickness used on the old trap, 1 percent of the surface area would be bare of gold within 3 months to 3 years [75]. The process is temperature dependent, so it likely only advanced during periods when the experiments were warm. Our trap electrodes had been in operation for ten years by the time the bare patches were discovered, with several months of that time in total warm due to shipping and repairs.

To avoid this problem in the new electrodes, we decided to use an electroplating process with thicker platings. In a typical commercial plating process, a nickel barrier layer is used to prevent this diffusion, but we avoid the use of nickel because it has a large magnetic susceptibility. The evaporation process used on the previous electrodes required specialized training and custom jigs that would be difficult to assemble in our highly constrained time frame, and had never been tested with the thicker coatings we now believed were necessary. Similar experiments in this lab in the 1990s used an electroplating process that could easily achieve much thicker coatings [68]. We decided to revert to the roughly 2 μ m thickness of those experiments with a similar electroplating process.

As we were also operating under a severe time constraint, we had most of the electrodes plated by a commercial plater. A 1-3 μ inch thick flash gold treatment was applied to each electrode before they were transferred to a soft pure gold bath where they were plated to 80 μ inch thickness (2 μ m). Measurements on a test piece indicated a possible 30 percent variation in plating thickness along the length of an



Figure 7.3: Split electrodes prepared for dipping in a bath at the commercial platers. These pieces were test pieces whose thickness was subsequently measured

electrode during batch plating, motivating a 2 μ m thick plating to guarantee adequate thickness along the length of the electrode.

The high thickness did result in grains visible to the naked eye. For some of the smaller electrodes, we carefully polished them by hand after plating with 1 micron paste until the surface looked homogeneous under a microscope, but for the 1.2 cm diameter electrodes we assumed that the particle would be far enough away that small grains would have a minimal effect (since previous traps of this size did not post-polish their electroplated electrodes [68]).

The flash gold solution contains about 2 percent cobalt, which is ferromagnetic.



Figure 7.4: Photograph of assembling the trap inside the positive-pressure cube. The jig is used to hold the stack together until the top plate can be attached with rods and nuts

We believe such a small amount of magnetic material should not impact the homogeneity of the magnetic field. Nevertheless, we plated the most sensitive electrodes of the precision trap in house using a sulfite gold solution (Technic 25) with no ferromagnetic material.

7.1.3 Trap assembly

To protect the electrodes from dust, we assembled the trap inside a positivepressure cube built from 80/20 aluminum, clear plastic, a fan, and a HEPA filter. The structure is pictured in 7.4. The trap is held together by beryllium copper rods between two endplates, which are secured with beryllium copper springs and molybdenum nuts. The two plates were measured to be parallel to within 0.1 mm after tightening all the screws, and then a second nut was added to prevent loosening on thermal cycling.

7.2 Wiring

Roughly doubling the number of trap electrodes doubles the number of electrical connections within the very limited space available at 4 K. In addition, noise or a poor electrical connection was a secondary hypothesis for problems detecting particles before the trap rebuild. This motivated the replacement of the old wiring with more reliable, better isolated connections. The full wiring from the electrodes to the hat is diagrammed for the precision trap in Figure 7.9, for the analysis and cooling trap in Figure 7.10, and for the loading trap in Figure 7.11. The final wiring choices represent a compromise between fitting all the necessary connections for 40 or so electrodes within the available space, reliability of the connection against shorts or breaks developing during thermal cycling, and isolating the most sensitive circuits from RF noise.

Each electrode needs a filtered line for applying a DC potential. All DC lines have five pole LC filter to block room temperature radiofrequency noise. While in the previous version all of these filters were in the 4K tripod region, space constraints and the large number of new lines necessitated moving most filters to the baffle above the dewar which supports the radiation shield at 20-50K (see discussion of the cryogenic system in 2.11.1). The 'high' and 'low' lines are separately filtered to ground, and ground plane of the boards is grounded to a copper support structure mounted on the baffle. The support structure was designed for dense packing of boards while still allowing easy access with a screwdriver to remove individual boards for repair (see Figure 7.5). To reduce ground loops, the low side of each twisted pair is only connected to ground the capacitors until the pinbase, which the common ground for



Figure 7.5: Photograph of filter boards mounted on top of the radiation shield baffle

the experiment.

The analysis trap DC voltages and the amp circuits that detect the signal are most sensitive to noise, so those filter boards are still mounted in the 4K tripod region. In the analysis trap we must keep the frequency stable enough to distinguish a 60 mHz spin flip. A 350 nV fluctuation in the ring potential would cause a comparable shift in the particle frequency, so we have a very stringent requirement on the stability of the trapping potential. The lower temperature reduces Johnson noise generated in the circuit itself, and its proximity to the shielded trap cap reduces opportunities for noise to be picked up after the filters. These lines from the hat pass through a separate, shielded tube in the dewar.

All rings, comps, and endcaps have an additional RC filter composed of the 1 Mohm resistor on the filter board and 1-10 μ F capacitor to ground in the tripod, with the largest filters on the rings and in the analysis trap. All lines have at least a 1 ms RC filter just above the vacuum feedthrough pin to filter out any noise radiated from other parts of the pinbase. All of the transfer and endcap electrodes have large resistors to ground at the pinbase, so that when not transferring particle between traps they can be simply grounded by disconnecting their bias lines at the hat.

The filter designs were modified in the rebuild to incorporate a diode circuit on those electrodes with time constants of 5 seconds or longer. The large filter capacitors used on our experiment are necessary given our extreme sensitivity to electrical noise, but they limit how quickly we can change the DC biases of each electrode. Since transfers between traps is done by changing the voltage of each intervening electrode in several steps, waiting 1-5 seconds adds up to an expected transfer time between the precision and analysis trap of more than two minutes. With the rebuild of the trap, the filter designs were modified to be smaller in space and to incorporate a diode circuit on those electrodes with time constants of 5 seconds or longer. We used an integrated transient voltage suppression chip (USBQNM50403E3/TR7CT-ND). The board was tested at 4 K with a 1 kHz square wave and found that the diodes successfully changed the time constant of the last RC filter when the voltage was raised above 10 V.



Figure 7.6: CAD model of the tripod region with five amplifiers and the five largest filter capacitors. Filter boards, coaxial cables, and the rest of the filter capacitors must fit around them.

Before entering the hat, each DC line is passed through one or more a relays controlled by a computer-controlled digital I/O card (Opto-22). The relays allow the electrode bias source to be switch or to be disconnected from the trap altogether for low-noise operation. For transfers, the DC bias voltage is supplied by a high voltage and relatively noisy SuperElvis source to a precision voltage source. For working with particles in a trap, a low-noise and high stability voltage source is used for the ring and compensation electrodes. The analysis and precision traps share a high-stability



Figure 7.7: Arrangement of cables through the support structure of the experiment

Stahl UM 1-14 source. A Fluke 5440b precision voltage calibrator was used to bias the loading trap, as the resonant voltage is above the 14 V range of the Stahl source. Each bias passes through vacuum feedthroughs at the hat and are carried to the tripod by constantan twisted pairs.

Weak or very high frequency RF signals are carried by stainless-steel coaxial cables (Microcoax UT-34-SS-SS). Stronger or low-frequency drives as well as temperature and magnetometer sensor connections are carried by constantant wisted pair, which can alter the signal phase but has better thermal isolation. The paths from the hat to the dewar were chosen to maximize the spatial separation and shielding of sensitive DC lines from noisier, sensitive DC lines, while minimizing the possible loop area between lines involved in the same circuit.

For the precision trap, Copper straps carry signals and biases from the pinbase



Figure 7.8: Photo of components and wiring inside the trap stack

to the electrodes (as was done in the old precision trap). We only used very low susceptibility materials within a few inches of the precision trap to avoid magnetic gradients. The electron g-2 experiment had tested common 'non-magnetic' capacitors and resistors and found residual magnetism, so we avoided placing any electrical components near the trap. Since the precision trap is closest to the pinbase, the straps are short and their impedance is low even at RF frequencies.

For the other traps, straps would need to be longer and would present a less

reliable return path for RF currents to pinbase ground. Those traps are less sensitive to small gradients. The number and density of connections was also now high enough large bare straps would need to be located uncomfortably close to each other, with a risk of shorting on thermal cycling. Instead, for the rest of the stack we used standard Teflon-coated 30 AWG wires that take up less space and have some protection against accidental shorts. RF currents are controlled with blocking resistors and shorting capacitors inside the trap can, which is essential for good performance of the sensitive RF detection circuits discussed in Chapter 8. Additionally, connections to adjacent electrodes pass through adjacent pins on the pinbase as much as possible in order to minimize the length across the pinbase of any current loops.

Every element inside the trap can must be completely reliable through thermal cycling to prevent the delay and disruption of opening the trap for repairs. Every grounding capacitor has a duplicated connected in parallel in case of a lead breaking during a thermal cycle. The same strategy was used for every blocking 1 MOhm resistor. We tried to avoid inductors inside the trap can because the nonmagnetic low-loss inductors we use are fragile and have on rare occasions had pads break on cooldown.

In some lines we must block cyclotron frequency currents while allowing axial frequency drives to reach the electrode. For most of those lines we made some simple inductors by wrapping the 30 AWG wire of the lead into a coil around a Teflon tube. This was not possible for the loading trap axial drive, because we also need to put a 100 ns pulse down the line for ejecting electrons. Instead, we used a 1/4-wavelength coaxial line for the cyclotron frequency. This transforms the impedance of the pinbase

grounding capacitor into a large impedance at the input of the coax and vice versa. Currents at the cyclotron frequency are forced to travel through the resonator circuit described in Section 8.3 rather than through the possibly lossy axial drive line.



Proton wiring 2018

Figure 7.9: DC and RF wiring for precision trap



Figure 7.10: DC and RF wiring for analysis and cooling trap



Figure 7.11: DC and RF wiring for loading trap

7.3 Thermal and magnetic shielding

We additionally made two changes to improve the thermal isolation and magnetic shielding of the experiment. The rest of the support structure of the apparatus is mostly the same as was used for the 2012 measurement and described in 2.11.1. See 7.12 for a complete section view of the magnet with experimental insert. The new apparatus has a modified thermal shield at the top of the experiment and a new thermal shield above the PPAC.

7.4 Thermal shield and magnetic shielding

In order to access antiprotons our experiment must be located in the magnetically noisy environment of a particle accelerator hall. The Antiproton Decelerator ramps the field in the magnets of its rings in several steps as the beam energy changes over its 120 second cycle, which causes a recognizable pattern of magnetic field fluctuations in our zone of about 40 mG peak to peak. That is about a 700 ppb variation relative to our 5.6 T field.

This problem has been known since the first precision antiproton experiments at CERN, and the solution was to install an extra compensation coil in the main solenoid [68]. The coil is a closed loop of superconducting wire that will conserve the flux through it – any external changes to the magnetic field will induce opposing currents in the loop. Through a careful choice of the coil geometry, we ensure at one special point, generally the center of the solenoid, conserving the flux means conserving the field at that particular point as well. Machining tolerances mean the


Figure 7.12: Solidworks model of apparatus in place inside the magnet

shielding is not perfect, but high shielding factors $S = \Delta B_i / \Delta B_o$ can be achieved. A ppb charge to mass ratio comparison in 1996 measured a shielding factor of 150 using

Figure 7.13: Magnetic field variation due to the AD cycle in the zone and at the center of the solenoid.

the same magnet [68].

That shielding factor drops rapidly when the precision trap is displaced from the center of magnet, as it was for the antiproton measurement. For that measurement, we decided we need to place the analysis trap above the precision trap so that the incoming beam could go straight into the precision trap without first passing through the smaller diameter and strong magnetic gradients of the analysis trap. The experiment hat was positioned as high as it could go in the bore before the thermal shield loses contact with the bore causing the hold time to drop rapidly. At that position the precision trap is still 44 mm below the center of them magnet. We measured a shielding factor of only 16 and an oscillation in the cyclotron frequency due to the AD field of 40 ppb from peak to peak. This was not significant at the ppm level of precision, but would have greatly complicated a sub-ppb measurement.

We designed an extension to the thermal shield that could fit over the radiation shield so we could raise the experiment until the precision trap was back at field center. The gap between the radiation shield and the walls of the bore is only 1/8"radially, so the extension piece needed to be no more than 1/16" thick. It was turned out of OFHC copper by the Harvard SEAS Scientific Instrument shop. Lower profile beryllium copper fingers (Laird Technologies 97-223) were soft-soldered around the bottom edge. A groove cut under the fingers allows them to bend inward freely when making contact. To ensure that the copper shield cannot press directly against the radiation shield when in position, 1/16" spherical divots were machined into the



Figure 7.14: Solidworks model of the new thermal shield extension that fits over the old radiation shield

radiation shield and mating nylon balls were glued into them with cryogenic glue. Although the relative position of the precision trap has changed with the four trap stack, we still need to keep the apparatus slightly higher in the magnet bore using the extension piece for reasons explained in the next section, and we use spacer electrodes to keep the precision trap positioned at field center. We noticed no change in hold time after switching to the new system and lifting the precision trap to field center.

In Figure 7.13, we can see that we can still detect the AD cycle in cyclotron decays,

but the peak to peak variation is now only 0.4 Hz or 5 ppb. This returns us to the high shielding factor we expected which will protect against incidental magnetic field noise. The AD cycle itself is still significant for our measurment and we will use the loading trap presented in the Section 8.6 to allow the delay of our most sensitive measurements until after the AD has shut down for the year if necessary.

7.4.1 New bottom radiation shield

A longer helium hold time would allow more stable operation between helium fills. To extend the hold time at CERN, we installed a new radiation shield below the experiment. This did not affect the hold time, disproving our theory about why the hold time has been consistently shorter at CERN than harvard, and the shield is still in place.

In 2015, we added a radiation shield below the bottom of the experiment. Before modification, the 50 K radiation shield around the helium dewar was directly facing the top of the PPAC, which is kept close to 300 K by a constant flow of N2 for stable operation. We added a small copper plate that mounts to the top of the PPAC through 1/2-inch nylon standoffs. A 0.625" diameter hole in the middle allows the beam to pass through. It is heat sunk to the 77 K bore of the magnet with large beryllium copper fingers.

Since radiated heat loads go as T^4 , this should have reduced the heat radiated onto the radiation shield from below by a factor of 9. Unfortunately, we did not see any improvement in the hold time when we reinstalled the system, which means that exhaust cooling of the radiation shield was sufficient to handle this extra radiation.



Figure 7.15: Photograph of the new shield mounted on the PPAC before re-installation

We are still searching for an extra heat source at CERN large enough to explain the discrepancy.

This does mean, however, that the experiment needs to sit at least 3/4-inches higher in the bore than it did previously, or it will hit the new shield on top of the PPAC. There are three spacer electrodes above the precision trap in the stack to position it within a few mm of field center when the bottom of the experiment is positioned just above the PPAC shield.

7.5 Magnetic shielding and stability

In order to minimize drifts in the magnetic field of the solenoid, we have to stabilize the pressure of the various cryogen spaces of the experiment. Controlling the pressure of the helium dewars should stabilize temperature as well most of the experiment is well heat sunk to the dewar whose temperature is set by the boiling point of liquid helium. Experimentally, controlling the pressure of the magnet and auxiliary nitrogen shield can also improve stability. A control system has been constructed that uses analog pressure transducers (Omega PXM409-880HB10V) and digital process controllers (Omega CN63300-R1-R2-F3-LV) to send an analog control signal. By changing the setpoint of a mass flow controller (MKS M100B01353CS1BV), the flow out of the experiment can be increased or decreased to stabilize the pressure. Nitrogen is released to atmosphere, but helium must eventually be released into the AD recovery line that varies between 0 and 200 mbar in pressure depending on flows of helium from other experiments in the AD hall. We added a 4 cubic meter gas bag as a buffer (Flexi-Liner) to prevent this variation from affecting the stability of the front end. Helium can be released into the bag at room pressure, and then when the pressure in the bag starts to rise helium is pumped out of the bag into the recovery line with an XDS scroll pump.

We also needed to ensure the spatial homogeneity of the field by adjusting the shim coils. While this was previously done in our magnet for the charge to mass ratio measurement [68], we adopted a slightly different procedure. Our solenoid is equipped with X, Y, 1Z, XY, YZ, XZ, and $X^2 - Y^2$ superconducting shim coils. The Z^2 gradient is adjusted by changing the ratio of the currents in the inner and outer sections of the main coil. We inserted a water NMR probe into the center of the magnet. Four nylon standoffs arranged around the probe support pushed the probe 1 cm off-center. The probe was rotated in 45 degree increments to scan the X,Y and Z directions, and moved vertically in 1 cm increments. The resulting data was fit to find the second order gradients, and shim coils were iteratively adjusted. Shimming was stopped when the largest gradients were $1.15 * 10^{-8}$ mm² for Z^2 and $1.75 * 10^{-8}$ mm⁻¹ for the Y gradient.

7.6 Alignment

A new sensor system has been constructed to measure the alignment of the electrode stack and magnetic field can be measured to high precision. Once measured, the alignment can be corrected using an adjustable spacer made from a bellows and three screws. The angle of the trap relative to the field is constricted by the small space between the radiation shield and the bore to \pm .15 degrees, but smaller variations in angle can cause heating or particle loss during transfer. The most precise way to check the misalignment is to measure all three particle frequencies and find the discrepancy between the predicted and measured magnetron frequency [76]. However, we require a method that does not rely on a trapped particle to quickly verify that the alignment is within the acceptable range for transfers.

Transfers have usually worked to some degree when the alignment was correct to $\pm .1$, so at a minimum we wish to measure the angle to that precision. Even higher precision would be desirable as misalignment as small as 0.01 degrees will cause noticeable magnetron heating during transfer due to the particles arriving off-center. Even a system with an offset, as long as it is consistent, would let us recover the 'good' alignment for transfers quickly after a cooldown instead of needing a painstaking process of trial and error every time.

Alignment measurement methods that had been used successfully at Harvard would interfere with the antiproton beam. Initial alignment at Harvard was done visually through a pinhole in the bottom flange of the magnet allowed visual alignment by looking up through the bottom of the magnet. Later a charge collecting target was mounted at the bottom of the stack opposite the FEP [35]. At CERN, the



Figure 7.16: Solidworks model of the magnetometer mounted on the endplate of the trap stack

antiproton beamline is mounted to the bottom of the magnet leaving no room for a pinhole. Additionally, the charge collecting target would block the antiproton beam from entering through the bottom of the trap.

Instead, we constructed a new system based on Hall effect magnetometers. The Hall effect refers to a voltage generated when current is passed through a conductor perpendicular to the magnetic field. A Lorentz force from the field on the moving charge carries will push them to one side of the conductor, causing a voltage to develop that can be measured with a multimeter. If the trap stack (and magnetometers) are perfectly aligned with the solenoid field, the magnetic field will be perpendicular to the current and there will be no force or voltage. We used HHP-NU sensors from Arepoc, which have a sensitivity of 70 mV/T. They are mounted in parallel with the trap stack to measure perpendicular field due to misalignment.

The magnetometers must be mounted inside the trap can. The bottom of the tripod region where most electronics are located is several inches away from precision trap, where the solenoid field is no longer homogeneous and beginning to decline. The spatial position of the sensor in the in-homogeneous field has a greater effect on the measurement than its angle, making it difficult to isolate the effect of the alignment.

The sensors are mounted on the stack endplate as shown in Figure 7.16, which pushes the stack into place against the pinbase, determining its alignment. The relative angle between the magnetometer and the center axis of the trap is then determined by only the joint between the mount and the baseplate, and the bottom spacer of the stack fitting into its socket. The magnetometers are pushed against their mounts using EDM-machined washers. This ensures alignment with the precisionmachined surface and allows them to be easily repaired if necessary.

The gradient from the analysis trap ring causes an offset in the magnetometer reading at this position. The iron ring bends the field by about .25 degrees at the center of the sensor. However, the relative position of the analysis trap and the sensor does not change as the apparatus angle is adjusted from the hat, so this should only add a constant offset. The new magnetometers typically measure angles of around 0.2 degrees when transfers are successful, which agrees with the offset predicted from this effect.

We find that adjustments of only 0.1 degree can make the difference between

success transferring particles from the loading trap to the precision trap or not. Most particle loss happens as we leave the narrowest part of the stack after ATCONE, which is not consistent with the particles running into the wall in an area of small trap diameter. Additionally, a 0.1 degree angle would imply the magnetron radius of the particle increases by only 0.3 mm as it moves from the loading to the precision trap.

We explain this by first noting that as particles leave the analysis trap they pass through a region with a strong magnetic gradient. If the particles are not centered in the trap at that point, either due to a large magnetron or cyclotron radius, gradients can force them far out of trap center. Alternatively, there may be a 'magnetic bounce' effect as particles are adiabatically transferred through a magnetic field. The cyclotron energy changes by 2 percent in this region, and this energy must be conserved through a commensurate change in the axial energy. If the cyclotron energy was not well cooled because the damping time for this motion is very long (see Section 8.3.2), the axial motion could become so excited that the particle leaves the trap.

The magnetometer reading associated with good transfers is very stable over short time periods. The scatter on the reading due to voltage noise is equivalent to about .03 degrees, or .07 degrees in the less sensitive low-current mode we were often running in to reduce heat load.

The reading from the magnetometers can drift over time, but we interpret this as a real change in position as these shifts are associated with declining success in transferring until the original readings are restored. We believe the top of the experiment, which is secured by spring fingers at the very top of the apparatus, may relax over time from its original position after an adjustment in position or an amplifier quench (the problem of amplifier quenching is discussed further in Section 8.2.3).

Chapter 8

Detecting and driving particles in a new apparatus

This chapter reports on the early commissioning of the trap. We have demonstrated trapping of antiprotons in the new loading trap, though we have not yet realized the hoped-for order of magnitude improvement in the loading rate. This chapter also presents the construction of new amplifiers for the experiment with a focus on improvements and new challenges encountered. The analysis, precision, and cooling trap amplifiers are working well enough for a measurement. However, there are challenges with the loading trap amplifiers that are limiting the progress of the experiment. We present an analysis of the difficulties we have encountered with detecting and counting small numbers of particles in the loading trap, and two proposals for resolving them.



Figure 8.1: Diagram showing two models of the how the cyclotron motion of the particle across a split compensation electrode induces currents in a tuned circuit: (a) series (b) parallel

8.1 Amplifier theory

The amplifier circuits used to detect the particle motion were a major focus of effort in the upgrade program. We present here the basic layout of an amplifier and define the main figure of merit, derive the parameters that affect it, and describe how we measure it. This provides context for the improved designs presented in the following sections.

We use an amplifier consisting of several components to detect the motions of the particle. The motion induces currents on the surface of the electrodes, which are converted into voltage by placing an inductive coil in parallel with the trap capacitance to form a tuned circuit with a large resistance. The coil is housed at the tripod inside a cylindrical shield that is grounded to the pinbase. One side of the coil is connected to the desired trap electrode, the other is grounded to the shield through a capacitor. The other trap electrodes are grounded through capacitors to the pinbase as well. The ungrounded end of the coil is connected in parallel to a field effect transistor (FET). The FET provides power gain and the output is matched to a 50 ohm cable to the hat.

The most important parameter of the tuned circuit is its effective resistance R. The detected signal voltage amplitude is proportional to this resistance through V = IR. In addition, the tuned circuit damps the particle's motion as described in Section 2.4, and the damping time constant is inversely proportional to R.

We can derive the parameters affecting R using a simple circuit model. In any practical circuit, there is some resistive loss r that limits the performance of the tuned circuit. Since the resistance of the wire in the inductor's coil is often the dominant source of loss, we usually model r in series with the coil inductance L as shown in figure 8.1(a). Around the resonant frequency ω_0 , we can transform the series impedance $\omega L + r$ to its equivalent parallel impedances [77]: $L_p \approx L$ and

$$R = \frac{(\omega_0 L)^2}{r} \tag{8.1}$$

where $\omega_0 = 1/\sqrt{LC_t}$ is the resonant frequency. This model is pictured in figure 8.1b. On resonance, L and C_t will cancel out and the impedance of the circuit is exactly R. When multiple sources of losses need to be considered, the resistance of each loss can be calculated and then transformed into a parallel resistance. The total resistance will then be limited by the smallest one or two contributions, which are often referred to as "loading down" the resonance. To achieve the best R at our target frequency, we want to minimize the loss r and maximize L by minimizing any extra capacitance in parallel with the trap.

We measure R by observing the noise spectrum of the amplifier. Johnson noise due to thermal motion of the electrons has a power proportional to the resistance [78]. The noise power spectrum of the amplifier is therefore proportional to the $R(\omega)$ of the tuned circuit. The resistance of this LCR circuit as a function of frequency has a Lorentzian shape with full width half max of $\Delta \omega$ (typically hundreds of Hz for axial amplifiers and a few hundred kHz for cyclotron amplifiers). The quality factor Q of the resonance is defined as

$$Q = \frac{\omega_0}{\Delta\omega} = \frac{\omega_0 L}{r} \tag{8.2}$$

This is a useful figure of merit since the resonant frequency and $\Delta \omega$ can both be directly measured from a noise spectrum scan. Additionally, from 8.1 and 8.2, they are both proportional to R:

$$R = Q\omega_0 L \tag{8.3}$$

. We typically measure L by measuring the resonant frequency of the coil as by changing a known capacitance. Therefore, a single measurement of Q and ω_0 from the noise spectrum allows us to calculate R.

8.2 Axial amplifiers

The new set of traps requires three axial amplifiers in order to keep the axial motion damped in each of the four traps. Two of the traps, the cooling and precision, share one axial amplifier. There is not enough physical space available in the tripod for separate amplifiers for each trap, and the disadvantage of an increased capacitive load is most acceptable in these traps because the main purpose of the axial amplifier in both traps to damp any heating of the axial motion rather than precise detection of the frequency. In contrast, we have a dedicated amplifier in the analysis trap because there it is critical to maximize the size of the axial signal as we are trying to detect frequency shifts of only 130 mHz (see Chapter 4 for more discussion of the challenges of analysis trap detection). We also use a dedicated amplifier to maximize the signal in the loading trap in order to compensate for the reduced signal there due to the



Figure 8.2: Photo of the new analysis axial amp with the top of the can open

large trap radius $\rho_0 = 1.2$ cm. Less space is available in the tripod region adjacent to the trap can because we must fit additional filter capacitors for new DC bias lines. The three new axial amplifiers together fit into the space previously occupied by the precision trap axial amplifier. In order to do this, we had to come up with a new design that was smaller in volume while maximizing the amplifier resistance and reliability. We present our design choices in the next section and Section 8.2.2 reports our resulting success of comparable resistance with the more compact design and presents a proposal for future improvement of the loading amplifier, while Section 8.2.3 reports our progress on solving a recurring issue with quenches of the superconducting materials.

8.2.1 Design and construction of the new axial amplifiers

The basic design of the new amplifiers is based on the design of the old amplifiers [35], but we also made a number of changes intended to reduce losses or improve reliability. We built the new amps to the same physical size as the previous small axial amp, which had achieved an R of 8 Mohm at 1 MHz and would be acceptable performance for trying a measurement. We use a rotated mounting post design to allow us to fit all three in between the legs of the tripod. A model of the new amplifier coils and mounting posts is shown in Figure 8.3. All three coils were intended to resonate with the trap around 1 MHz and so were built for an inductance of 700 uH. We used NbTi superconducting wire to eliminate resistance due to the large number of turns necessary to achieve that inductance, as was also done for the previous axial amps [35] and for similar amplifiers in our lab since [68]. Formvar-coated .0039" diameter wire was hand-wrapped around a custom-machined Teflon form in a single layer. One of the new amplifiers is shown in Figure 8.2 with the can top removed so that the coil is visible. For the new design, we changed the coil shape and position, the trap materials, the construction of the leads, and the design of the FET circuit board.

The new amplifiers were wound in a toroidal shape on a Teflon form. The previous large axial amplifier had been tested with both a toroid and solenoid, and the toroid had slightly higher resistance. This was attributed to lower magneto-resistance of a toroid due to its more confined field [35]. It was also hoped that a toroid would have a higher effective inductance due to less mutual inductance between the coil and the shield. Unfortunately, careful measurements of the coils after construction still

Amp	L (uH)	Test Q	Experiment Q	f (MHz)	R (MOhm)
Loading	570	15000	2000	1.069	5.6
Analysis	450	25000	8000	1.225	28
Precision	420	8000	4000	1.337	14

Table 8.1: Properties of the three axial amps

indicated an inductance about 30 percent lower than calculated for a free standing toroid with the same parameters. The inductance of the amplifier was estimated by measuring the resonant frequency with four different test capacitors and fitting to find the inductance and capacitance of the coil using

$$f_{res} = \frac{1}{2\pi\sqrt{L_{coil}\left(C_{coil} + C_{test}\right)}} \tag{8.4}$$

The calculated inductance of the amplifier used for the precision trap was only 420 uH \pm 30 uH, with 7 \pm 2 pF of coil capacitance. This turned out to be the smallest of the three amplifiers, as shown in Table 8.1. The 25 percent variation between inductors is likely explained by differences in turn density.

We also moved all the axial amplifiers into a region of relatively low field. Tests conducted at Harvard shortly before the new amplifiers were constructed indicated that magneto-resistive loss could have limited the analysis axial amplifier [58]. The new axial amplifiers were intended to sit at the position of the old precision amplifier, which was located further above magnet center than the old analysis amp. The solenoid field drops from 3.25 T to 1.26 T between those two positions, lowering magneto-resistive losses by 30 percent.

Additionally, we replaced the Macor electrode spacers in the analysis and precision trap with fused quartz spacers. We have known for a while that the Q of the amp



Figure 8.3: SolidWorks model of the new axial amplifier design, showing the coil support mounted on Teflon bottom support. A.Top view. B. Side view

when tested with a low loss capacitor is much higher than it is when the amp is mounted on the experiment. We suspected the main source of loss in the trap is dielectric losses in the Macor spacers, because at room temperature it is a relatively lossy material [79]. Tests at Harvard indicated that changing the material of the spacers could reduce the losses by as much as a factor of two [58].

We also changed the construction of the lead between the coil and the trap. The

new leads going to the experiment were constructed from 0.05" diameter OFHC copper wire, twisted together with Teflon tubing for insulation. The thicker, high purity wire was hoped to reduce loss relative to the 28 AWG wire leads that had been previously used, and using twisted pair was hoped to reduce noise pickup on the roughly 10 cm long leads needed to reach the pinbase from the new higher position of the amplifiers. Following similar logic, the grounding strap that connects the inductor's shield to the pinbase was attached using electron beam welding rather than spot welding.

We also redesigned the FET circuit boards. The new board design minimizes capacitive coupling to ground at the FET input, a suspected source of loss, by removing the copper grounding plane beneath it. The final board design is shown in figure 8.4). We ordered the circuit boards printed at a commercial shop rather than custom milling them. This was hoped to increase the reliability and interchangeability of the amplifier boards. For ease of manufacturing and durability, we switched from a Teflon-based weave material to a ceramic material (Rogers 4003b) that has similar dielectric properties but is much stiffer. This turned out to have a disadvantage because the larger thermal conductivity of the material made it more difficult to swap components in the matching networks and bias lines without overheating the FET (which has a maximum temperature of 175 C). In the end, we needed a combination of keeping the FET soaked with isopropanol, clamping a metal tool over the FET as a heat-sink, and simply replacing the FET as well when adjacent components needed replacement.



Figure 8.4: Layout of the new axial amplifier board, top view. Copper is shown in red.



Figure 8.5: Noise spectrum from the new analysis amplifier with a test capacitor.

8.2.2 Results and future steps

Both the precision and axial amplifier surpassed our goal of significantly higher resistance than the former analysis axial amplifier. The properties of the three axials are listed in Table 8.1. A noise scan of the analysis amplifier is shown in Figure 8.5. The analysis amplifier in particular has nearly four times as much resistance, which should translate directly to additional signal. However, the loading amplifier resistance was lower than hoped for and it was not possible to detect small numbers



Figure 8.6: Noise signal with a dip due to roughly 50 particles.

of antiprotons, which limited progress in all three traps. We report the loading amplifier performance, present a theory about our difficulty detecting particles, and make a proposal to improve its performance

The loading amplifier had difficulty detecting small numbers of particles, which is only partially explained by its low resistance. The resistance is about forty percent smaller than the former analysis trap amplifier, and in combination with the large trap radius, we expect a coupling constant that is only about half the size of the former precision trap. A lower coupling constant makes it difficult to detect the 'dip' signal from small numbers of particles (see Section 2.4). However, we were able to detect a signal from a single particle in the precision trap, so this alone does not explain our trouble seeing dips smaller than 50 particles (one of the smallest dips detected is shown in Figure 8.6).

This is likely because the lower than expected inductance forces us to use a trap voltage source with a higher maximum voltage than the ultra-low-noise DC biases used for the other traps. The trap voltage must be set to 22.7 V to bring the antiprotons into resonance with the tuned circuit, which is outside the maximum range of our highest stability voltage source (a new Stahl 1-14 with a maximum voltage of 14 V). Instead, we used 30 year old Fluke 5440B voltage calibrators to bias the loading ring and compensation electrodes. While we have used such supplies for many years, we have had increasing difficulty finding such supplies whose noise is still in specification as they have not been manufactured for decades.

Replacing the coil with a larger inductance would both increase the signal and allow us to use our new voltage source. Inductances of higher than 1 mH have been achieved in amps of this size with good Q by using multiple layers of windings [80] and in amplifiers on the next-generation apparatus at Northwestern by using thinner wire [58]. Increasing the inductance by a factor of two would increase the effective resistance by at least that much. It would also bring the tuned circuit frequency down to 735 kHz and the trapping voltage to 10.7 V, much less than the 14 V maximum voltage of a Stahl high stability voltage source.

8.2.3 Heat sinking

There is one other major problem with the amplifiers as currently constructed. In 2016, the superconducting coil of our precision amplifier started periodically quenching. When we designed the new amplifiers, we tried to prevent this by including several features that improved the heat-sinking of the superconducting material. However, the problem persists and the pattern of quenching with the new amplifiers may indicate the problem lies with the bore vacuum.

The signature of an amplifier quench is the sudden disappearance of the noise

resonance signal of the tuned circuit. At the same time, roughly 20 percent of the liquid helium in the experiment dewar rapidly boils off, accompanied by spikes in the pressure of the helium exhaust line and the bore vacuum. The amount of helium boiled off agrees within a factor of two with the calculated heat from the exothermic reaction of the superconducting shields going normal¹.

A major design goal for the new amplifiers was to make an extremely robust thermal connection between the tripod and the NbTi amplifier can. The old amplifier had a rectangular post attached directly to the cylindrical amplifier can with only two screws. The new post design has a flat circular base brazed onto the rectangular post. The base is bolted to the bottom of the amplifier can with six screws, increasing the surface area for heat transfer. We added a .005" thick layer of indium foil between the base and the can for even better heat transfer.

We also needed to heat sink the coil itself to the can. In previous toroidal amps, the coil inside the can had an alumina cylinder push-fit through its central bore for heat sinking. This resolved some initial issues with quenching before the 2012 measurement [35]. However, we were concerned by the high dielectric loss factor of alumina. We designed a pair of Teflon supports with a large surface area of contact to make up for the reduced thermal conductivity of Teflon. Unfortunately, in our tests, the outer support dramatically lowered the amplifier Q. We only included the inner support on the amps installed on the experiment.

$$Q = \frac{V_s H_c T}{4\pi} \left(\frac{\partial H_c}{\partial T}\right)_p \tag{8.5}$$

¹ The latent heat of transition for a superconductor going normal is given by [81]:

Given the critical field and slope of NbTi titanium[82], a quench should produce about 1200 J of heat, which is enough to boil off about half a liter out of our 3.6 L dewar.

The new amplifiers also suffer from regular quenches when installed on the experiment. While quenches in 2016 sometimes had no identifiable cause, quenches with the new amps are always correlated with some extra heat load. For example, a quench might occur after applying a strong drive or after applying a high bias voltage to the trap stack as we do for particle transfer or loading. Up until 2016, all of these had been standard procedures without any quenches. Reducing the voltages used during loading and transfer from 130 V to 50 V did lengthen the average time between quenches, but even at the lower voltage eventually a transfer attempt would cause a quench. In fact, the longer it had been since the last quench, the less heating was needed to trigger the next quench and the longest period between quenches was about a week.

Because of this, my hypothesis is that we may have had a small leak in the bore vacuum on our last cooldown. Over time, gas from the leak would cryogenically pump onto the electronics of the tripod region. Eventually, small local heating from a drive or a high voltage through a resistor would release cryosorbed gas, softening the vacuum and triggering a quench. This hypothesis explains the gradually increasing susceptibility to quenches that we observed. It would also explain why a next-generation apparatus at Northwestern was able to operate without quenching even though their amps used a closely related design with even fewer heat-sinking features [58]. The main evidence against a leak is that we performed several thorough leak checks of the bore and did not find it. While we could try adding heat sinking material to the coil, we may benefit more from searching for a leak from spaces that are difficult to check with the experiment in place.

8.3 Cyclotron amplifiers

We also made two new cyclotron amplifiers for the cooling and loading trap. The cooling amplifier is very similar to the former precision trap cyclotron amplifier [35] with several modifications intended to improve the Q and the frequency stability. The loading amplifier design is very different because the loading trap presents new challenges for sensitive radiofrequency (RF) circuits. We first describe the design and performance of the cooling amplifier, with a focus on the general considerations of a design at this frequency and the new ideas implemented. Next we describe the unique challenges of the loading trap, our proposal for meeting that challenge, the performance thus far, and a proposal for further improvement.

8.3.1 The cyclotron amp in the cooling trap

A cyclotron amplifier consists of a shielded inductor connected to a FET like an axial amplifier, but the much higher frequency has several consequences for the design. The 100 times higher frequency for the cyclotron motion means that we need a 10,000 times smaller inductance to resonate with the trap capacitance, which has three important consequences. The resistance of copper at 4 K is actually smaller than magneto-resistive losses in a NbTi coil of similar geometry [58]. We generally find that copper solenoid coils have a larger effective resistance and are easier to wind from thick low-loss copper wire. We also have an additional goal of maximizing the frequency stability. It is possible to adjust the particle cyclotron frequency through changing the current in the magnet, but afterwards we must wait several months for the magnet to settle in order to minimize drifts in the field [68]. Instead we adjust the coil inductance until the particle's frequency is within the roughly 150 kHz bandwidth of the tuned circuit. This is challenging because the frequency of the tuned circuit shifts by multiple linewidths on cooldown. We would like to both increase the tuned circuit resistance, and standardize the cooldown shift so that the amp is reliably centered at the correct frequency after cooldown. For the cooling amplifier, we tried several new ideas to increase the resistance and stability of the amplifier: we changed the spacer materials, we added a form to the amplifier, and we removed the varactor. These changes successfully doubled the resistance, but the frequency stability was not enhanced as much as we hoped.

The first change we made was fabricating the spacers of the cooling trap out of fused quartz instead of Macor like the precision trap spacers. As discussed in the previous section, the dielectric loss of Macor was the largest component of the loss on our axial amplifiers. We had also observed that loss on the experiment was about twice as large as the loss measured with an RF capacitor in place of the trap on a test stand. That indicated that our leading source of loss was not in the amplifier itself, and we suspected the spacers were limiting the Q for the cyclotron amplifiers as well.

Next, we added a thin Teflon form supporting the coil. The main motivation was to improve the frequency stability of the amplifier. Previous versions of the precision amp were constructed by winding the coil around a form, then twisting the coil off and mounting it free-hanging with just a teflon spacer around the bottom lead to support it. The thick copper wire used is stiff enough to retain its shape, and it originally was thought better to avoid unnecessary dielectric material which is a potential source of loss. However, the higher Q of the new cooling amp indicates that the Teflon form is not a leading source of loss. In fact, it may prevent losses from uneven current distributions by keeping the coil centered and evenly spaced. This was the only major change we made that could explain why the Q of the cooling cyclotron amplifier was higher on the test stand as well as the experiment.

Finally, we did not include a varactor in parallel with the coil. Varactors are diodes whose capacitance changes based on the reverse bias voltage applied. We had used a such a diode to add 0.6-0.8 pF of capacitance to the circuit, allowing us to tune the frequency by 150 kHz. This allowed us to adjust the amplifier frequency in situ. Varactor diodes did not reduce the Q of the amplifier on the test stand. However, we did observe that the Q of the old precision amplifier on the experiment depended on the varactor bias voltage, which was attributed to magnetoresistance [35]. We expected that the new cyclotron form would make the cooldown shift sufficiently repeatable that the varactor was no longer needed, so we removed the varactor to avoid these losses.

The average size of the cooldown shift has been reduced, but the variation from cooldown to cooldown is still larger than the amplifier bandwidth. While the precision cyclotron amplifier shifted on average each cooldown by roughly 1.2 MHz, the cooling cyclotron amplifier only shifted by about 700 kHz. However, over a half-dozen cooldowns where I targeted the same frequency, the actual frequency at 4 K varied from 86.4 to 87 Mhz. That 600 khz range is four times larger than the amplifier bandwidth, making it likely that the effective resistance at the particle frequency was significantly reduced on some cooldowns. However, we should be able to compensate for this using the gradient in the cooling trap. We can shift the particle up and down by asymmetrically biasing the trap endcaps, which changes the average magnetic field seen due to the significant linear gradient present in cooling trap. Changing the particle position by 100 μ m would already shift the frequency by 20 kHz.

However, the effective resistance of the amplifier has improved by nearly a factor of two. The Q of the cooling amp with a test capacitor was 3000, while the previous cyclotron amplifier had a Q of 1200 last time it was tested on the test stand. As the inductance and resonant frequency is similar, the Q is directly proportional to the resistance. The Q on the experiment improved by a similar factor, going from 700 to roughly 1500. The fact that the Q dropped by a factor two going from the test stand to the experiment with both the old precision amplifier and trap and the new cooling amplifier and trap may suggest that losses from the magnetic field are responsible for the lower Qs on the experiment. If the leading source of loss is magnetoresistance in the coil, they should scale in a consistent way as each coil is moved into the same magnetic field. Alternatively, the relationship is a coincidence and we are still limited by loss in the spacers. Measurements at 5 K and 10 GHz indicate a loss tangent consistent with a Q of 1900 for our resonator [83], though dielectric losses are often frequency dependent and thus may be very different at 86 MHz. In any case, together with the design improvements of the cooling trap discussed in Chapter 6 we should have a cooling trap cyclotron damping time of only 3 seconds, which is more than sufficient.



Figure 8.7: Photograph of the loading cyclotron amp mounted inside the trap can next to the loading trap

8.3.2 Loading cyclotron amplifier

The loading amplifier design has the additional challenge that the distance between the trap and the tripod is very large, and the long leads needed add capacitance and loss. In addition, the loading trap itself has about twice the capacitance of the other traps in the stack. The tuned circuit needs a much smaller inductance relative to the old precision amplifier to resonate at the cyclotron frequency. Instead we constructed a smaller loading cyclotron amplifier intended to be mounted next to the trap itself. This allowed us to use a coil with inductance comparable to the old precision cyclotron amplifier. The rest of this section describes the motivation for moving the amplifier inside the trap can: to greatly reduce effects from the leads in order to compensate for an unexpectedly large trap capacitance. Section 8.3.3 discusses the new design inside the trap can, and Section 8.3.4 discusses its performance.

The first problem is that the parasitic inductance of the long leads affects the behavior of the tuned circuit at this frequency. We typically connect the electrodes of the trap to the pinbase using copper straps. The flat shape maximizes the surface area of the signal, minimizing the inductance per unit length. Even so, the 10 cm strap used for the old precision amplifier added about 60 nH of inductance. This inductance does not increase the signal across the coil, but it adds to the inductance of the coil as it resonates with the trap capacitance. A strap long enough to reach the new loading trap would be 15 cm and an inductance of at least 100 nH, requiring 50 percent less inductance in the coil. We considered using a coaxial lead that would have much less inductance than a strap because much of it is cancelled out by the mutual inductance of the inner and outer conductor. However, the effect of the extra capacitance between the inner and outer conductor more than compensated for the reduced inductance, even for a custom coaxial lead constructed from 6 mm copper tubing and teflon spacers supporting a copper inner conductor. The total effect of the extra distance to the loading trap is a 33 percent reduction in the amplifier performance.

We also found that the loading trap itself had about twice the expected interelectrode capacitance. We had decided to keep the gap between electrodes in the loading trap as small as possible to make the electrostatic trapping potentials as accurate as possible. The same spacer thickness was used as in the precision trap in order to keep the mechanical fit to the electrode equally tight in the new trap. An estimation of the expected capacitance based on approximating the flat surfaces as parallel plate capacitors shows that the loading trap should have about 13 pF of trap capacitance, nearly double the 7 pF estimated capacitance of the old precision trap. That agrees with the experimentally measured value of 16 pF.

We would not be able to damp or detect particles effectively with our initial design

due to these extra sources of capacitance, combined with the larger trap radius. The damping and signal is proportional to R/ρ_0^2 as described in Chapter 6. We would need a coil inductance and thus R at least three times smaller than the old precision trap, so given that the trap radius is two times larger the signal from the particle would be twelve times smaller. Most importantly for controlling the particles, the cyclotron damping rate would increase by the same factor to 36 minutes. We feared it could be even longer in practice as extra losses in the longer lead might further reduce the resistance. With such a long damping time, cooling excited particles after loading or transfer would take an impractical amount of time.

8.3.3 Loading cyclotron amplifier design

We decided instead to mount the amplifier coil inside the trap can to eliminate the contribution from the long leads. The new design has a larger inductance coil, a FET located separately from the coil outside the can, and a varactor diode to tune the resonant frequency.

The new coil has a higher inductance and a smaller diameter. Eliminating the lead inductance allowed us increase the coil inductance by a factor of three to 150 nH. We needed to make the diameter of the coil and its shield much smaller so that they could fit next to the loading trap. The amplifier mounted inside the trap can is pictured in Figure 8.7. The coil is connected across and resonates with the split halves of the loading bottom compensation electrode, as shown in Figure 8.8. Overall, we expected this tuned circuit would have about a 12 minute damping time.

The FET remains outside the trap can for several reasons, with a homemade rigid



Figure 8.8: Wiring diagram of the loading trap resonator and FET

coaxial cable carrying the signal from the coil to the FET. First, the FET package is unavoidably magnetic and we wanted to maximize its distance from the precision trap. Additionally, the FET can be easily damaged by errors in biasing, and its circuit board has an impedance matching circuit on the output that sometimes needs adjustment. Keeping those in the relatively accessible tripod region minimizes the number of times the trap can would need to be opened for repairs. There are also no free electrical feedthroughs on the pinbase to pass the FET bias into the trap can. We expected that only about 5 percent of the signal would reach the FET through the long coax. However, since we usually detect peaks from excited particles we intended to overcome a lack of signal with a stronger drive to larger radii. In any case, it was considered more important to maximize the front-end performance so particles could be cooled even if they could not be detected.

A varactor diode was placed in parallel with the coil inside the trap can to allow adjustment of the frequency without opening the trap vacuum space. Its capacitance can be tuned from 1.2 to 1.7 pF by changing its bias voltage from 0 to 9V, which alters the frequency by about 600 kHz.



Figure 8.9: Noise signal from loading cyclotron amp, with noisy temperature sensors that are usually unplugged left connected to drive resonance. Blue is with varactor biased to 0V, yellow with the varactor biased to +9 V

8.3.4 Loading cyclotron amplifier performance

The loading cyclotron amplifier was not able to detect particles. In addition to a low signal to noise, the resonant frequency varied from cooldown to cooldown by more than we could compensate with the varactor.

The signal amplitude coming out of the amplifier was at least a factor of two lower than we had predicted. Our best estimate for the damping time with the loading cyclotron amplifier as installed on the experiment is 22 minutes. We were not able to verify that, as we never detected cyclotron signals. The lack of cyclotron signals made it much more difficult to get to small numbers of particles in the loading trap and slowed progress in all traps. The problem is likely due to a combination of complications from the loading gradients, the low R of the tuned circuit, and the very small signal making it to the FET outside the trap can. The cooldown shift of the loading cyclotron amp was too variable to fully compensate for it with the varactor. The cooldown shift was typically twice the shift in the old precision trap, and instead of reaching the final frequency within an hour or so of the first helium fill it took half a day to reach the final frequency. The long time to cool down is probably due to poorer heat-sinking inside the trap can. The larger cooldown shift is probably explained by the different geometry of the coil, which had more turns, less separation from the grounded shield, and lacked a Teflon form due to time constraints. Even with a form, a varactor is likely essential for tuning the circuit to the correct frequency without needing to open the trap can.

8.3.5 Future steps

To remedy the issues of the last section, we propose two possible approaches that require modifying the trap stack but may be necessary to see cyclotron signals. The first is to rearrange the existing traps in the stack so that the precision trap is at the center of the stack just above the loading trap. This would reduce the gradients in the loading trap by moving it closer to the center of the solenoid field. Section 8.4 makes the case that driving and detecting cyclotron signals is much more difficult in the presence of these gradients, and eliminating the gradient may be enough by itself to let us see signals from highly excited particles. If not, we could add a cyclotron amplifier back to the precision trap and use that for counting individual particles after transfer from the loading trap. In that case, we would only need the loading cyclotron amplifier for cooling particles enough after loading that they can be safely stored, which has already been demonstrated with the current amplifier. There is a technical problem of finding space for a third amplifier in the tripod, but otherwise this would be a simple reversion to a precision trap amplifier design that has effectively counted particles on the old experiment.

Alternatively, we could modify the split electrode of the loading trap and the adjacent spacers to reduce the trap capacitance, allowing us to move the amplifier back outside the trap can. This would address the three most likely sources of extra loss in the system, and also increase the coupling signal from the coil out to the FET by an order of magnitude. First, reducing the trap capacitance reduces the effect of any loss in the electrodes. Second, loss in the coil is inversely proportional to its physical volume [84], and outside the trap can there is space for a coil that is twice the size. Finally, the coil would be more accessible outside the trap can, making it feasible to reduce the coupling of the varactor diode which has been suspected of magneto-resistive losses [35].

The proposed trap modification would affect the split compensation electrode and the two spacers on either side. The proposed trap is shown side by side with the current trap in Figure 8.10. To reduce trap capacitance, the new spacers should be twice as thick. The compensation electrode would also need to be modified with a deeper lip for the spacer to sit on, so that the larger spacers are still mostly screened from the particle. However, we would increase the gap between the inner surface of the bottom compensation electrode and the ring from 0.0035" to 0.007". This should also reduce the capacitive coupling through the gap. A 0.007" gap was used in a trap of this radius for the charge to mass ratio measurement [68], so any imperfections of the potential this would cause are presumably tolerable. We could also replace the sap-


Figure 8.10: Proposed revision of the loading trap for better RF performance

Modification	Current C (pF)	Proposed C (pF)
Across split	1.5	.75
Through spacing balls	1.6	.6
Through spacer	9.2	3
Increasing gap to electrodes above and below	4.3	2.2
Total	13	6.5

Table 8.2: Trap capacitance sources in current and proposed revised trap

phire balls used for spacing the halves of the split with balls made from fused quartz, which has a lower dielectric constant. Finally, the width of the cut separating the two halves should be doubled. All these modifications could reduce the capacitance by a factor of two, as shown in Table 8.2. The lower trap capacitance would allow us to place a 100 nH coil outside the trap can, which is not predicted to improve the damping time but would considerably simplify troubleshooting, repair, and unwanted pickup of the drive signal.



Figure 8.11: Cyclotron peaks from antiprotons during the ppm measurement. From left to right, these show signals from (a) Dozens of antiprotons (b) 4 antiprotons (c) one antiproton

8.4 Sweeping the cyclotron drive in a magnetic gradient

We use the cyclotron signal to count the number of particles in the trap. This is an essential step for a measurement as we must ensure we have only a single particle for the most accurate measurement [35, 36]. We will review how the cyclotron signal distinguishes individual particles and the usual method of detecting the cyclotron signal before introducing following sections that discuss the difficulties of detecting this signal in the presence of strong magnetic gradients.

The cyclotron signal is particularly suitable for this task because each particle develops a different relativistic frequency shift depending on their individual energy. This shift is due to the relativistic increase in the particle mass by the Lorentz factor

$$\gamma = \frac{1}{\sqrt{1 - v^2/c^2}}$$
(8.6)

The cyclotron frequency is inversely proportional to the mass, so higher energy par-

ticles have a lower cyclotron frequency. The fractional frequency shift is proportional to the energy of the particle:

$$\frac{\Delta\omega_c'}{\omega_c'} = -\frac{E_c}{E_c + m_p c^2} \approx -\frac{E_c}{m_p c^2} \tag{8.7}$$

Even though v/c < .001 for a trapped particle, we are able to distinguish this shift as the peak from the excited particle has an extremely narrow width (typically 5 mHz in the former precision trap and even narrower in the new loading trap).

To excite the particles, we sweep a drive with decreasing frequency so that the drive comes into resonance multiple times as the particle's energy increases. After turning off the drive, an excited response can be immediately detected on the cyclotron amplifier, as shown in Figure 8.11 from the former precision trap. With multiple protons in the trap, they often end up at different energies causing a peak to be visible for each particle. If a single peak is seen, we repeat the sweep several times to make sure it is not two particles that are "locked" into a similar orbit. We can see a distinctive exponential decay of the particle energy and thus frequency in the minutes after the drive, which damps into the cyclotron amplifier on a timescale of minutes (Section 2.4).

However, this does not account for the effect of a magnetic gradient on the cyclotron frequency and the response of a particle to a swept drive. First, the magnetic gradients greatly increase the parameter space that must be searched to find the initial response. Gradients can also alter the shape of the exponential decay which is usually the distinctive sign of a true cyclotron signal. Finally, they may interfere with the coherent response to the swept drive preventing us from exciting the particles in the first place. Taken together, these effects explain why initial attempts at using the loading and cyclotron trap to count particles were unsuccessful. The final section presents a modification to the trap stack that should eliminate these problems, allowing us to quickly verify we have a single particle.

8.5 Frequency uncertainty due to magnetic gradients

The magnetic gradients in our cooling and loading traps are both much smaller than the gradients of the analysis trap and include a linear gradient that causes a much wider range of possible cyclotron frequencies. The loading trap gradients are due to its position near the edge of the main solenoid magnet, where the field is beginning to fringe outward. They cause an 800 khz span of possible frequencies that must be searched to find the resonance. The cooling trap gradient is due to its position next to the iron ring of the analysis trap, but the lower positional uncertainty results in lower frequency uncertainty.

The magnitude of the solenoid field is changing rapidly at the position of the loading trap. The solenoid field is only relatively flat over a 10 cm range around the the center of the magnet, and at the position of the loading trap there are both large linear and quadratic gradients. This is larger than would be expected from the length of the solenoid coil, because the magnet was designed to maximize the homogeneity of the field near its center by using two coils with precisely defined geometries. We used a Hall sensor magnetometer to measure the magnetic field over a 20 cm range along the central axis of the magnet, shown in Figure 8.12. The NMR

method used for shimming [68] is more precise, but the field is too inhomogeneous for a detectable NMR response more than ± 2 cm from magnet center. The loading trap is roughly 11.6 cm below the center of the magnetic field, where the field is starting to fringe outwards and the magnitude is declining. This is due to its position at the opposite end of the trap stack from the precision trap which must be placed at trap center for the most homogeneous and stable field (see Chapter 7). The exact distance between the precision trap and field center is only known to ± 0.5 cm due to machining tolerances. As the field magnitude varies rapidly in this region, this range of positions corresponds to a frequency between 84.298 and 85.080 MHz.

As a result, it takes an extremely long time to search this 800 kHz span for the cyclotron frequency. The width of the particle response is broadened by the gradient as well as the damping to the amplifier, but it is still less than a Hz in both traps. The drive sweep must cover steps of no more than tens of Hz at a time to have a chance of resonating with the particle. In the old trap, a typical sweep would take a 10 Hz step every 0.5 s. We also need to pause regularly to check for a peak from the particle in the amplifier signal, typically every 1-5 kHz.

The magnetic gradient in the cooling trap is dominated by the effect of the iron analysis ring that is only 6.8 mm away. The cooling trap is only about 5 cm away from the precision trap, so magnetic field gradients due to main magnet coil are negligible The measured gradient of the loading trap and the calculated gradient in the cooling trap are compared in Table 8.3. The calculated cyclotron frequency around the cooling trap as a function of position was shown in Figure 6.6.

The cooling trap gradients are significantly larger than those of the loading trap,



Figure 8.12: Estimated cyclotron frequency as a function of position, based on measured magnetic field strength along the central axis of the magnet bore. The position of the loading trap is shown by the green shaded region, which includes uncertainties due to machining tolerances.

but the tight machining tolerances on the intervening electrodes and spacers mean we only need to search a 100 kHz frequency range corresponding to the \pm 100 μ m positional uncertainty. This is time consuming, but likely the main problems in the cooling trap are the other effects of the gradients covered in the next few sections.

8.5.1 Gradient effects on decay shape

The magnetic field gradient changes the optimum drive sweep parameters. In a gradient, the average magnetic field also changes as the particle's radius increases. We discussed these shifts briefly in Section 6.2.3, where we found the shift for a particle of radius ρ and amplitude A in Equation 6.10. Assuming that $A \ll \rho$, we can write

Trap	$\mathbf{B}_1(T/m)$	$B_2(T/m^2)$	linewidth from $B_2(Hz)$
Old precision	.022	1	0.9
Cooling	12	1360	22.7
Loading	-5	-75	5.8

Table 8.3: Gradients in the loading and cooling traps

the shift as

$$\Delta f_c = \frac{q}{m} \left(\sqrt{\left(B_0 - \frac{B_2}{2} \rho^2 \right)^2 + \left(B_1 2 \rho \right)^2} - B_0 \right)$$
(8.8)

This depends on ρ^2 to first order, which is proportional to the energy just like the relativistic shift, but the constant scaling factor is different. This swamps the effect due to relativity, changing the optimal sweep span and introducing uncertainty into the parameters. We calculate the effects in the loading and cooling traps before summarizing the effect on our measurement.

In the loading trap, the gradient shift is on the same order as the relativistic shift, but the sign is opposite so the frequency actually increases with radius. This is expected because the magnitude of a solenoid field would increase as you move slightly away from the center as the field lines should be bending outwards to circle back around the outside of the solenoid. The frequency shift as a function of radius is shown in Figure 8.14, with shaded blue indicating uncertainty due to error in the measurement of the gradients. For comparison, the shift due to relativistic effects at 1 mm is -139 Hz, and the shift due to the gradients is about +600 Hz at 1 mm.

For the cooling trap, the frequency still decreases as the radius increases, but the size of the effect is much larger than for the relativistic shift. This means we need to cover larger ranges with each drive sweep in order to reach a large radius. The



Figure 8.13: Projected change of cyclotron frequency with radius due to cooling trap gradients. Red is the average frequency, and the shaded blue area is the uncertainty due to machining tolerances

frequency shift as a function of radius for a particle at the center of the cooling trap is shown in Figure 8.13.

While this general relationship and sweep span can be calculated, fine details like the best step size, drive strength, and time per step are usual found through trial and error. With the sweep characteristics changing, they may need to be re-optimized before we can see a strong response. This adds even more time to the initial search for a cyclotron peak.

8.5.2 Linewidth broadening and cyclotron sweeps

Magnetic gradients also change the cyclotron transition lineshape, which may make the cyclotron drive less effective. We review the change to the cyclotron line-



Figure 8.14: Projected change of cyclotron frequency with radius due to loading trap gradients. Red is the average frequency, and the shaded blue area is the uncertainty due to machining tolerances

shape, demonstrate the effect on the signal size, and then estimate the size of this effect.

The second-order gradient causes a broadening of the cyclotron lineshape, as discussed in Section 2.8. Johnson noise from the axial amplifier drives random fluctuations in the amplitude of the axial motion, and because of magnetic gradients, the average field seen by the particle is changing on the timescale of the axial damping time. The average axial amplitude and standard deviation are both given by [37]:

$$\left\langle z^2 \right\rangle = \frac{k_B T}{m\omega_z^2} \tag{8.9}$$

and the resulting standard deviation in the transition frequency is

$$\Delta\omega = \omega_0 \frac{B_2}{B_0} \left\langle z^2 \right\rangle \tag{8.10}$$

Trap	B2 (T/m^2)	$z_{rms}~(\mu m)$	linewidth from B2 (Hz)
Old precision	8	71	0.6
Loading	-75	38	2
Cooling	1360	33	23

Table 8.4: Linewidths in the loading and cooling traps

This standard deviation is referred to as the linewidth parameter $\Delta \omega$. The gradients and resulting linewidths are shown in Table 8.4. The much larger gradients in the new traps are partially compensated by the smaller axial amplitudes due to the choice of larger axial frequencies. Assuming we have kept the axial temperatures to 8 K, the loading trap linewidth has increased by a factor of three and the cooling trap linewidth is 23 times larger. When $\Delta \omega >> \gamma$, the lineshape is approximately [37]

$$\chi(\omega) = \frac{\Theta(\omega - \omega_0)}{\Delta\omega} \exp\left[-\frac{\omega - \omega_0}{\Delta\omega}\right]$$
(8.11)

which is a valid approximation for the old precision trap as well as both the new cooling and loading traps.

The decrease in the height of a peak implies a possible decrease in the final amplitude after the drive which could be a problem because we are limited in the drive strength we can use. As described in Section 8.2.3, strong drives can cause local heating that leads to quenches, and we have avoided increasing the drive by more than a factor of 20 dB because beyond that we have seen more quenches. We want to estimate with slightly more accuracy to see if the effect of the drive has decreased by more than we can compensate with 10 times more drive voltage.

To estimate the effective attenuation of the drive, we start from the steady state solution for the damped driven anharmonic oscillator [37] and examine its behavior. The relativistic motion contributes an amplitude dependent term to the frequency, resulting in a multi-valued response with an upper and a lower branch illustrated in 8.15, where $w = |u^2| / |u_{max}|^2$, u is the relativistic velocity in the rotating frame, and u_{max} is the maximum velocty given by $|u_{max}|^2 = (2e/m\gamma_c)^2 |E_0|^2$. The same equations can be used for the relativistic bottle and a magnetic gradient as long as we replace the term proportional to u^2/c^2 with B_2u^2 , so we have a similar response curve. The maximum velocity u_{max} does not change since it is determined by the equilibrium between the drive and the damping force of the amplifier, but the slope of the curve can get steeper, shallower, or reverse direction. For a given amplitude, the distance between the two branches is

$$\omega_1 - \omega_2 = \gamma \sqrt{\frac{1}{w} - 1} \tag{8.12}$$

This distance is independent of N and thus the source of the anharmonicity. As the amplitude increases, the distance between the branches gets smaller until they converge to a single solution at the peak, where w = 1.

When the amplitude of the particle becomes large enough that the distance between the two branches is less than the linewidth, we assume it will tend to stop following the drive sweep. Such behavior was observed with trapped electrons when dephasing of the particle was due to drive noise rather than magnetic field noise [85].

We can solve for the fractional amplitude where the distance between the branches is equal to the linewidth. This occurs at

$$w_{max} = \frac{1}{1 + \left(\frac{\Delta\omega}{\gamma}\right)^2} \tag{8.13}$$

$$|u_{final}|^2 = w_{max} |u_{max}|^2 = \frac{1}{\gamma^2 + \Delta\omega^2} \left(\frac{2eE_0}{m}\right)^2$$
 (8.14)



Figure 8.15: Steady state response of a particle to a cyclotron drive due to magnetic gradients or relativistic shift

This would indicate that the final amplitude after the drive is limited by the larger of the cyclotron amp coupling and the bottle linewidth. But in all three traps, $\Delta \omega \gg \gamma$, so the contribution from the cyclotron amp can be ignored. In that case, the largest radius achievable from a sweep at a given drive power should be smaller in the new traps by a factor of $\Delta \omega$.

In principle we should have been able to overcome this with stronger drives in the loading trap, but the factor of 23 in the cooling trap would significantly reduce the final radius. If the axial motion is hotter than the 8 K assumed here, the linewidth would be even broader.

8.5.3 Next steps for cyclotron detection

The last section presented three ways in which the cyclotron motion is more difficult to excite or detect in the presence of a strong magnetic gradient like the ones in the cooling and loading traps. None of these effects clearly rule out seeing a driven response, but they each make it more difficult. Given that we have not been successful in finding a resonance with our initial searches, a different approach should be used going forward. We propose a method to eliminate the gradients in the loading trap so we can use it for particle counting, and an alternate technique for performing the most important cyclotron frequency measurements in the cooling trap.

For fast particle counting in the loading trap, it is easiest to rearrange the stack so that the loading trap is no longer in a region of strong magnetic gradients. A new apparatus is being tested at Northwestern with the precision trap in the center of the stack, so that the loading trap is only 5 cm from field center. We will probably also need a more sensitive cyclotron amp in the loading trap to see driven responses there, as discussed in Section 8.3. This stack arrangement would let us revert to using the precision trap to count particles if we still have difficulty detecting the signal in the loading trap.

For the cooling trap, we cannot avoid magnetic gradients but we propose detecting the cyclotron frequency using the analysis trap instead. The magnetic gradients are inevitable since the cooling trap must be located adjacent to the analysis trap for easy transfers. However, in the cooling trap our main consideration is finding the cyclotron frequency so we can make sure that the tuned circuit is effectively damping the particle motion. We can test this directly by alternating alternate driving in there with transferring the particle back to the analysis trap to check if the cyclotron radius has increased. If we suspect that we can't efficiently drive particles in the cooling trap due to the large gradients, we could instead excite the particle in the analysis trap. Then we could move the particle to the cooling trap, wait for several damping times, then move back to the analysis trap to check that the cyclotron radius has damped. This will verify that the cyclotron frequency is at least within the bandwidth of the amplifier, while simultaneously measuring the damping time. If the damping time is significantly extended because of mistuning, we can try to vary the cyclotron frequency of the particle by biasing the endcaps of the cooling trap to shift the particle's vertical position.

8.6 Degrader and loading

The loading trap was added to the experiment to allow us to work with particles during the relatively quiet periods of the long shutdown when the antiproton beam is not available. A secondary goal was to increase the antiproton loading rate by an order of magnitude. While in the end we will do the measurement with a single particle, the large signals from thousands of trapped antiprotons could greatly speed up early commissioning and troubleshooting. In practice, loading with the new trap is only slightly improved over the old trap, which indicates we must look to other parts of the experiment for the desired improvement.

The loading trap was designed with a larger radius because our experiences with loading in 2012 indicated we were limited by the small radius of the precision trap. Section 3.1 describes the loading procedure in the old trap, which was at most able to

catch a hundred antiprotons per shot through the usual procedure of rapidly switching the voltage at the trap entrance after the beam has entered the stack. Early generations of antihydrogen loading experiments had successfully caught thousands of antiprotons using this basic technique, but in a trap with twice the diameter [86]. Even in those experiments, the number of antiprotons caught was found to increase with the depth of the initial trapping well only up to 830 V or so before levelling off. This effect was blamed on high energy particles hitting the trap walls. If particles are assumed to have roughly equal axial and radial energy, particles with larger axial energies would also have larger radii. As the trap radius increases particles of higher and higher energy can be caught. At some point higher energy particles have such large radii that they hit the walls, and the number caught would stop increasing [86]. In our experiment in 2012, we saw a similar pattern (see 3.1), but with the turning point around 300 V. Since we expect this effect to scale as the radius squared, this is the right relationship for a 6 mm diameter trap [36]. Based on this evidence, we assumed that designing the loading trap with a 12 mm radius would restore the higher loading rates of the earlier experiments.

We also changed the design to minimize the gap between the degrader and the extension electrode, as we suspected that the previous several millimeter gap may have been interfering with loading. The electric fields are not well-defined in the intervening space and particles might have gone astray. The gap was designed to ensure that they could not interfere with each other despite the uncertainty on their relative positions. The old degrader was mounted to a separate support structure attached to the bottom of the can, and the relative spacing between it and the bottom



Figure 8.16: SolidWorks model of the degrader of the new loading trap and degrader assembly mounted inside the trap can

of the trap stack was set by two indium seals. The new loading trap was designed with the degrader mounted directly to the bottom of the trap stack. A SolidWorks model of the new trap and degrader is shown in Figure 8.16. Macor spacers electrically isolate the degrader mount from the endplate holding the rest of the stack in place. Following the rule of thumb of 500 V volts per 0.001" of separation in vacuum [47], and allowing for machining tolerances and shrinking on cooldown, the 0.005" gap from the degrader to the electrode stack should allow us to safely put at least 1000 V on the degrader.

Loading in the new trap has been similar to loading in the old trap despite the change. The "turning point" is even lower in voltage now, as the number caught at 200 V is the same as at 400 V. As we were also trying minimize heat loads on the trap for reasons explained in Section 8.2.3, after this was established we generally defaulted to using only 200 V. In December 2017 we typically caught about 100 antiprotons



Figure 8.17: The number of antiprotons loaded as a function of SF6 fraction in the energy tuning cell

per shot. In the next year's beam run starting May 2018, that had dropped to 50 or less for reasons that are still not explained. Even at these low voltages, if we had been able to load with similar efficiency as the antihydrogen experiment we would have caught several thousand antiprotons.

Unlike the old precision trap, we are able to stack multiple shots because we are now able to load directly into the main trap 3.1. After the particles are caught in the high voltage trap by the switching voltage, we sympathetically cool the antiprotons with electrons into a shallow well on the loading trap. After 200 seconds of cooling, we open the potential at the bottom of the stack again so the next shot can enter the trap, and we repeat the process. We were able catch 3-10 shots in a row and stack them until we had dozens of antiprotons trapped in the loading trap. Several dozen stored antiprotons would be adequate to perform a measurement with a wellcharacterized experiment. After this was achieved, we decided to focus on cooling antiprotons and using protons to finish commissioning the trap as the most efficient way to use limited beam time.



Figure 8.18: Antiprotons caught as a function of time to close the trap

However, we can conclude that the smaller trap radius was not responsible for low antiproton loading rates during the ppm measurement. We also replaced the degrader and most of the cold wiring when the new trap was installed, making it unlikely that any of these elements are the cause. Instead, we should investigate the trapcan bottom, radiation shield, PPAC, and bore vacuum as these elements were not substantially modified and could be potentially blocking the beam. However, they will all need to be replaced and likely significantly redesigned before the next beam run anyway because of a planned upgrade to the AD. A second deceleration ring will further cool antiprotons to 100 keV before they are sent to experiments, which should improve trapping efficiencies. However, this adds an experimental challenge because such a low energy beam will be completely stopped by the 10 μ m thick vacuum windows currently used to separate our bore vacuum from the beamline vacuum and other spaces in the beam path. A new interface to the beamline vacuum will be needed, likely requiring a complete redesign of the beam position detector and the energy tuning system. Great care should be taken to ensure the entire beam can still pass through the new components.

8.7 Conclusion

This chapter reports the detection of particles particles in the new apparatus. It presents the construction of five new amplifiers for detecting particles signals in all traps, with improved performance relative to their volume. Our experience shows that a Teflon support coil doubled the performance of a high frequency cyclotron amplifier. The loading trap amplifiers are not as sensitive as we would like, but they are most likely limited by the high capacitance and loss of the loading trap. Two plans to adapt the trap electrodes for better performance are presented. We also present an argument that large magnetic gradients are preventing the detection of cyclotron peaks (our usual detection method for counting small numbers of particles), and plans for resolving them. Finally, we report the successful loading of antiprotons in the new loading trap.

Chapter 9

Conclusion

This thesis reports two milestones in antiproton measurements [1, 2], and then presents progress in method and apparatus for a higher precision measurement. The first milestone is the ppm measurement of the antiproton magnetic moment [1]. The second was the demonstration of single proton spin flip detection [2]. To make a ppb measurement feasible, we designed and implemented an upgrade program intended to reduce electromagnetic noise in the analysis trap and reduce the time needed to cool the cyclotron motion. These improvements open the possibility of using a new measurement technique based on separated oscillatory fields to achieve subppb precision. The new trap and supporting equipment has been constructed and commissioning began during the last antiproton beam run.

Our measurement of the antiproton magnetic moment was the first to use a single particle [1], representing a 680-fold improvement in precision over previous measurements [14]. The antiproton g-factor is

$$-\frac{g_{\overline{p}}}{2} = -2.792845(12) \left[4.4\text{ppm}\right]. \tag{9.1}$$

This allowed the first direct comparison of antiproton and proton moments, using the same apparatus and method for both particles, yielding the ratio

$$\frac{\mu_{\overline{p}}}{\mu_p} = -1.000000 \pm 0.000005. \tag{9.2}$$

This comparison agrees with the Standard Model prediction.

The g-factor measurement relied on measurement of the spin and cyclotron frequencies of a single antiproton in our Penning trap. After proton measurements were carried out, the trap had to be adapted and transported to the Antiproton Decelerator facility at CERN and then commissioned for antiproton loading. Through repeated trials at several frequencies, we found the resonant frequency for the spin and cyclotron motion by detecting the increased Allan deviation of the axial frequency due the transitions caused by the drive. The ratio of the two resonant frequencies is equal to g/2.

Our observation of individual spin flips [2], is a key step towards a higher precision measurement. The demonstration dataset showed an axial frequency stability of 60 mHz between points, making the 130 mHz shift due to individual spin flips detectable with 96 percent fidelity and 25 percent efficiency. This allows the use of double trap methods where the frequency is measured in a separate, low-gradient trap and the high gradient analysis trap is only used for spin state is detection. We proposed to increase the efficiency to near unity using adiabatic fast passage.

Experience gained during the accomplishments of the previous sections told us an upgrade was needed for a part-per-billion measurement to be feasible. Upgraded methods and apparatus were were designed and partially implemented. Our approach was to simultaneously reduce the cyclotron cooling time with a new dedicated trap and increase the acceptance of the analysis trap with a modified design. The coupling of the new analysis trap to the thermal reservoir has been enhanced by a factor of 72. The new design of the analysis trap shields the particle from electromagnetic noise that could induce cyclotron transitions, increasing the stability by as much as a factor of 10.

The new electrodes and supporting equipment have been constructed and initial commissioning was performed during the 2018 antiproton beam run. In addition to the new cooling and analysis traps, a new precision trap designed for lower magnetic gradients and a new loading trap for enhanced antiproton loading and storage were constructed. Additionally, we installed a new alignment system, new thermal shields, and new cryogenic electrical connections arranged for better RF shielding. New amplifiers and detection electronics were installed that achieved equivalent performance despite smaller volumes, thanks to advancements in coil design and understanding of loss in the trap. Challenges remaining to be addressed in the loading trap are also analyzed, mainly relating to the detection circuits in that trap. Two possible remedies are proposed, and when those are implemented all pieces will be in place for a new measurement.

A new measurement technique based on a separated oscillatory field method has been analyzed and shown promising for achieving a measurement at the 0.1 ppb level, and it becomes much more feasible with a shorter cyclotron cooling time. A simultaneous sequence of two pulses separated by an evolution time is used for both the spin and cyclotron frequency. The final state is determined by an interference pattern between the phase accumulated by the particle and the drive over the evolution time, allowing for a resonance as narrow as a few parts in 10^{10} . We have analyzed the expected limitations of the method in the environment of CERN, where magnetic field variation is a major concern.

Once this apparatus and the new techniques can be fully demonstrated, they promise to allow measurements with sub-ppb linewidth of both the proton and antiproton magnetic moments.

Bibliography

- J. DiSciacca, M. Marshall, K. Marable, G. Gabrielse, S. Ettenauer, E. Tardiff, R. Kalra, D. W. Fitzakerley, M. C. George, E. A. Hessels, C. H. Storry, M. Weel, D. Grzonka, W. Oelert, and T. Sefzick, Phys. Rev. Lett. **110**, 130801 (2013).
- [2] J. DiSciacca, M. Marshall, K. Marable, and G. Gabrielse, Phys. Rev. Lett. 110, 140406 (2013).
- [3] R. Frisch and O. Stern, Z. Phys. A: Hadrons Nucl. 85, 17 (1933).
- [4] I. Estermann and O. Stern, Z. Phys. A: Hadrons Nucl. 85, 4 (1933).
- [5] I. I. Rabi, J. M. B. Kellogg, and J. R. Zacharias, Phys. Rev. 46, 157 (1934).
- [6] Peter Franken and Sidney Liebes, Phys. Rev. **104**, 1197 (1956).
- [7] D. J. Collington, A. N. Dellis, J. H. Sanders, and K. C. Turberfield, Phys. Rev. 99, 1622 (1955).
- [8] J. H. Gardner and E. M. Purcell, Phys. Rev. 76, 1262 (1949).
- [9] H. Sommer, H. A. Thomas, and J. A. Hipple, Phys. Rev. 82, 697 (1951).
- [10] P. Frank Winkler, Daniel Kleppner, Than Myint, and Frederick G. Walther, Phys. Rev. A 5, 83 (1972).
- [11] Peter J. Mohr, David B. Newell, and Barry N. Taylor, Rev. Mod. Phys. 88, 035009 (2016).
- [12] Georg Schneider, Andreas Mooser, Matthew Bohman, Natalie Schön, James Harrington, Takashi Higuchi, Hiroki Nagahama, Stefan Sellner, Christian Smorra, Klaus Blaum, Yasuyuki Matsuda, Wolfgang Quint, Jochen Walz, and Stefan Ulmer, Science **358**, 1081 (2017).
- [13] C. Smorra, S. Sellner, M. J. Borchert, J. A. Harrington, T. Higuchi, H. Nagahama, T. Tanaka, A. Mooser, G. Schneider, M. Bohman, K. Blaum, Y. Matsuda, C. Ospelkaus, W. Quint, J. Walz, Y. Yamazaki, and S. Ulmer, Nature 550, 371 EP (2017).

- [14] T. Pask, D. Barna, A. Dax, R.S. Hayano, M. Hori, D. Horváth, S. Friedreich, B. Juhász, O. Massiczek, N. Ono, A. Sótér, and E. Widmann, Physics Letters B 678, 55 (2009).
- [15] A. Kreissl, A. D. Hancock, H. Koch, Th. Köhler, H. Poth, U. Raich, D. Rohmann, a. Wolf, L. Tauscher, A. Nilsson, M. Suffert, M. Chardalas, S. Dedoussis, H. Daniel, T. von Egidy, F. J. Hartmann, W. Kanert, H. Plendl, G. Schmidt, and J. J. Reidy, Zeitschrift für Physik C Particles and Fields **37**, 557 (1988).
- [16] F. Bloch and C. D. Jeffries, Phys. Rev. 80, 305 (1950).
- [17] J. H. Gardner, Phys. Rev. 83, 996 (1951).
- [18] S. H. Koenig, A. G. Prodell, and P. Kusch, Phys. Rev. 88, 191 (1952).
- [19] Robert Beringer and Mark A. Heald, Phys. Rev. **95**, 1474 (1954).
- [20] Sidney Liebes and Peter Franken, Phys. Rev. **116**, 633 (1959).
- [21] E. Richard Cohen and J. W. M. DuMond, in *Nuclidic Masses*, edited by Walter H. Johnson (Springer Vienna, Vienna, 1964), pp. 152–186.
- [22] R. Honerjäger and E. Klein, Zeitschrift für Physik **169**, 32 (1962).
- [23] J. H. Sanders, K. F. Tittel, J. F. Ward, and Brebis Bleaney, Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences 272, 103 (1963).
- [24] Than Myint, Daniel Kleppner, Norman F. Ramsey, and Hugh G. Robinson, Phys. Rev. Lett. 17, 405 (1966).
- [25] J. DiSciacca and G. Gabrielse, Phys. Rev. Lett. **108**, 153001 (2012).
- [26] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, Phys. Rev. 105, 1413 (1957).
- [27] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, Phys. Rev. Lett. 13, 138 (1964).
- [28] G. Lüders, Ann. Phys. 2, 1 (1957).
- [29] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov, New Journal of Physics 14, 095012 (2012).
- [30] Gerhart Lüders and Bruno Zumino, Phys. Rev. **106**, 385 (1957).

- [31] E. Abouzaid, M. Arenton, A. R. Barker, M. Barrio, L. Bellantoni, E. Blucher, G. J. Bock, C. Bown, E. Cheu, R. Coleman, M. D. Corcoran, B. Cox, A. R. Erwin, C. O. Escobar, A. Glazov, A. Golossanov, R. A. Gomes, P. Gouffon, J. Graham, J. Hamm, Y. B. Hsiung, D. A. Jensen, R. Kessler, K. Kotera, J. LaDue, A. Ledovskoy, P. L. McBride, E. Monnier, H. Nguyen, R. Niclasen, D. G. V. Prasad, X. R. Qi, E. J. Ramberg, R. E. Ray, M. Ronquest, A. Roodman, E. Santos, P. Shanahan, P. S. Shawhan, W. Slater, D. Smith, N. Solomey, E. C. Swallow, S. A. Taegar, P. A. Toale, R. Tschirhart, Y. W. Wah, J. Wang, H. B. White, J. Whitmore, M. J. Wilking, B. Winstein, R. Winston, E. T. Worcester, T. Yamanaka, E. D. Zimmerman, and R. F. Zukanovich, Phys. Rev. D 83, 092001 (2011).
- [32] R. S. VanDyck, P. B. Schwinberg, and H. G. Dehmelt, Phys. Rev. Lett. 59, 26 (1987).
- [33] Elise M. Novitski, Ph.D. thesis, Harvard University, 2017.
- [34] T.G. Myers G.Gabrielse, S.E. Fayer and X. Fan, .
- [35] N. Guise, Ph.D. thesis, Harvard University, 2009.
- [36] Jack Michael DiSciacca, Ph.D. thesis, Harvard University, 2013.
- [37] L. S. Brown and G. Gabrielse, Rev. Mod. Phys. 58, 233 (1986).
- [38] D. Wineland, P. Ekstrom, and H. Dehmelt, Phys. Rev. Lett. **31**, (1973).
- [39] G. Gabrielse, L. Haarsma, and S. L. Rolston, Intl. J. Mass Spec. Ion Proc. 88, 319 (1989), ibid. 93, 121 (1989).
- [40] D. J. Wineland and H. G. Dehmelt, J. Appl. Phys. 46, 919 (1975).
- [41] B. D'Urso, B. Odom, and G. Gabrielse, Phys. Rev. Lett. **90**, 043001 (2003).
- [42] B. D'Urso, R. Van Handel, B. Odom, D. Hanneke, and G. Gabrielse, Phys. Rev. Lett. 94, 113002 (2005).
- [43] L. S. Brown and G. Gabrielse, Rev. Mod. Phys. 58, 233 (1986).
- [44] 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).
- [45] 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).
- [46] G. Gabrielse, Adv. At. Mol. Opt. Phys. 50, 155 (2005).

- [47] P. Yesley, Ph.D. thesis, Harvard University, 2001.
- [48] James F. Ziegler, M.D. Ziegler, and J.P. Biersack, Nuc. Inst. and Meth. B 268, 1818 (2010).
- [49] 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, and E. A. Hessels, Phys. Lett. B 548, 140 (2002).
- [50] 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).
- [51] D. S. Hall, Ph.D. thesis, Harvard University, 1997.
- [52] D. Hanneke, S. Fogwell, and G. Gabrielse, Physical Review Letters 100, 120801 (2008).
- [53] Q. A. Turchette, Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland, Phys. Rev. A 61, 063418 (2000).
- [54] G. Gabrielse, A. Khabbaz, D. S. Hall, C. Heimann, H. Kalinowsky, and W. Jhe, Phys. Rev. Lett. 82, 3198 (1999).
- [55] S. Ulmer, C. Smorra, A. Mooser, K. Franke, H. Nagahama, G. Schneider, T. Higuchi, S. Van Gorp, K. Blaum, Y. Matsuda, W. Quint, J. Walz, and Y. Yamazaki, Nature 524, 196âĂŞ199 (2015).
- [56] Glen McPherson, Applying and Interpreting Statistics : A Comprehensive Guide, Springer Texts in Statistics, second edition. ed. (Springer New York : Imprint: Springer, New York, NY, 2001).
- [57] F. Bloch, Phys. Rev. **70**, 460 (1946).
- [58] Mason Marshall, Ph.D. thesis, Harvard University, 2019.
- [59] Norman F. Ramsey, Science **248**, 1612 (1990).
- [60] V P Chebotayev, A V Shishayev, B Ya Yurshin, and L S Vasilenko, Applied physics. 15, 43 (1978).
- [61] J. M. Raimond, M. Brune, and S. Haroche, Rev. Mod. Phys. 73, 565 (2001).
- [62] Ch. J. Bordé, Ch. Salomon, S. Avrillier, A. van Lerberghe, Ch. Bréant, D. Bassi, and G. Scoles, Phys. Rev. A 30, 1836 (1984).
- [63] Norman F. Ramsey, Phys. Rev. **76**, 996 (1949).

- [64] Vasant Natarajan, Kevin R. Boyce, Frank DiFilippo, and David E. Pritchard, Phys. Rev. Lett. 71, 1998 (1993).
- [65] Vasant Natarajan, Ph.D. thesis, Massachusetts Institute of Technology, 1993.
- [66] Norman Ramsey, Molecular Beams (Oxford Clarendon, ADDRESS, 1986), Vol. 2.
- [67] D. Hanneke, Ph.D. thesis, Harvard University, 2007.
- [68] David Phillips, Ph.D. thesis, Harvard University, 1996.
- [69] G. Gabrielse and F. Colin MacKintosh, Intl. J. Mass Spec. Ion Proc. 57, 1 (1984).
- [70] B. D'Urso, Ph.D. thesis, Harvard Univ., 2003.
- [71] S. Fogwell Hoogerheide, Ph.D. thesis, Harvard University, 2013.
- [72] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt, Rev. Mod. Phys. 87, 1419 (2015).
- [73] J.B. Koo and D.Y. Yoon, Metallurgical and Materials Transactions A 32, 1911 (2001).
- [74] A.K. Bhargava and M.K. Banerjee, in *Comprehensive Materials Finishing*, edited by MSJ Hashmi (Elsevier, Oxford, 2017), pp. 398 – 420.
- [75] S. P. Pucic, in 1993 IEEE Instrumentation and Measurement Technology Conference (IEEE, ADDRESS, 1993), pp. 114–117.
- [76] L. S. Brown and G. Gabrielse, Phys. Rev. A 25, 2423 (1982).
- [77] C. Bowick, *RF Circuit Design*, *RF Circuit Design* (Elsevier Science, ADDRESS, 2011).
- [78] J. B. Johnson, Phys. Rev. **32**, 97 (1928).
- [79] MACOR, Corning Incorporated Lighting and Materials.
- [80] H. Nagahama, G. Schneider, A. Mooser, C. Smorra, S. Sellner, J. Harrington, T. Higuchi, M. Borchert, T. Tanaka, M. Besirli, K. Blaum, Y. Matsuda, C. Ospelkaus, W. Quint, J. Walz, Y. Yamazaki, and S. Ulmer, Review of Scientific Instruments 87, 113305 (2016).
- [81] K. H. BennemannJohn B. Ketterson, Superconductivity, 3 ed. (Springer-Verlag Berlin Heidelberg, ADDRESS, 2008), Vol. 1.
- [82] Henry John Muller, Ph.D. thesis, University of Wisconsin Madison, 1989.

- [83] Jerzy Krupka, Krzysztof Derzakowski, Michael Tobar, John Hartnett, and Richard G Geyer, Measurement Science and Technology 10, 387 (1999).
- [84] W.W. Macalpine and R.O. Schildknecht, Proceedings of the IRE 47, 2099 (1960).
- [85] H. Dehmelt G. Gabrielse and W. Kells, Phys. Rev. Lett. 537 (1985).
- [86] P. Oxley, Ph.D. thesis, Harvard University, 2003.