New Apparatus and Methods for the Measurement of the Proton and Antiproton Magnetic Moments

A thesis presented
by

Mason Claflin Marshall

to
The Department of Physics
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of

Physics

Harvard University
Cambridge, Massachusetts
May 2019

Dissertation advisor

Gerald Gabrielse

Author

Mason Claflin Marshall

New Apparatus and Methods for the Measurement of the Proton and Antiproton Magnetic Moments

Abstract

The first direct measurement of the antiproton magnetic moment was performed using a single particle in a Penning trap. The result, $\mu_{\bar{p}}/\mu_N = 2.792\ 845\ (12)$ [4.4 ppm], is 680 times more precise than the previous best measurement. Together with a prior measurement of the proton magnetic moment in the same apparatus, this was the first direct comparison of the proton and antiproton magnetic moments. This stringent test of CPT invariance gave $\mu_{\bar{p}}/\mu_p = -1.000\ 000\ (5)$ [5.0 ppm]. The observation of individual spin flips of a single proton, also reported here, opened the possibility of further improving measurement precision by orders of magnitude.

Improving this result by a factor of $\sim 10^4$ requires measuring μ outside the large magnetic field gradient needed to detect spin flips. Two methods are proposed to avoid the leading uncertainties in such a high-precision two-trap measurement. One of these is to measure single spin-flips of a single proton or antiproton. The other is to induce multiple spin flips in the presence of a spin-cyclotron coupling drive, and observe the resulting change in the cyclotron energy. The design, construction, and commissioning of an appropriate apparatus, with high-quality particle detection and newly designed Penning trap electrodes, is reported.

Contents

List of Figures

List of Tables

Acknowledgments

I would first like to thank Jerry Gabrielse. During my time in the lab, I had the rare opportunity to experience the full gamut of precision measurement physics, from making a measurement on the first-generation apparatus to designing and constructing the second generation of the experiment. The understanding and experience with all stages of experimental science this provides is, I believe, one of the great advantages of working in the Gabrielse lab. Throughout my time in the group Jerry has been a great source of suggestions, advice and encouragement, as well as intellectual challenge, pushing me to develop my own ideas and myself as a scientist.

I owe a great deal to Jack DiSciacca, with whom I worked for the first two years of my Ph.D. In addition to teaching me practical skills I needed to contribute to the experiment (and, with Nick Guise, constructing the first-generation apparatus), Jack has been a role model for me, both as a scientist and a graduate student.

I am grateful for eight years of collaboration with Kathryn Marable. Having a colleague of the same year made my time working on the first-generation experiment both more successful and more enjoyable. In the years since I began work on the second-generation experiment, Kathryn has been a great source of both advice and support, and I believe our collaboration has improved both of the experiments.

I am also grateful to the younger students now taking over the project, Andra Ionescu and Geev Nahal. Working with them over the past few years has been a great pleasure, and I am comfortable the future of the experiment is in good hands.

I also would like to thank many other graduate students and postdocs in the Gabrielse lab. I benefited greatly from conversations with the electron team; my thanks to Elise Novitski, Josh Dorr, Shannon Fogwell, Ron Alexander, Maryrose

Barrios, Tom Myers, Sam Fayer, and Xing Fan. At CERN, we benefitted from work with members of the ATRAP collaboration, including Stefan Ettenauer, Eric Tardiff, Rita Kalra, Nate Jones, and Tharon Morrison. I also would like to thank Cris Panda, Cole Meisenhelder, and Daniel Ang of the ACME collaboration for experimental discussions and advice, as well as keeping the office lively after the move to Northwestern.

I have also been privileged to work with and mentor several talented undergraduate students, and am proud that many have continued in physics (including two who have joined this group). Thanks to Makinde Ogunnaike, Jan Makkinje, Andra Ionescu, Tom Myers, Jonah Philion, and Elizabeth Choi for their work on the experiment.

Outside the Gabrielse lab, I would like to thank Ed Myers for collaboration and discussions on precision measurement methods, and Dan Fitzakerly and Cody Storry for assistance establishing the experiment at CERN and interfacing with the AD.

I also thank my committee members, Professors Melissa Franklin and Ron Walsworth, for their investment in me and my work, and for support and helpful advice as I have made my way through grad school.

At Harvard, I have benefitted from the expertise of many talented specialists, especially Stan Cotreau, Mike McKenna, and Jim MacArthur. I am also grateful for the support of Jan Ragusa, Patricia McGarry, and Laura Nevins.

Finally, I would like to thank my parents Deb Claflin and Eric Marshall, my brother Duncan Marshall, and my wife Erin Hutchinson for constant support and encouragement as I navigated the ups and downs of graduate school. I owe so much to their ongoing love and support.

Chapter 1

Introduction

Penning traps have been used in some of the most precise measurements of particle and atomic properties [1, 2, 3]. Penning trap studies owe this remarkable precision to the simplicity of the system – a single charged particle is trapped in three dimensions by harmonic, static potentials, and kept isolated for months at cryogenic temperatures. Experimental techniques have been developed over several decades which allow delicate control over the trapped particle.

The program to measure the magnetic moments of a single proton and antiproton in a Penning trap has been under way in the Gabrielse lab since 2005. This was inspired by the remarkable success of the measurement of the electron magnetic moment with quantum jump spectroscopy [4, 5], the most precise measurement of any property of a fundamental particle. While the 700-fold smaller magnetic moment and 2000-fold smaller charge-to-mass ratio of the proton presented significant challenges, extending these techniques to the proton and antiproton opened a new avenue to test the standard model prediction of symmetry between matter and antimatter.

The work presented in this thesis can be broadly divided into three categories. A substantial contribution is reported to the culminating results in the first-generation apparatus, including the first single-particle measurement of the antiproton magnetic moment [6], as well as the first identification of single spin flips of an individual proton [7] (reported simultaneously with [8]). Second, we present and evaluate detailed proposals for two methods to achieve a further factor of 10^4 increase in precision, one of which was entirely developed during this work. Finally, an improved apparatus capable of implementing either of the proposed methods was designed, constructed, and commissioned with single trapped protons.

1.1 Proton and antiproton magnetic moments

The intrinsic magnetic moment is a property of a particle which describes the interaction of its spin with a magnetic field. This is usually expressed in terms of a dimensional constant and a dimensionless scaling factor. For the proton and antiproton, these are the nuclear magneton $\mu_N = e\hbar/2m_p$ and the g-factor $g_{p,\bar{p}}$. In terms of the dimensionless Pauli spin operators, $\boldsymbol{\sigma} = \boldsymbol{S}/(\hbar/2)$,

$$\begin{aligned} \boldsymbol{\mu_p} &= \frac{g_p}{2}\frac{e\hbar}{2m_p}\boldsymbol{\sigma} = \frac{g_p}{2}\mu_N\boldsymbol{\sigma} \\ \boldsymbol{\mu_{\bar{p}}} &= -\frac{g_{\bar{p}}}{2}\frac{e\hbar}{2m_p}\boldsymbol{\sigma} = -\frac{g_{\bar{p}}}{2}\mu_N\boldsymbol{\sigma} \end{aligned} \tag{1.1}$$

For a classical spinning charge distribution, g would be exactly 1, while for the spin moment of a Dirac point particle g would be exactly 2. Interactions with the QED vacuum give the electron its g value $g_e \approx 2.001$ (in terms of the Bohr magneton $\mu_B = e\hbar/2m_e$). The proton has $g \approx 5.59$; the substantial difference from the Dirac

particle value is due to the QCD substructure not present in the electron.

1.2 Measurement history

The history of proton magnetic moment measurements extends back to the 1930s, and encompasses conceptual and technical advances which contributed to the awarding of several Nobel prizes. The first measurements were performed in molecular beams during the 1930s by Stern et al. [9, 10]. The development of nuclear magnetic resonance by Rabi allowed Purcell and Bloch to improve upon this measurement in the 1940s [11]. The final major improvement, prior to single-particle Penning trap measurements, was the use of the hydrogen maser, developed by Ramsey in the 1950s [12]. The most precise determination of the proton magnetic moment before 2000 was made by Kleppner et al. during the 1970s, combining the hydrogen maser with several other measurements and theoretical corrections [13].

No analogue to any of these techniques exists for the antiproton magnetic moment – an antihydrogen maser or molecular-beam measurement would require cold antimatter amounts well beyond current global production capabilities. The previous best measurements of the antiproton magnetic moment were performed with approximately part-per-thousand precision. In these experiments, an antiproton beam collided with a lead [14, 15] or helium [16] target to form an antiprotonic atom (where an antiproton temporarily replaces the outer valence electron) and the hyperfine structure was interrogated to determine the magnetic moment.

In 2012 and 2013, our experiment contributed to this history with the first single-particle Penning trap measurements of the proton [17] and antiproton [6] magnetic

moments, at the few parts-per-million level. This was the first time the proton and antiproton magnetic moments could be directly compared in analogous experiments, and represented a factor of 680 improvement in precision over the previous best antiproton magnetic moment measurement.

Since our 2013 measurement, during the development of our new higher-precision methods and apparatus, parallel research has continued to advance the state of the art in proton and antiproton magnetic moment measurements. The international BASE collaboration has most recently measured the antiproton magnetic moment with a full width at half maximum of 32 ppb and a relative precision of 1.5 ppb [18], and the proton magnetic moment with a 3 ppb width and relative precision of 0.3 ppb [19]; the methods which led to those measurements are discussed in section 5.1.

Meanwhile, our team has been developing new methods to improve the ultimate measurement precision. Using the techniques and apparatus described here, we hope to achieve a significant further increase in precision – a measurement with a full width at half maximum of as little as 0.2 ppb.

Throughout this thesis, we evaluate prospective new measurement methods in terms of the lineshape full width at half maximum. We evaluate the width rather than the final measurement precision because that final precision will depend not only on the method used, but also on the time spent taking data to fit a lineshape center and on the level at which systematic effects can be controlled or corrected. A full analysis of systematic effects is beyond the scope of this thesis, and decisions about data collection will be made based on future experimental considerations. We choose the full width at half maximum rather than the Gaussian linewidth because the new

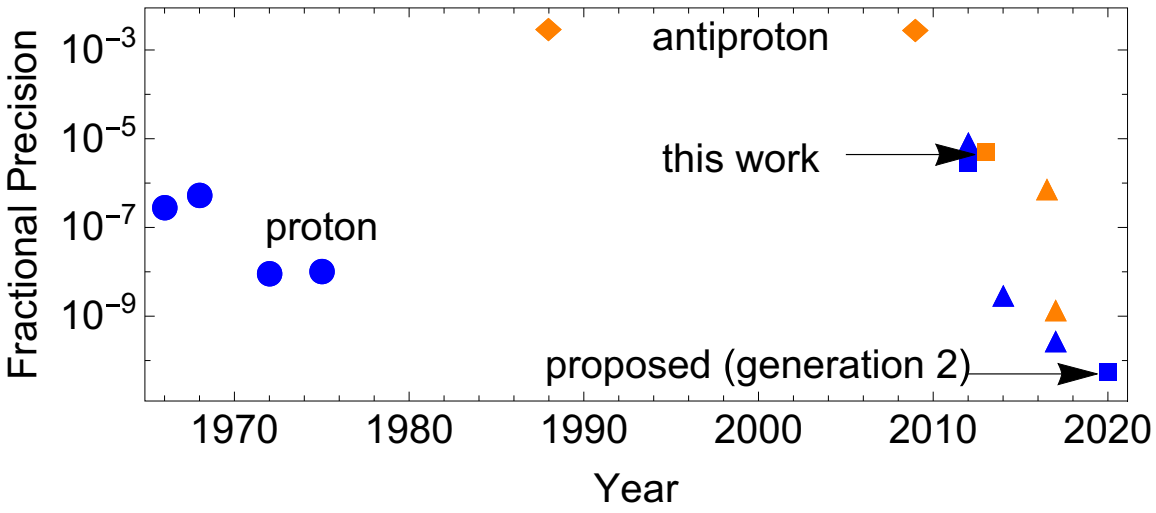

Figure 1.1: History of proton (blue) and antiproton (orange) magnetic moment measurements. Measurement methods: Circle – atomic beam; diamond – exotic atoms; square – Harvard single-particle Penning trap; triangle – BASE collaboration single-particle Penning trap.

measurement methods proposed will not yield normally distributed data, and the full width at half maximum provides a well-defined way to compare methods.

1.3 CPT symmetry

The standard model of particle physics predicts that the proton and antiproton magnetic moments should have equal amplitude and opposite sign. This is a consequence of CPT symmetry – the invariance of the standard model under simultaneous charge conjugation, parity inversion, and time reversal transformations.

Charge conjugation implies the replacement of particles with the corresponding antiparticles; parity inversion flips the sign of spatial coordinates; and time reversal inverts the time coordinate or, equivalently, reverses the direction of all motions. Applied to a particle in a vacuum, charge conjugation gives $q \to -q$ and thus a factor

of -1 to μ. The magnetic moment is even under a parity transformation, and time reversal inverts both the spin and all magnetic fields, giving a net factor of +1. The net CPT transformation thus gives $\mu_{\bar{p}} = -\mu_p$, and comparison of the proton and antiproton magnetic moments constitutes a test for CPT violation.

While no experimental evidence of CPT violation has yet been discovered, the history of fundamental symmetry studies in the 20th century motivates searches for it. Parity was thought to be an exact symmetry of nature until its violation was proposed by Lee and Yang [20] in 1956 and demonstrated by Wu [21] in 1957, in the beta decay of cobalt atoms. The simultaneous action of charge and parity (CP) was subsequently believed to be an exact symmetry, but violation was discovered in neutral kaon decays by Cronan and Fitch in 1964 [22].

A strong theoretical backing exists for CPT as an exact symmetry of nature. CPT symmetry is conserved by any quantum field theory which is both local and Lorentz invariant [23], including the Standard Model of particle physics. However, the Standard Model is known to be incomplete, lacking (for example) a reconciliation between quantum mechanics and general relativity at small length scales and large energies [24]. CPT violation can be found in quantum field theories allowing Lorentz violation, or in some higher-energy models like certain string theories [25, 26, 27].

Precise comparisons of CPT-invariant quantities are thus searches for physics beyond the Standard Model. Many such searches have been performed in different systems [28]. Lacking an a priori reason to prefer any specific beyond-standard-model theory, these CPT violation searches perform extremely precise measurements on simple systems across a broad range of energy scales and particle sectors.

One of the major shortcomings of the Standard Model is the inability to explain a matter-dominated universe – a Big Bang under the Standard Model would have produced nearly equal amounts of matter and antimatter, which would have annihilated [29]. This baryon asymmetry can be explained in the presence of CP violation, as well as baryon number violation and thermal non-equilibrium; however, the CP violation present in the Standard Model is insufficient to explain the observed asymmetry. CPT-violating interactions, together with baryon number violation, suffice to explain baryon asymmetry without thermal non-equilibrium [30]. Precision comparisons of properties of matter-antimatter partners, including the proton and antiproton magnetic moments, are thus a particular focus of searches for CPT violation.

The longstanding campaign to measure the electron magnetic moment has been extraordinarily successful. One of the main motivations of that work is as a test of theory – the comparison to QED calculations can determine the fine structure constant α, and comparison to other measurements of α sets the most precise test of QED itself [31] (as well as a strict the Standard Model itself, through couplings to virtual hadrons and the weak force). Because of its gluon substructure, the proton magnetic moment arises from QCD rather than QED interactions. While the precision of lattice QCD calculations of nucleon magnetic moments has improved rapidly in recent years [32, 33], the best calculated values for the proton magnetic moment remain 8-9 orders of magnitude less precise than the experimental precision, making comparison to theory a less compelling motivation than the test of CPT.

1.4 Overview of work presented

The work presented in this thesis includes a precision measurement of the antiproton magnetic moment; the development of a new method for the proton and antiproton magnetic moment measurements in a double Penning trap; and construction and commissioning of an apparatus to implement that method. These steps do not follow the standard order of a precision measurement Ph.D (which would culminate rather than begin with the measurement) because my work bridged two generations of the experiment. I joined the lab in the sixth year of the project, and spent the first three years of my Ph.D. working with the first-generation apparatus. During this period, I contributed to the first single-particle precision measurement of the antiproton magnetic moment [6] as well as the first identification of individual spin flips of a single proton [7] (reported simultaneously with [8]). This work was performed on an apparatus constructed by previous graduate students Nick Guise and Jack DiSciacca and described in detail in references [34] and [35].

After these successful results, I remained at Harvard and began work on a second-generation, higher-precision version of the experiment. In my fourth year, I designed a new experimental method and an improved apparatus, with the goal of sub-part-per-billion precision. During the subsequent two and a half years, I constructed, tested, and assembled all parts of the new apparatus. I then spent the last months before the move to Northwestern working with trapped protons to commission the apparatus.

In this thesis, I will describe the first-generation apparatus only insofar as necessary to report the results achieved between 2011-2014, or to explain the upgrades implemented on the new apparatus. The new apparatus, including the upgraded and

redesigned elements, is discussed in detail.

The presentation of work in this thesis will follow the trajectory of my Ph.D, reporting first the culminating results of work in the first-generation apparatus, followed by the progress made towards a sub-ppb measurement in the new methods and second-generation apparatus.

Chapter 2 develops the principles of particle trapping and control in a Penning trap. The basic theory of particle motions in a Penning trap is reviewed, as are the methods used to excite and measure the proton's motions and transition frequencies.

Chapter 3 presents our antiproton magnetic moment measurement. The chapter details methods used to trap and cool a single antiproton. The procedure used to measure the antiproton magnetic moment is also discussed, as are the lineshapes of the data. Finally, the measurement result is reported.

Chapter 4 focuses on identifying the spin state in the analysis trap, a pre-requisite for precision measurements using quantum jump spectroscopy. The chapter first discusses the procedure for identifying spin flips, as well as a framework for evaluating the reliability and efficiency of that identification. Our experimental demonstration of spin-state identification in the first-generation apparatus is presented. Finally, two methods are discussed to improve spin-state identification for future measurements.

Chapters 5 and 6 investigate methods to achieve sub-ppb precision in the second-generation experiment. Chapter 5 focuses on methods using single spin flips. This chapter first summarizes previously demonstrated techniques used in other groups for two-trap magnetic moment measurements. A method for measuring the cyclotron frequency with separated oscillatory fields is then discussed. Finally, chapter 5 presents

a proposed method to achieve sub-ppb precision with separated oscillatory fields and single spin flip identification.

Chapter 6 presents a new measurement method which I developed during this thesis. This method aims to improve statistical precision over single-spin-flip methods by accumulating the information from multiple spin transitions as changes in the proton's cyclotron state. The chapter presents a procedure for this measurement, including a treatment of the proton's quantum mechanical response to the procedure. The predicted lineshape and signal-to-noise ratio are calculated for different possible experimental scenarios.

Both chapters 5 and 6 include numerical estimates of data collection rates and achievable precision, as a guide for the future of the experiment. The results of these estimates and their sensitivity to various experimental parameters are compared at the end of chapter 6.

Chapters 7 and 8 present the improved design and construction of the apparatus and electronics. Chapter 7 focuses on the cryogenic detection electronics. The axial detectors were substantially improved during this thesis; the changes which enabled these improvements are discussed. This chapter also analyzes a previously insignificant effect which is found to limit the new detectors, feedback through the cryogenic FET; this limit is explained and a method to control the effect is proposed. Finally, the chapter reviews the behavior and potential improvements for the other detectors.

Chapter 8 presents the design and construction of the rest of the apparatus. The theory of Penning trap electrode design is detailed, including the improvements implemented in the new precision trap. The rest of the apparatus which was constructed

during this thesis is also discussed, with a focus on aspects which either saw significant improvement or posed new challenges, including the structure, cryostat, and DC and RF electronics. Finally, an in situ cryogenic alignment system is presented.

Chapter 9 discusses results with the new apparatus, focusing on several months of commissioning work which took place before the apparatus was moved to Northwestern University. Signals from single trapped protons are shown, and the status of work in the experiment is discussed.

Finally, chapter 10 provides detailed proposals for next steps towards a sub-ppb measurement of the proton and antiproton magnetic moments, as well as further improvements to the apparatus which could be implemented in the future.

Chapter 2

Proton in a Penning Trap

Penning traps are used for precision measurements because of the simplicity of the system and the precision available. A detailed review of single-particle Penning trap methods can be found in [1]. In this chapter, we will summarize the confinement and control of a single proton or antiproton as is needed for magnetic moment measurements.

2.1 Penning Trap Principles

An ideal Penning trap consists of a large, static, spatially homogeneous magnetic field and a superimposed, weaker electric field. The magnetic field provides radial confinement of charged particles, as they orbit field lines in a cyclotron motion. The superimposed electric field near the particle takes the form of a quadrupole, $V \propto 2z^2 - \rho^2$, providing a harmonic restoring force along the magnetic field axis. A charged particle acted on by these fields undergoes three separate harmonic motions. The

magnetic field induces a fast cyclotron orbit. The $\hat{z}$ component of the electric field induces a nearly harmonic axial oscillation. The combined action of the magnetic field and the radial electric-field component induce a slow magnetron $\mathbf{E} \times \mathbf{B}$ drift. The combination of these motions is illustrated in figure 2.1.

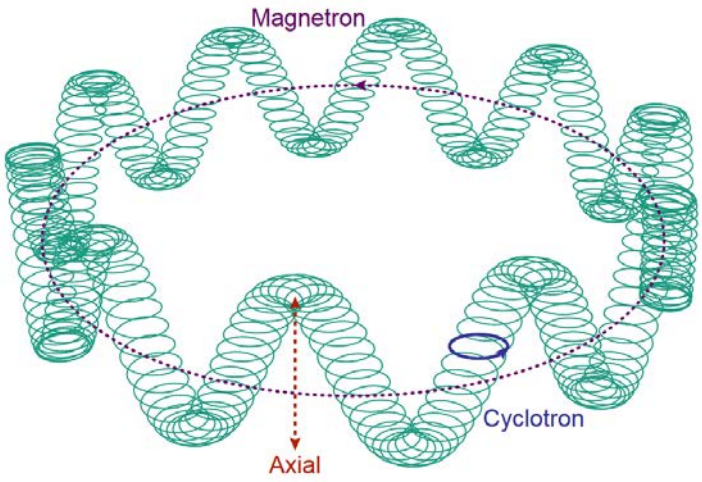

Figure 2.1: Motions of a charged particle in a Penning trap. Frequency differences are compressed for clarity. Figure from [36].

2.1.1 Axial motion

The field from an ideal quadrupole axial trapping potential can be written

$$V(z,\rho) = V_0 \frac{z^2 - \frac{\rho^2}{2}}{2d^2} \tag{2.1}$$

where V_0 is the trapping potential and d is the trap size, $d = \frac{1}{2}\left(z_0^2 + \frac{1}{2}\rho_0^2\right)$ [37] (using the coordinate system defined in figure 2.2). To create this field, we apply voltages to trapping electrodes surrounding the particle. To create a perfect quadrupole field, these voltages would be applied along hyperbolic equipotential lines; however, the

hyperbolic electrodes required would not allow access for antiprotons (and would have to be infinitely long). Instead, in this and all work with cold trapped antiprotons, we use open-access cylindrical Penning trap electrodes [38]. This allows access for the antiproton beam coming from the CERN Antiproton Decelerator (see section 3.1). To approximate a quadrupole field near the center of the trap, we use five electrodes – a central ring electrode, two endcap electrodes which are long compared to the radius, and two compensation electrodes. The trapping potential V_0 is applied to the ring, the endcaps are grounded through cryogenic resistors, and the compensation electrodes are held at V_c determined by the compensation ratio V_c/V_0. Figure 2.2 shows the precision trap electrodes for the new apparatus.

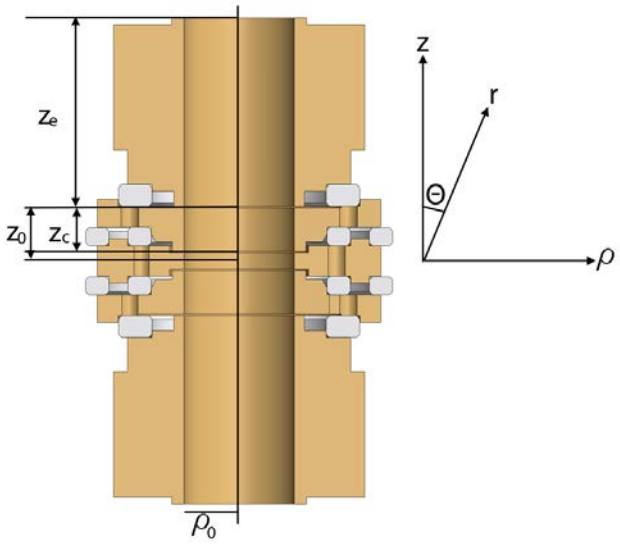

Figure 2.2: Coordinate system and trap dimensions for the precision trap.

The electrostatic potential can be expanded in Legendre polynomials, taking only

even terms to enforce reflection symmetry across the xy-plane:

$$V(\mathbf{r}) = \frac{V_0}{2} \sum_{\substack{k=0 \\ \text{even}}}^{\infty} C_k \left(\frac{r}{d}\right)^k P_k(\cos\theta) \tag{2.2}$$

The lowest nontrivial term, k=2, corresponds to the ideal harmonic z^2 potential. Higher-order terms are thus referred to as "anharmonic". The frequency of a harmonic oscillator is independent of its amplitude; these anharmonic terms introduce an amplitude dependence to the axial frequency. Amplitude dependence makes precise measurements of the particle's frequency difficult, because the axial motion is in equilibrium with a thermal bath, leading to changes in axial energy (see sections 2.3 and 7.2). The compensation electrodes allow these anharmonic terms to be minimized. With a compensation voltage V_c, the C_k coefficients can be written as

$$C_k = C_k^{(0)} + \frac{V_c}{V_0} D_k \tag{2.3}$$

and the ratio V_c/V_0 can be chosen to minimize anharmonic terms in $V(\mathbf{r})$.

In addition to allowing compensation for a range of thermal amplitudes, the trap design must fulfill the criterion of "orthogonality" [37], expressible as $D_2 \approx 0$. This implies that the axial frequency is relatively independent of the potential applied to the compensation electrodes. This is important experimentally – because of machining tolerances and other imperfections, the optimal ratio V_c/V_0 differs from its calculated value, and must be tuned empirically. Doing so would be much more difficult if the trapping potential and compensation ratio had to be tuned simultaneously.

The values of $C_k^{(0)}$ and D_k are determined by the trap geometry and calculated

in detail in [37, 38].[1] The expression for axial frequency at the center of a perfectly tuned, cylindrical Penning trap, with $\rho \to 0$ and all $C_{k>2} \to 0$, is thus

$$\nu_z = \frac{1}{2\pi}\sqrt{\frac{qV_0C_2}{md^2}} \tag{2.4}$$

In practice, trap imperfections mean that all higher-order terms cannot be simultaneously tuned to 0 – at the ratio where $C_4 = 0$, C_6 or higher may be nonzero. The axial frequency is modified by these anharmonic contributions, for axial amplitude A, as

$$\bar{\nu_z^2} = \nu_z^2\left(1 + \frac{3C_4}{2C_2}\left(\frac{A}{d}\right)^2 + \frac{15C_6}{8C_2}\left(\frac{A}{d}\right)^4 + ...\right) \tag{2.5}$$

Empirical tuning thus consists of finding the ratio at which the amplitude-dependent contributions for protons at thermal equilibrium most closely sum to zero.

At thermal energies, the axial quantum number is large enough to treat the axial motion classically. Because the motion is approximately harmonic, we can obtain the relationship between amplitude and energy from the classical expression for a harmonic oscillator. A quantum mechanical description is also useful for determining coupling rates and limits, and follows the form of the quantum harmonic oscillator.

$$E_z = \frac{1}{2}m\omega_z^2 z^2, \quad E_z = \hbar\omega_z\left(k + \frac{1}{2}\right) \tag{2.6}$$

[1] Note that these coefficients are calculated for a trap with applied voltage $\pm\frac{V_0}{2}$ on the endcaps and ring electrodes, and V_c on the compensation electrodes. In practice, the potential is more stable when the ring and comps are connected to a high-precision voltage source, with the endcaps grounded through cryogenic resistors. The voltages actually applied in the experiment, V_{ring} and V_{comp}, are related to the voltages in the above expansion by $V_0 = -V_{ring}$, $V_c = V_{comp} - \frac{1}{2}V_{ring}$.

2.1.2 Cyclotron and magnetron frequencies

The frequency of the cyclotron motion is set by the magnetic field strength. In the absence of the electric field, with a magnetic field **B**, it is given by:

$$\nu_c = \frac{1}{2\pi}\omega_c = \frac{qB}{2\pi m} \tag{2.7}$$

The cyclotron frequency is modified by the radial component of the quadrupole. The particle in both fields undergoes two circular motions, with frequencies

$$\nu_\pm = \frac{1}{2}\left(\nu_c \pm \sqrt{\nu_c^2 - 2\nu_z^2}\right) \tag{2.8}$$

$\nu_+ = \nu_c'$ is referred to as the "trap-modified cyclotron frequency", while $\nu_- = \nu_m$ is referred to as the "magnetron frequency". For an ideal trap, they are related by the expressions

$$\nu_+ = \nu_c - \nu_m, \qquad \nu_m = \frac{\nu_z^2}{2\nu_+} \tag{2.9}$$

Each of these circular motions also can be described as a simple harmonic oscillator. Accounting for the frequency hierarchy $\omega_+ \gg \omega_z \gg \omega_-$, the energies are

$$E_+ = \frac{1}{2}m\omega_+^2\rho_c^2 \qquad E_- = \frac{1}{2}m\left(\omega_-^2 - \frac{1}{2}\omega_z^2\right)\rho_m^2 \approx -\frac{1}{4}m\omega_z^2\rho_m^2 \tag{2.10}$$

$$E_+ = \hbar\omega_+\left(n + \frac{1}{2}\right) \qquad E_- \approx -\hbar\omega_m\left(l + \frac{1}{2}\right) \tag{2.11}$$

Note that the magnetron energy per state is negative, because the magnetron motion is unbound. Spontaneous decay and black-body coupling rates are small enough that it is effectively stable [1], but to get a proton at the trap center the magnetron motion must be "cooled" via a sideband coupling (section 2.6) to the axial motion.

2.1.3 Measuring g: the Larmor frequency

In addition to the physical motions described above and shown in figure 2.1, the proton's spin precesses in the magnetic field at the Larmor or spin-flip frequency

$$\nu_s = \frac{1}{2\pi}\omega_s = \frac{1}{2\pi}\mu \cdot \mathbf{B} = \frac{1}{2\pi}\frac{g}{2}\frac{qB}{m} \tag{2.12}$$

Comparing with equation 2.7, we see that ν_c and ν_s share the same B dependence. We can then write g as a ratio of the two frequencies:

$$\frac{g}{2} = \frac{\nu_s}{\nu_c} \tag{2.13}$$

In chapter 3 we report a g-factor measurement performed by taking the ratio of the two measured frequencies. In chapters 5 and 6 we analyze higher-precision approaches relying on simultaneous measurement of both frequencies.

To measure g we need to know ν_c, but can only directly measure ν_+. We take advantage of the Brown-Gabrielse invariance theorem [1], which relates the motions robustly despite distorted potentials, machining imperfections, or misalignments:

$$\nu_c = \sqrt{\nu_+^2 + \nu_z^2 + \nu_-^2} \tag{2.14}$$

Precisely measuring ν_c and g thus requires measuring all three frequencies. However, the frequencies are widely separated, with the hierarchy $\nu_+ \gg \nu_z \gg \nu_-$. A sub-ppb measurement of g therefore only requires 1 ppm accuracy in ν_z and 10% in ν_-.

2.1.4 Experimental frequencies and parameters

The axial and cyclotron frequencies are set by the electric and magnetic field strengths, respectively. We can set their values to avoid nearby radio and TV stations,

optimize the measurement linewidth (see chapters 5 and 6), and match narrow-band detection electronics (chapter 7). Both the first- and second-generation apparatus incorporate at least two open-endcap Penning traps – a copper precision trap with a homogeneous magnetic field, and an analysis trap with a large quadratic gradient, generated by an iron ring electrode and used to read out the spin state (as described in section 2.2). The new apparatus incorporates two additional Penning traps for particle loading and cooling, described in section 8.1.2. Figure 2.3 shows the Penning traps used in the 2012 measurement and the new apparatus, while table 2.1 shows experimental parameters for the precision and analysis traps in each apparatus.

2.2 The magnetic gradient

The frequencies in equation 2.14 represent physical motions of the proton and can be directly measured. Measuring the spin frequency, on the other hand, requires observing changes in the spin state. The saturated magnetism of a ring electrode made from high-purity iron applies a magnetic gradient of the form

$$\Delta\mathbf{B} = B_2 \left[\left(z^2 - \frac{\rho^2}{2} \right) \hat{B} - \left(\hat{B} \cdot \mathbf{z} \right) \boldsymbol{\rho} \right] \tag{2.15}$$

For a cold particle near the center of the trap, $\rho \approx 0$ and equation 2.15 reduces to $\Delta\mathbf{B} \approx B_2 z^2 \hat{B}$. See section 8.1.1 for further discussion of the effect of magnetic materials on the field in the context of Penning trap design.

The magnetic gradient adds a term to the potential energy of the proton's magnetic moment in the magnetic field, $\boldsymbol{\mu} \cdot \Delta\mathbf{B} = \mu B_2 z^2$. This adds to the z^2 term in the axial

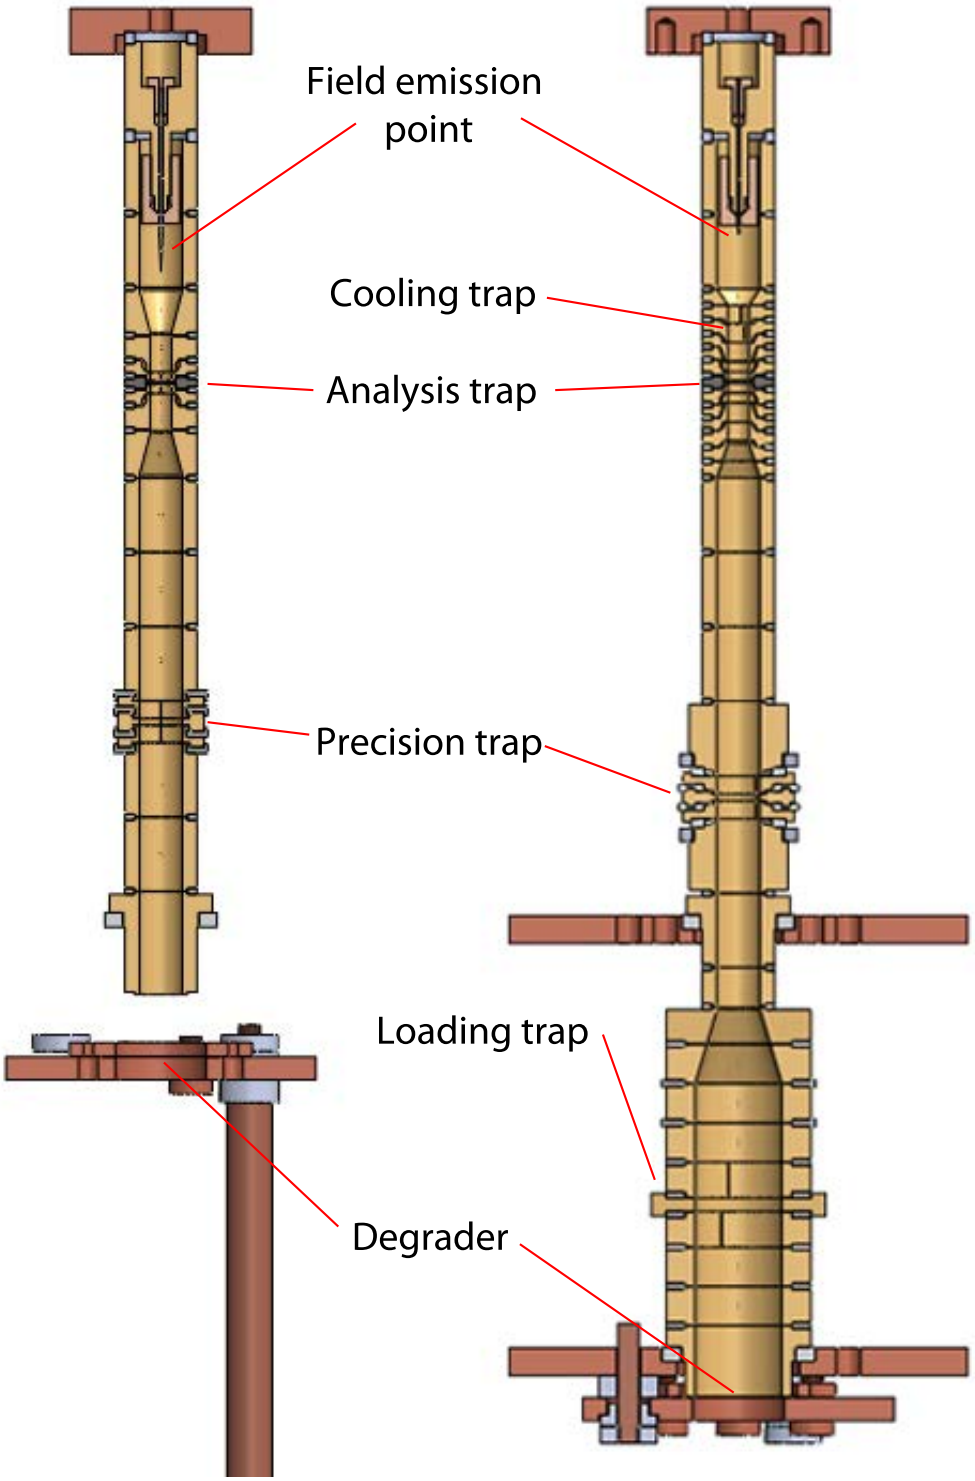

Figure 2.3: Penning trap stacks used in the 2012 antiproton measurement (left) and new apparatus (right), described in this thesis.

Table 2.1: Frequencies and parameters for the precision (P.) and analysis (A.) traps in the experiments described in this thesis. Magnetron motion parameters are given in the axial sideband cooling limit described in sec. 2.6.

Parameter	2012 P. Trap	2012 A. Trap	New P. Trap	New A. Trap
Magnetic Field	5.7 T	5.2 T	5.7 T	5.2 T
Trap Voltage V_0	-1.5 V	-1.2 V	-4.2 V	-1.5 V
ν axial	570 kHz	920 kHz	919 kHz	1075 kHz
ν cyclotron	86 MHz	79 MHz	87 MHz	$\sim$79 MHz
ν magnetron	1.9 kHz	5.0 kHz	4.8 kHz	$\sim$7.2 kHz
ν spin	240 MHz	220 MHz	243 MHz	$\sim$220 MHz
Trap radius (ρ_0)	3.0 mm	1.5 mm	3.0 mm	1.5 mm
Trap height (z_0)	2.93 mm	1.47 mm	2.93 mm	1.47 mm
Trap size (d)	2.58 mm	1.29 mm	2.58 mm	1.29 mm
Axial amplitude	z $\approx$ 70 μm	z $\approx$ 40 μm	z $\approx$ 44 μm	z $\approx$ 38 μm
Cyclotron ρ	$\rho_c \approx 0.5$ μm	$\rho_c \approx 0.5$ μm	$\rho_c \approx 0.5$ μm	$\rho_c \approx 0.5$ μm
Magnetron ρ	$\rho_m \approx 6$ μm	$\rho_m \approx 5$ μm	$\rho_m \approx 5$ μm	$\rho_m \approx 4$ μm
Axial Q.N.	k $\approx$ 146,000	k $\approx$ 91,000	k $\approx$ 91,000	k $\approx$ 77,000
Magnetron Q.N.	k $\approx$ 146,000	k $\approx$ 91,000	k $\approx$ 91,000	k $\approx$ 77,000
Cyclotron Q.N.	n $\approx$ 970	n $\approx$ 1050	n $\approx$ 960	k $\approx$ 1050
B2 (Tesla/m^2)	1	270,000	$\sim$0.1	270,000

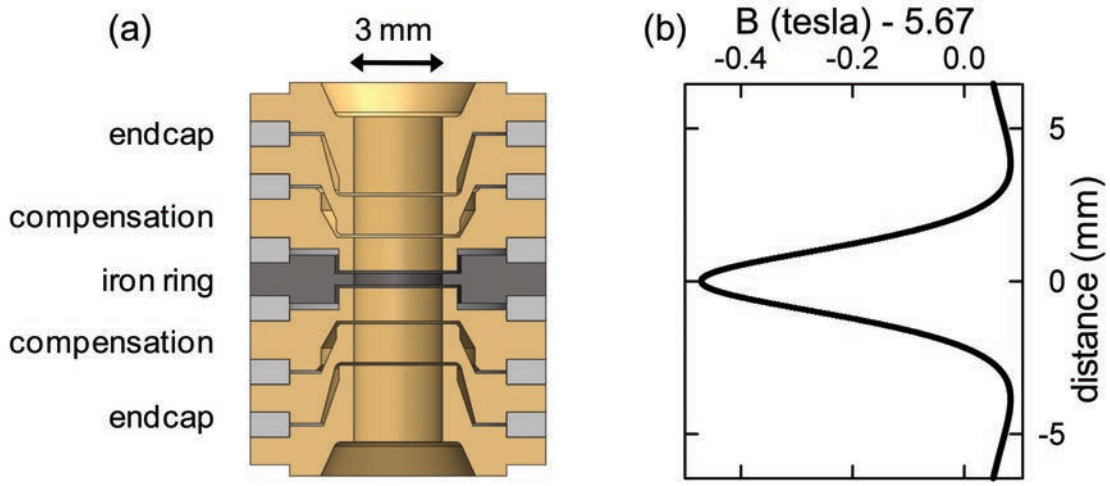

Figure 2.4: (a) Analysis trap in the new apparatus, including high-purity iron ring. (b) Magnetic bottle field caused by the iron ring.

potential of equation 2.1, shifting the axial frequency by an amount proportional to the total magnetic moment of the particle. The total magnetic moment includes contributions from the spin and the orbital spin and magnetron motions. The axial frequency shift due to the magnetic bottle is proportional to these three quantum numbers:

$$\Delta\nu_z \propto \left[\frac{gm_s}{2} + \left(n_c + \frac{1}{2}\right) + \frac{\nu_-}{\nu_+}\left(l + \frac{1}{2}\right)\right] \tag{2.16}$$

For our calculated bottle gradient of 290,000 T/m^2 and analysis trap axial frequency of 1.075 MHz, the frequency shift due to a spin flip would be approximately 120 mHz, with a cyclotron transition giving 41 mHz and a magnetron transition giving 3μHz. With an axial frequency of 919 kHz using the higher-SNR detector (see chapter 7), the frequency shifts would be 135 mHz, 49 mHz and 3 μHz. A precise and stable measurement of the axial frequency in the magnetic gradient lets us read out

changes in the magnetic moment. The gradient is therefore essential to the experiment. However, it also broadens the transition lineshapes, as discussed in sec. 3.4.

2.3 Image current detection

As a proton oscillates in the trap, it induces a current in the nearby trap electrodes. The induced current completes its circuit as it follows the particle's motion between electrodes. By placing a large resistance between adjacent electrodes, we can both damp the particle's motion and convert the miniscule current to a detectable voltage. In order that Johnson noise not excite undesired transitions, this large resistance is implemented as a resonant LC tank circuit; the principle of operation of these circuits are discussed in section 7.2.

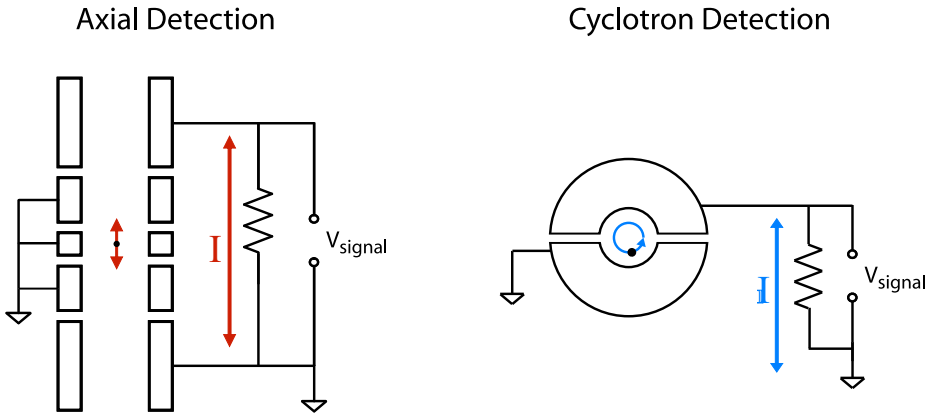

Figure 2.5: Simplified schematic of circuits for axial and cyclotron image current detection.

We take the axial motion as our example. The induced voltage as the image current passes across the resistance acts back on the particle with an electric force

$$f = -q\kappa\frac{IR}{2z_0} \tag{2.17}$$

where κ is a dimensionless geometrical constant, calculated in [34], and z0 is the trap height. This dissipates power in the resistor at the rate

$$-\dot{z}f = I^2R \tag{2.18}$$

Combining these equations gives

$$\dot{z}\frac{q\kappa IR}{2z_0} = I^2R \rightarrow I = \frac{q\kappa\dot{z}}{2z_0} \tag{2.19}$$

In terms of the axial amplitude A and frequency ω_z, we can write this as

$$I = \frac{q\kappa}{2z_0}A\ \omega_z \sin(\omega_z t) \tag{2.20}$$

At 4K in the precision trap, this current is approximately 2 fA. Because this current is so small, low-noise and high-quality detection is essential to the experiment.

The axial equation of motion of the particle, incorporating the restoring force from the trapping potential and the damping due to dissipation in the resistor, becomes

$$\ddot{z} + \gamma_z\dot{z} + \omega_z A^2 z = 0, \qquad \gamma_z = \left(\frac{q\kappa}{2z_0}\right)^2 \frac{R}{m} \tag{2.21}$$

We can model the interaction of the particle and detector as a series LC circuit [39, 40], which is useful for understanding the motion of a particle at equilibrium driven by thermal noise. Substituting equation 2.19 into equation 2.21 gives

$$\frac{2z_0}{q\kappa}\dot{I} + \frac{q\kappa}{2z_0}\frac{R}{m}I + \omega_z A^2 \frac{2z_0}{q\kappa}\int I dt \tag{2.22}$$

$IR = V$ gives the detected signal, and we can set the other coefficients equal to the coefficients in a series LC circuit to get the approximate values

$$L_{eff} = m\left(\frac{2z_0}{q\kappa}\right)^2 \approx 2\times 10^7 H, \qquad c_{eff} = \frac{1}{L_{eff}\omega_z^2} \approx 1.5\times 10^{-21} F \tag{2.23}$$

2.4 Measuring and driving the axial frequency

At 1 MHz, a superconducting coil is used to form the effective resistance for axial detection; the resonant impedance is inversely proportional to dissipation in the circuit, so the low-loss superconductor enables a large effective resistance. Additionally, the amplitude at thermal equilibrium is much larger for the axial motion than the cyclotron (table 2.1), creating larger image currents. Because of the large resistance and amplitude, the axial motion has a higher signal-to-noise ratio and is therefore the main tool used for trap diagnostics. Several techniques for axial frequency measurement are used at different points in the experiment.

2.4.1 Driven axial detection

For a motion to be observed it must be near the LC circuit's resonant frequency, where the resonator has a high effective resistance. Changing the trapping potential changes the frequency of the axial motion; the potential must be adjusted to bring it into the detector's bandwidth. When the exact trapping potential required is uncertain, this may require a broad search across a range of voltages. Examples include finding a response after cooling down; tuning the voltage ratio for anharmonicity compensation; or finding a response in the analysis trap after a transfer, when the cyclotron and magnetron radii cause an unknown shift in the magnetic bottle. These broad searches are usually performed with an axial drive.

Two RF signals are applied to a trap electrode, at frequencies which sum to ν_z. (The specific frequencies are chosen to avoid mixing with AM radio stations or other ambient signals to create noise within the range of the amplifier.) The lower of the

two frequencies acts as a modulation to the trapping potential [1], and the particle sees the combination of the drives as a signal at its axial frequency. This modifies equation 2.21 with the addition of a forcing term:

$$\ddot{z} + \gamma_z \dot{z} + \omega_z A^2 z = F(t)/m \tag{2.24}$$

Appropriately chosen drive amplitudes increase the current of equation 2.20 above the resonator's Johnson noise. Applying two drives which sum to ν_z avoids direct feedthrough into the amplifier, which would otherwise overwhelm the proton signal.

When finding the resonant trapping voltage, the drive frequency is held fixed while the voltage is swept. Changing the voltage changes the particle's axial frequency (eqn. 2.4), exciting the axial motion when it reaches the fixed drive frequency. Sweeping the voltage lets us cover a larger range than the frequency, where we would be limited by the resonator bandwidth. The high signal-to-noise axial drives can cover a broad parameter space, and are often used for characterization of a trap after cooling down. This is usually done with several trapped protons, as the bandwidth of the response increases with particle number.

Once an axial response has been found, the driven axial signal can also be used to measure the anharmonicity and adjust the compensation electrode voltage. This procedure has been discussed in [34]. Figure 2.6 shows driven axial responses for different compensation ratios $\frac{V_c}{V_0}$ in the new apparatus, demonstrating anharmonicity tuning using axial drives.

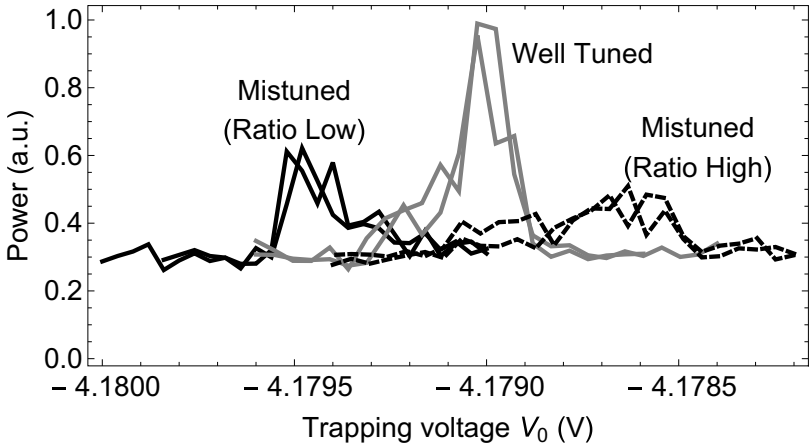

Figure 2.6: Driven axial responses in the new apparatus precision trap. Comp-to-ring voltage ratio differs by 0.5% between each pair of scans.

2.4.2 Axial dips

In the absence of a drive, eqn. 2.22 shows that the particle behaves like a series LC circuit, which passes current at its resonant frequency. The particle thus acts as a short to ground for Johnson noise at the axial frequency. When it is resonant with the detector, we can amplify the Johnson noise and observe the proton as a "dip" in the power spectrum. Figure 2.7 shows a single-particle dip in the new apparatus.

The width of the dip when well tuned is set by the damping width γ_z from equation 2.21. In the presence of anharmonicity, the dip is broader and shallower. As the particle is driven by Johnson noise, it constantly selects different energy states from a Boltzmann distribution. When C_4 or greater in equation 2.2 are nonzero, the terms in equation 2.5 give an amplitude-dependent shift in the axial frequency. The observed dip in the noise spectrum is thus a convolution of the anharmonic potential with the distribution of axial states. Figure 2.8 shows a set of dips in the new apparatus at different compensation ratios for anharmonicity tuning.

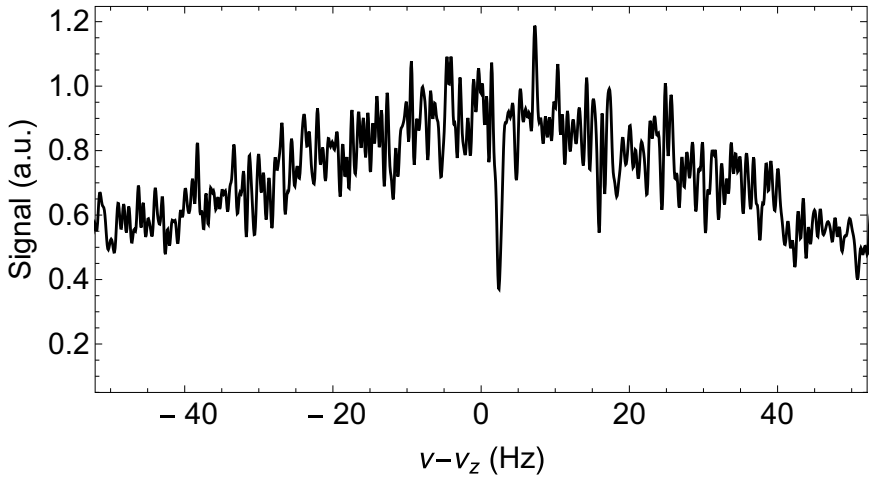

Figure 2.7: One-particle dip in the new apparatus precision trap. Dip width $\gamma_z \approx 2\pi * 0.7$ Hz is consistent with $\kappa \approx .335$ [34] and R$\approx$ 90 MΩ (table 7.3) in eqn. 2.21.

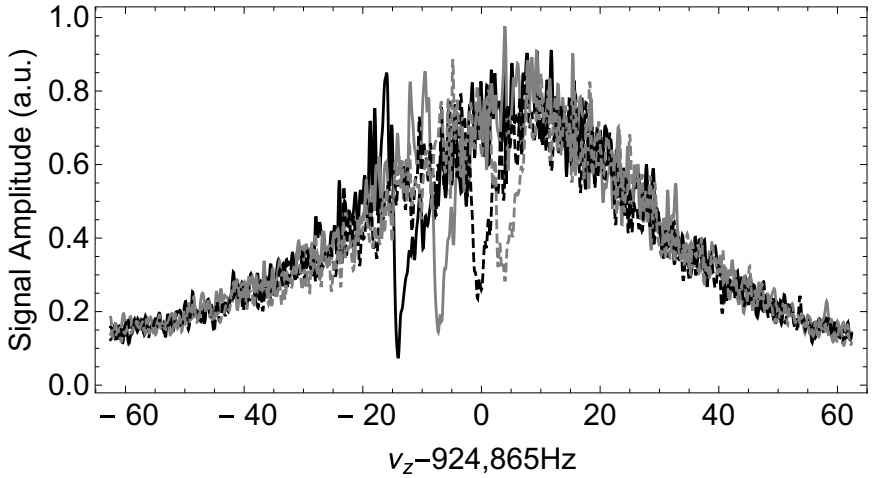

Figure 2.8: Dip spectra at different ratios V_c/V_0 (for 3 protons), taken while tuning the trap in the new apparatus. The more harmonic the trap is, the deeper and narrower the dip; the width of the dips pictured here ranges from 3.3 Hz (rightmost dip) to 1.8 Hz (leftmost dip). The ratio differs by 0.06% between scans.

2.4.3 Feedback and the self-excited oscillator

The undriven axial signal can be amplified, phase-shifted, and applied back to the particle using the axial drive line. This modifies the equation of motion, replacing the forcing term of equation 2.24 with a term proportional to I (or $\dot{z}$):

$$\ddot{z} + (1 - G\cos(\phi))\,\gamma_z \dot{z} + \omega_z A^2 z = 0 \tag{2.25}$$

This feedback can be applied to modify the damping rate and change the effective electronic temperature of the damping resistance, demonstrated with electrons in [41, 42] and with protons in [34, 43]. A detailed treatment of the effect of feedback on axial temperature is presented in [35].

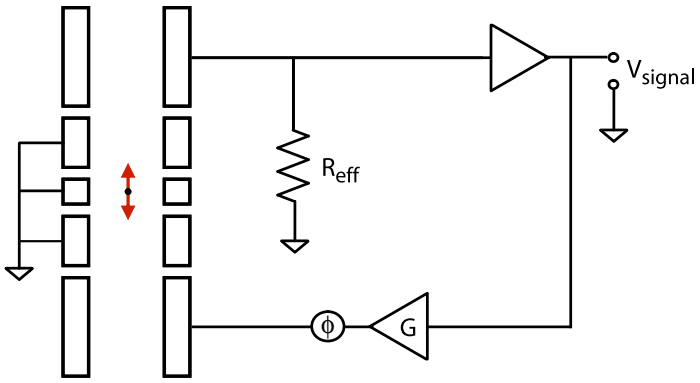

Figure 2.9: Schematic of circuit for axial feedback

If the gain and phase G, ϕ are adjusted such that $G\cos(\phi) = 1$, equation 2.25 becomes an undamped simple harmonic oscillator. When this condition is met, the particle oscillates freely at its natural resonant frequency. This is equivalent to a driven signal where the proton itself provides the drive – a self-excited oscillator (SEO) [42, 44]. The proton can then oscillate undamped with a signal over the

amplifier noise floor, giving the high signal-to-noise and fast measurement time of a driven signal without forcing the proton's frequency to match a drive.

Stably achieving this condition requires actively locking the proton's amplitude. Without a mechanism to control the strength of feedback, noise fluctuations would cause $G\cos(\phi)$ to temporarily differ from 1, resulting in exponential growth or damping of the amplitude. We therefore employ a digital signal processor (DSP) to control the feedback strength and hold a stable axial amplitude.

The DSP Fourier transforms the signal from the detector and determines the maximum amplitude at a single frequency (presumed to be the proton's motion). It then adjusts a control voltage to a voltage variable attenuator (VVA). The VVA attenuates the feedback signal. The DSP and VVA acting together increase feedback strength if the axial amplitude is too small, and reduce it if the amplitude is too large. This lock loop holds the proton at a constant amplitude above the noise level. Stability of the lock loop is key not only to avoid runaway damping or signal growth, but also to avoid frequency fluctuations in the presence of anharmonicity.

The SEO enables fast, precise detection of the proton signal. Figure 2.10 shows the SEO signal from a single particle, compared to a dip, in the apparatus used for the 2013 antiproton measurement. The high signal-to-noise and fast measurement time with the SEO are most useful when determining the spin and cyclotron states in the analysis trap, as this requires resolving very small frequency differences. The SEO is not suitable to determine ν_c using the invariance theorem – all three frequencies must be measured under the same conditions to avoid systematic effects. However, the axial frequency resolution needed to determine ν_c is significantly lower than that

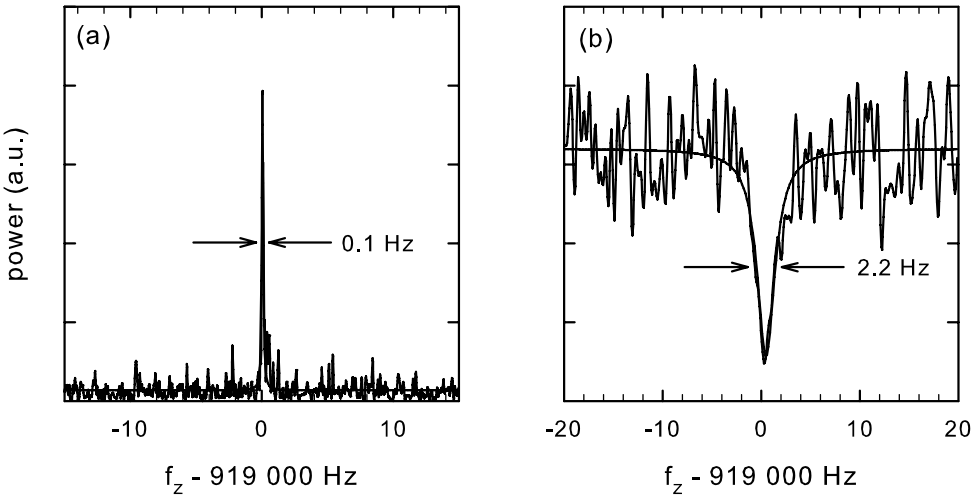

Figure 2.10: Axial signals in the analysis trap of the apparatus used for the 2013 antiproton magnetic moment measurement. (a) SEO signal averaged for 16 seconds. (b) Dip signal averaged for 80 seconds. Figure from [35].

needed to identify the magnetic moment in the analysis trap.

2.5 Cyclotron driving and detection

Similarly to the axial motion, the cyclotron motion is detected by forcing the image current through an effective resistance (figure 2.5). The smaller amplitude of the cyclotron motion and lower effective resistance[2] achievable at 87 MHz limit direct detection of the cyclotron motion to driven signals. However, compared to the axial motion, the cyclotron motion sees a stronger trapping force from the magnetic field; we can drive the particle's energy up to several keV before risking particle loss.

[2]See section 7.4 for details on the cyclotron resonator. In general, the effective resistance of an LC circuit on resonance is proportional to the inductance; a much smaller inductance is needed for the higher-frequency cyclotron resonator, giving weaker damping.

The cyclotron equation of motion is a damped harmonic oscillator [1]:

$$\ddot{v}_x + \gamma_c \dot{v}_x + \omega_+^2 v_x = 0 \tag{2.26}$$

where $v_{x,y}$ are the radial coordinates and γ_c is the cyclotron damping rate, given by

$$\gamma_c = \left(\frac{e\kappa_c}{2\rho_0}\right)^2 \frac{R}{m} \tag{2.27}$$

with a dimensionless geometrical constant κ_c, calculated in [45]. This expression for the damping rate is similar to that for the axial motion, but the effective resistance R is several orders of magnitude smaller. Because of the lower resistance and larger accessible energy, the optimal way to observe the cyclotron motion is to apply a drive and increase the radius, then turn off the drive and measure the image current as energy damps into the resistor. Figure 2.11 shows excited cyclotron peaks for antiprotons observed during the 2013 measurement.

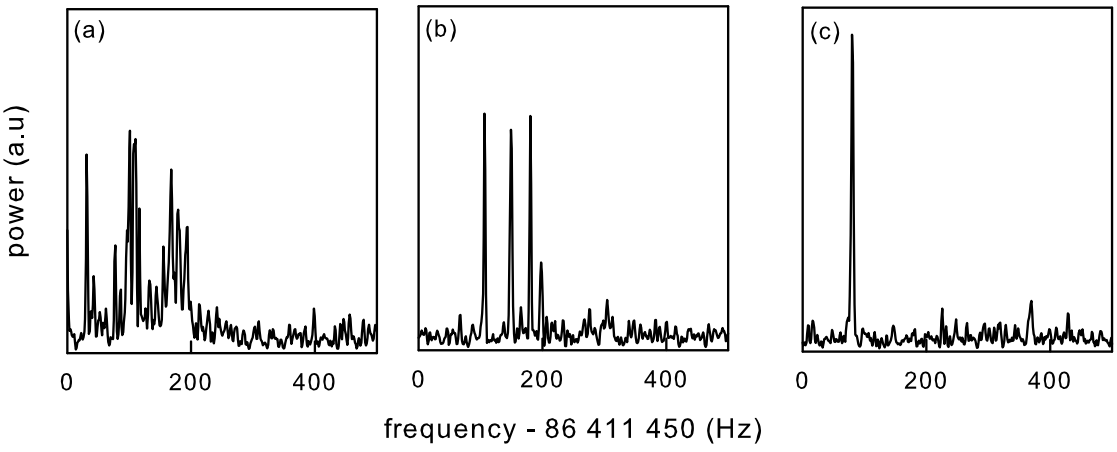

Figure 2.11: Excited cyclotron signals for different particle number: (a) <100 antiprotons, (b) 4 antiprotons, (c) a single antiproton. Figure from [35].

The peaks in figure 2.11 are separated by the different energies of the individual

particles, with a frequency shift due to special relativity. When a proton is driven to a large amplitude, the cyclotron motion acquires a relativistic mass shift, proportional to the Lorentz factor

$$\gamma = \frac{1}{\sqrt{1 - v^2/c^2}} \tag{2.28}$$

This results in a frequency shift proportional to the energy of the motion:

$$\frac{\Delta\nu_+}{\nu_+} = -\frac{E_c}{E_c + m_p c^2} \approx -\frac{E_c}{m_p c^2} \tag{2.29}$$

For a cyclotron energy of 1 keV, the frequency shift is approximately 1 part per million, or almost 100 Hz out of the 87 MHz cyclotron frequency.

This shift must be accounted for when exciting the cyclotron motion. Rather than applying a drive at a set frequency, we sweep the drive from above to below the zero-energy cyclotron frequency. At each step, the motion acquires enough energy for the relativistic mass increase to bring it out of resonance with the drive. The drive is then stepped down again into resonance, building up a large excitation.

Once the drive is turned off, the cyclotron motion dissipates energy into the resistor, with power $I^2R \propto (\gamma_c \rho_c \omega_c)^2 R$. Equation 2.10 shows that $E_c \propto \rho_c^2$; the power dissipation is thus proportional to the cyclotron energy, resulting in an exponential decay $E_c = E_0 e^{-t/\tau_c}$, with time constant $\tau_c = 1/(2\gamma_c)$. We can observe this exponential decay and fit the data to obtain the zero-energy cyclotron frequency. Figure 2.12 shows a single-particle cyclotron decay in the new apparatus.

This technique was used with high precision for comparison of the proton and antiproton charge-to-mass ratios [46]; however, the 2013 antiproton magnetic moment measurement (chapter 3) and proposed sub-ppb measurements (chapters 5 and 6) all rely on other methods than image current detection to obtain ν_+ and determine g.

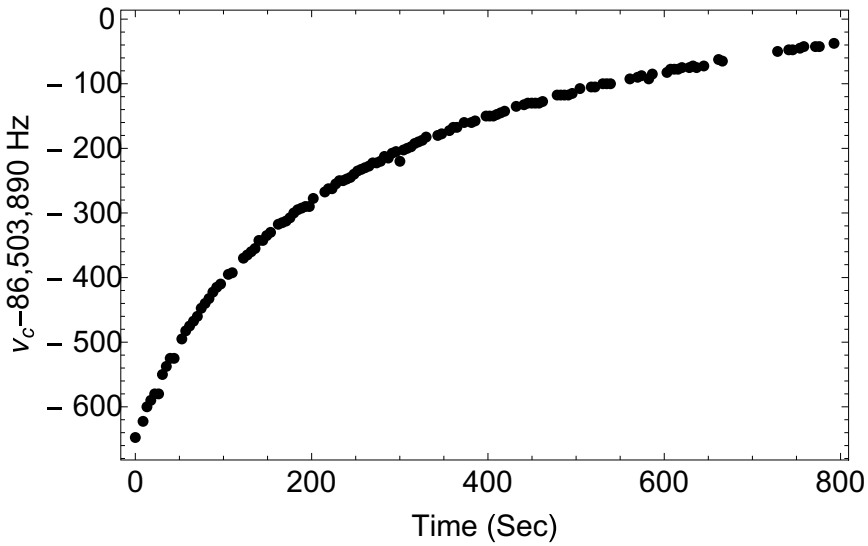

Figure 2.12: Cyclotron decay in the new apparatus, demonstrating the relativistic mass shift as the excited particle damps energy into the detector.

That said, measuring the zero-energy cyclotron frequency using decays is important for diagnostics of the magnetic field – see section 10.1.7.

2.5.1 Obtaining a single proton

In addition to measurement of the zero-energy cyclotron frequency, the excited cyclotron signal is useful for obtaining a single proton or antiproton. Protons are loaded by excitation and ionization of adsorbed gas using an electron beam, fired from a field emission point, followed by a filtered noise drive [47]. Antiprotons are loaded by trapping a pulse from CERN's Antiproton Decelerator (see ref. [48] or sec. 3.1). In either case, between 1 and 100 particles are loaded at a time. Sweeping a drive through the cyclotron resonance excites the motion of each particle. Collisions and interactions between particles separate them at different radii with different relativistic shifts, observable as individual peaks on Fourier transformed detector signal.

While monitoring the number of particles using these individual peaks, we can iteratively reduce the axial trapping potential and spill protons out of the trap. This lets us reliably obtain a single proton or antiproton. Figure 2.11 demonstrates this – the same cloud of particles is repeatedly excited and measured after reducing the trap depth. It is worth noting that several drive sweeps are required to prove a single proton is trapped, as multiple protons can present a single peak – a collision between particles early in a drive sweep could push one out of resonance with the drive; or multiple particles can move with the same radius. Figure 2.13 shows simultaneous decays from two protons.

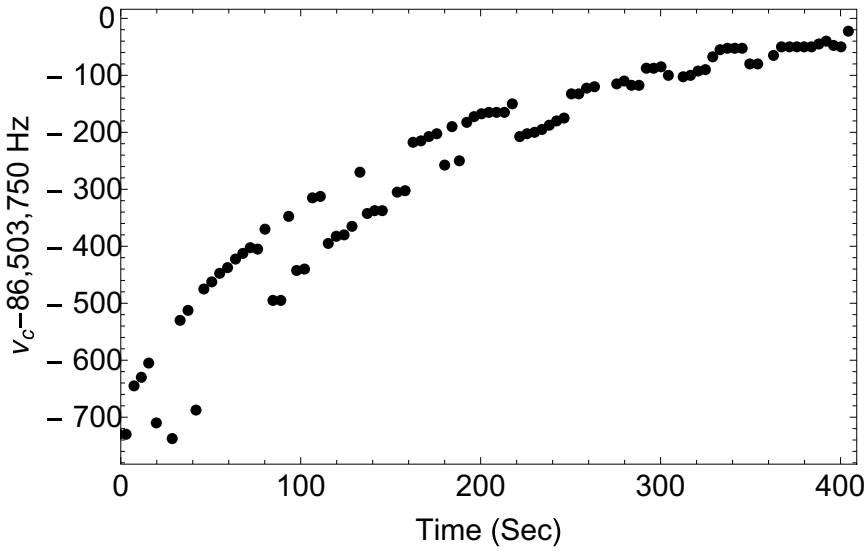

Figure 2.13: Tracking the frequency with the highest amplitude on the cyclotron detector – two particles at different energies present peaks at different frequencies (similar to fig. 2.11(b)).

2.6 Magnetron sideband cooling

As discussed in section 2.1.2, the magnetron motion is undamped but effectively stable. However, the presence of noise near the magnetron frequency – either Johnson noise from resistances between split electrodes, or low-frequency RF which might be picked up on a ground loop – can cause growth in the magnetron radius. A large magnetron radius can put the proton far from trap center in a region with an anharmonic electric field, or even lead to particle loss. To keep the magnetron radius small, as well as to initially center the particle after loading, we "cool" the magnetron motion by sideband coupling it to the damped axial motion[3].

An RF drive applied with xz symmetry at the cooling frequency $\nu_z + \nu_m$ transfers energy between the two motions at a rate proportional to the quantum numbers [1]. The axial motion is coupled to the effective resistance of the resonator, acting as a thermal bath. The limit of sideband cooling is therefore to equalize the magnetron and axial quantum numbers. The magnetron motion after sideband cooling thus follows a Boltzmann distribution, characterized by

$$T_m = \frac{\nu_m}{\nu_z} T_z, \qquad \rho_m = \sqrt{\frac{4k_B T_m}{m\omega_z^2}} \tag{2.30}$$

A heating drive with xz symmetry at $\nu_z - \nu_m$ increases the energy in both motions.

2.6.1 Using sideband drives for initial particle detection

As both the heating and cooling sideband drives put energy into the axial motion, they can both be used to excite a signal in the axial detector. These signals are useful

[3]We call the process which reduces the quantum number "cooling" as it brings the particle closer to trap center, even though this technically increases the potential energy in the magnetron motion.

in a few ways during operation of the experiment. With large particle numbers, especially when different ions are present, the damping from the axial detector can exceed the coupling to the axial drive [34]. Strong axial drives can mix in the FET and feed into the signal chain, overwhelming particle signals. Sideband drives can be made strong enough to excite the axial motion of a cloud of particles with a large signal. Figure 2.14 shows a voltage scan with a sideband heating drive – note that the response is millivolts wide, compared to hundreds of microvolts for the axial scans in figure 2.6. Figure 2.15 shows a time-series of the detector power spectrum. Two protons are visible, excited at different voltages during the scan. The shift in frequency is typical for sideband heating – as the magnetron radius grows, the proton gets closer to the electrode and sees a larger potential, increasing the axial frequency.

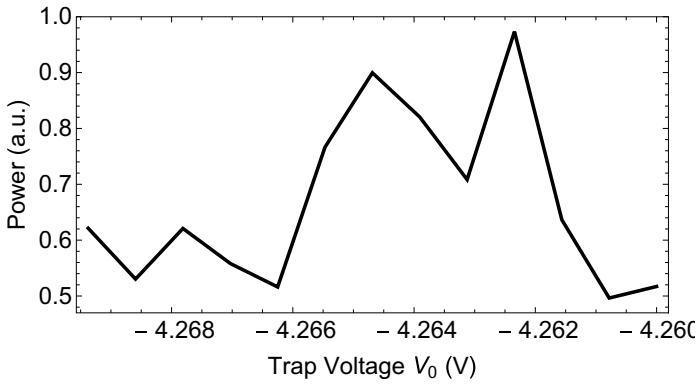

Figure 2.14: Axial response as trap voltage is scanned while a sideband heating drive is applied.

Additionally, because of the large available drive strength, applying a short set of heating and cooling drives is the fastest way to confirm a small number of protons is present in the trap. Figure 2.16 shows a pair of heating and cooling peaks for a

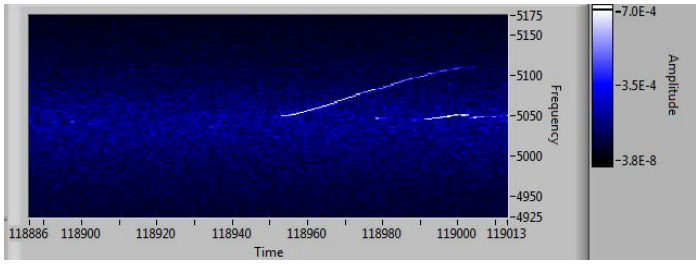

Figure 2.15: Time trace of the axial detector output, demonstrating sideband heating peaks during a voltage scan. Color intensity represents signal amplitude; each column is an FFT of the detector signal, averaged for two seconds.

proton in the new apparatus.

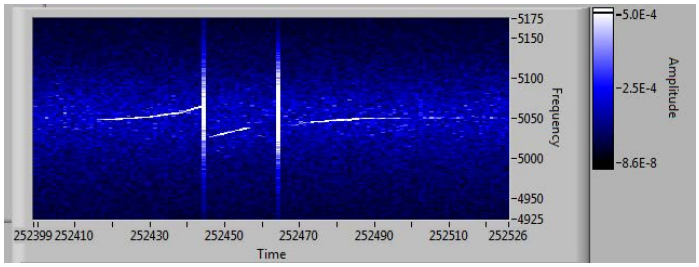

Figure 2.16: Heating and cooling peaks on the axial detector are an efficient way to confirm a proton is present and near the center of the trap. With the heating drive, the frequency grows away from the center of the amplifier resonance; with the cooling drive, it damps back to the original value.

After transferring between traps, a particle often requires sideband cooling to return it to the center of the trap. When the particle enters the analysis trap with an undetermined cyclotron quantum number, the axial frequency is not precisely known because of the bottle shift (eqn. 2.16). A voltage sweep with a cooling drive applied is useful to both find the axial motion and re-center the proton in the trap.

2.6.2 Magnetron frequency measurement – "split dips" and avoided crossing

At the limit of sideband cooling, when the axial and magnetron quantum numbers are equal, the sideband coupling no longer increases the axial energy. The axial motion is thus effectively undriven, and we can measure a thermal dip signal. However, the presence of the sideband coupling drive results in a "split" dip, with two separate features whose frequency depends on the drive strength and detuning. Figure 2.17 shows a sideband-coupled split dip from a single proton in the new apparatus.

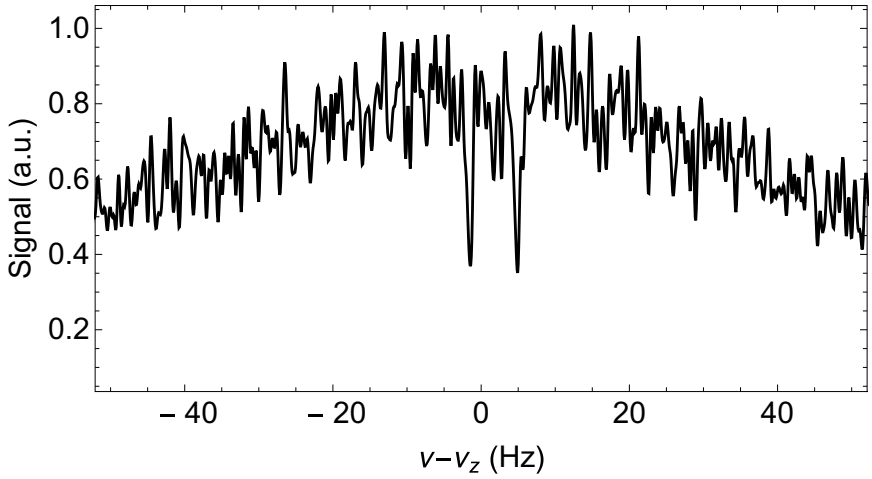

Figure 2.17: Split dip demonstrating avoided crossing from axial-magnetron sideband drive.

This behavior has been analyzed as a classical avoided crossing [49] and demonstrated with trapped ions [49, 50] and protons [34, 51]. The frequencies of the split dips can be used to precisely measure the magnetron frequency, assuming the undriven axial frequency is well known, by solving the equations [49]

$$\epsilon_{\pm} = -\frac{\delta}{2} \pm \frac{1}{2}\sqrt{\delta^2 + |V|^2} \tag{2.31}$$

where $\epsilon_\pm$ is the detuning between either dip and the axial frequency, δ is the detuning between the the applied drive and the true cooling frequency $\nu_z+\nu_m$, and $|V|$ is related to the strength of the coupling drive. Fitting these two equations gives δ and $|V|$, and δ gives ν_m. Additionally, because a split dip only appears when the magnetron and axial quantum numbers are finally equalized, its presence can be taken to show the sideband cooling limit has been reached.

2.7 Spin flip and anomaly drives

The spin precession frequency does not describe a physical motion of the particle, and thus cannot be detected using image currents; instead we use the bottle coupling (section 2.2). To measure g, we must therefore drive spin flips. This is done by applying current at the Larmor frequency through an electrode as pictured in figure 2.18 (a), creating an oscillating perpendicular magnetic field.

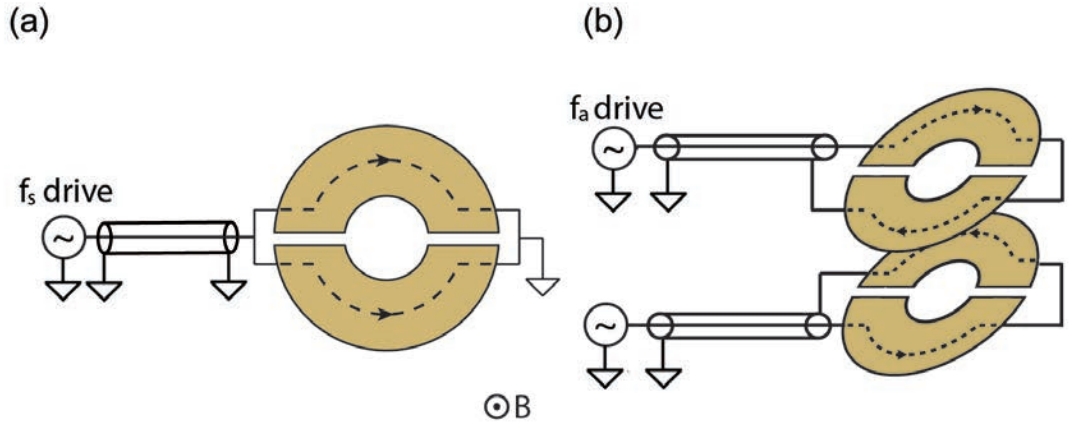

Figure 2.18: Current paths used for applying drive fields to the spin (a) and anomaly (b) transitions.

A field applied with a co-rotating component $\mathbf{b} = b_0 \left(\hat{x}\cos(\omega t) + \hat{y}\sin(\omega t)\right)$ (and thus total magnitude $2b_0\sin(\omega t)$) will give a Rabi frequency

$$\Omega_{sf} = \frac{g}{2}\frac{qb_0}{m_p} = \frac{2b_0\mu_p}{\hbar} \tag{2.32}$$

In the presence of a magnetic bottle field, the transition lineshape due to field inhomogeneity and thermal amplitude fluctuations modifies the effective transition rate [1]. Applying a weak drive for a time Δt will give a transition probability

$$P = \frac{1}{2}\left(1 - \exp\left[-\pi\Omega_{sf}^2\Delta t\chi_{sf}[\omega]\right]\right) \tag{2.33}$$

where χ is the lineshape parameter [1, 52] discussed in section 3.4. When the drive is applied for a long enough time Δt at a high enough Rabi frequency, the transition probability saturates and approaches $\frac{1}{2}$.

In chapter 6 we propose a measurement method using an anomaly drive as well as a spin-flip drive. The anomaly frequency is the sideband simultaneously driving a spin flip and a cyclotron transition. This requires a magnetic gradient oscillating at the frequency $\omega_{ad} = \omega_s - \omega_c$, giving a field at the particle's radial position $\boldsymbol{\rho}$:

$$\mathbf{b}(t) = b_1\boldsymbol{\rho}(t)\cos(\omega_{ad}t) \tag{2.34}$$

where b_1 is the magnitude of the field. This gradient is applied with counter-rotating loops of current applied to the compensation electrodes, as pictured in figure 2.18 (b). This field gives a Rabi frequency

$$\Omega_a = \frac{qg}{2m_p}\frac{b_1\rho_c}{2} \tag{2.35}$$

This has a similar lineshape to the spin-flip transition; the transition probability after a time Δt also takes the form of equation 2.33, with Ω_a replacing Ω_s.

In both of these cases, we require a relatively large current at a relatively high radio frequency. For the spin-flip drive in the analysis trap, the large magnetic gradient means that $\chi_{sf} \rightarrow \frac{1}{\Delta\omega_{sf}}$ on resonance; with $\Delta\omega_{sf}/2\pi \approx 25$ kHz, a large Rabi frequency is required to saturate the transition in a reasonable time. For the anomaly drive, the Rabi frequency is proportional to the small cyclotron radius ρ_c. To achieve reasonable transition rates, we use tuned-circuit drive lines to maximize the current in the electrodes generating the driving field. The design and implementation of these transmission line resonators are discussed in detail in section 8.5.

Chapter 3

Antiproton Magnetic Moment Measurement

In 2011, the experiment completed the first single-particle measurement of the proton magnetic moment [17]. Having successfully demonstrated a single-particle method, we decided to measure the antiproton magnetic moment with the same apparatus. We had one year in which to install the experiment in the Antiproton Decelerator (AD) at the European Organization for Nuclear Research (CERN), outside Geneva, and make a precision measurement before the antiproton beam would be shut off for accelerator upgrades. We moved the experiment and replicated the Harvard setup in an accelerator hall at CERN; demonstrated antiproton trapping in a smaller trap than used before; and successfully made the first single-particle measurement of the antiproton magnetic moment, improving the precision by a factor of 680 over the previous best measurement.

In this chapter we first describe the methods used to slow and trap antiprotons.

While these methods, developed in the TRAP collaboration [48, 53, 54], have been described previously [55], trapping antiprotons for the 2013 measurement was a non-trivial effort, due to complications associated with using a smaller electrode diameter than any previous antiproton trap.

We then discuss the methods for the parts-per-million measurement of the antiproton magnetic moment, where we used a time-averaged procedure to detect the axial frequency shifts due to a resonant spin-flip or cyclotron drive. We finally discuss the spin and cyclotron lineshapes, before presenting the result of the measurement.

3.1 Trapping and cooling antiprotons

The world's only source of cold antiprotons is the CERN AD. Protons from the CERN Proton Synchrotron collide with an iridium target, converting their kinetic energy to mass and creating antiprotons (among other particles). These antiprotons are magnetically guided to the AD, where they are slowed by RF fields and collisions with cold electrons. Antiprotons are then delivered to the experiments in bunches of $\sim 3 * 10^7$ with 5 MeV of kinetic energy, and must be slowed and cooled a further 10 orders of magnitude before reaching thermal equilibrium with the trap.

Before entering the trap, the antiprotons are further slowed by interactions with normal matter. As a high-energy antiproton passes through a thin foil or gas cell, it loses energy through elastic collisions via the Coulomb interaction. If an antiproton's energy drops to zero, it will be captured in the foil and annihilate. When material thickness and density is optimally tuned, half the antiprotons will annihilate inside the beryllium degrader; the remainder will enter the trap at low kinetic energy, including

a significant fraction with a few hundred eV or less. The energy loss depends on material properties and thickness, which were chosen using the SRIM[1] code [56]. A gas cell, described in [57] and consisting of an adjustable mixture of SF_6 and He, allows fine tuning of the average kinetic energy loss.

Antiprotons in or around the trap are detected using methods adapted from particle physics. When an antiproton annihilates with a proton, the resulting reaction produces energetic charged pions. These pass through the trap and magnet, and through a scintillating charged-particle detector surrounding the apparatus. This generates a photon signal in the scintillator, which is detected with a photomultiplier tube. This system detects approximately 50% of annihilations in the trap [58], allowing detection of 10 or fewer antiprotons. This was important in optimizing our small trap for consistent antiproton loading. The beam is steered into the trap using a parallel plate avalanche counter [59] – a gas cell containing large-area position-sensitive electrodes which detect ionization in the gas due to the antiproton bunch.

Figure 3.1 shows the sequence for catching antiprotons once they enter the trap. Antiprotons with low kinetic energy are reflected by a 300V potential applied to an electrode partway up the stack (fig. 3.1a). Before they return to the degrader, a 300V pulse is applied to it, closing the potential well (fig. 3.1b). The timing and amplitude of this pulse are chosen to maximize the number of antiprotons trapped. Antiprotons from the AD come in a 500 ns window, preceded by a trigger signal; we use a custom high-voltage switch (described in [60]) to attain sufficient timing resolution to catch the antiprotons. The 300V pulse amplitude is chosen because we see no increase in antiprotons trapped at higher energies.

[1]Stopping and Range of Ions in Matter

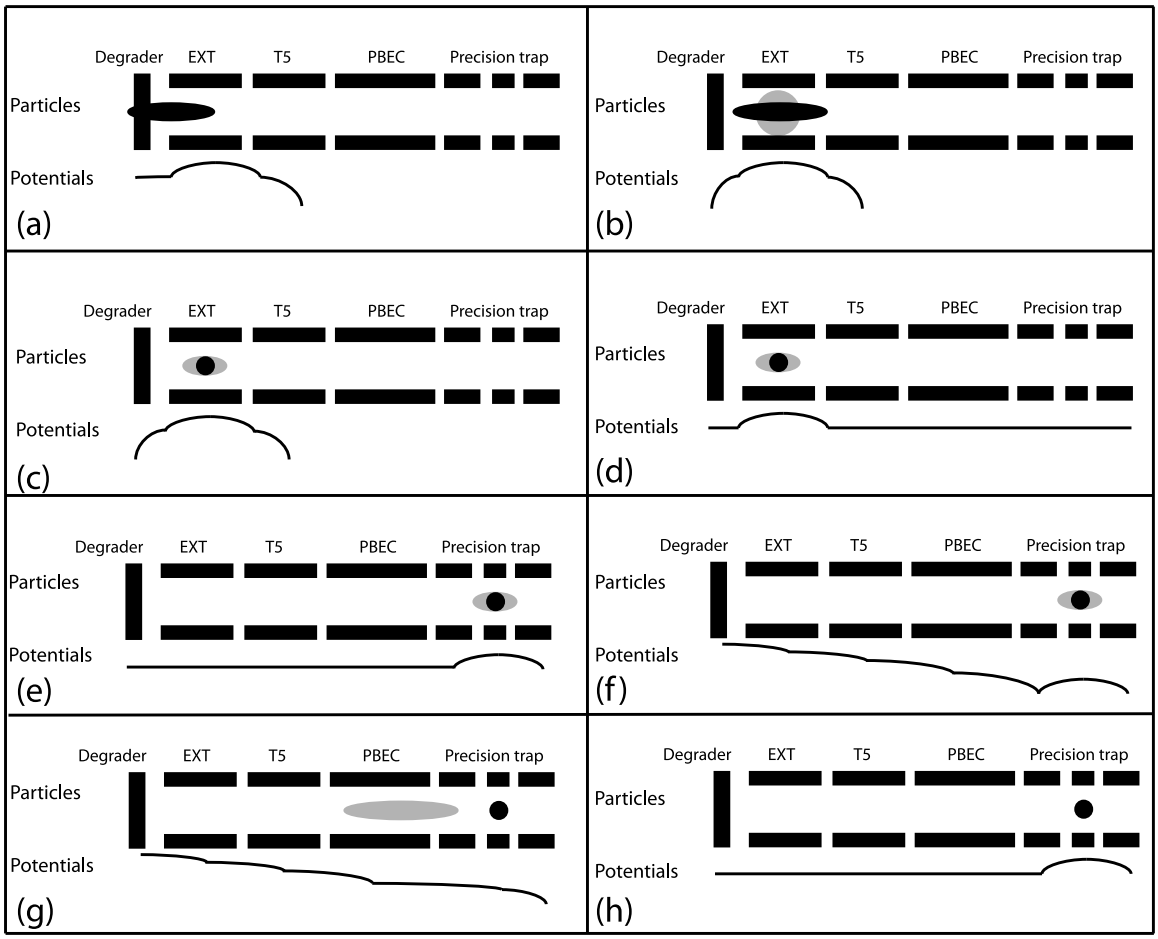

Figure 3.1: Schematic of the process for catching and cooling antiprotons.
(a) Antiprotons (black) enter the trap through the degrader.
(b) A fast switch pulses 300V on the degrader. Electrons (grey) are loaded in a low voltage well with the antiproton shot, or from the FEP.
(c) Antiprotons pass many times through the electron cloud, losing energy.
(d) HV well is switched off.
(e) Antiprotons and electrons are adiabatically transferred to the precision trap. Sideband cooling condenses the electrons at the center of the trap.
(f) A sloped potential is established to the degrader.
(g) RF pulses open the trap long enough for electrons to leave the trap.
(h) Only low-energy antiprotons remain in the trap.

Antiprotons trapped in this HV well still have kinetic energies up to 300 eV. Before they can be trapped in our low voltage well, this must be reduced by a factor of 100; and before they can be used for a measurement, cooling by an additional factor of 10,000 is required. Electron cooling is used to reduce this kinetic energy [48].

Electrons are introduced to the trap during the loading process, either by firing a field emission point or by secondary scattering as the antiprotons pass through the degrader (fig. 3.1b). These electrons are trapped in a smaller potential, nested within the antiproton HV well. The antiprotons pass through the electrons and collide with them, transferring energy to the electrons; collisions between the electrons distribute their energy between all of their motions. The electron cyclotron motion quickly radiates energy via spontaneous emission, keeping the electrons in thermal equilibrium with the black-body background. As the antiprotons pass thousands of times through the electron cloud, they lose their kinetic energy and reach thermal equilibrium with the trap (fig. 3.1c-d). Figure 3.2 demonstrates the progression of electron cooling with antiprotons in the HV well and electrons in a nested low-voltage well.

In other experiments in our group with trapped antiprotons, an HV well several electrodes long has been used to maximize the number of antiprotons trapped while easing the strict pulse timing requirements. A similar long well failed to catch any antiprotons in our low-voltage trap; instead, they annihilated during electron cooling. Only when we shortened the the HV well to a single electrode were we able to reliably trap and cool antiprotons. A shorter HV well also increased the number of high-energy antiprotons trapped.

Our trap is the smallest ever used for antiproton studies, at half the diameter

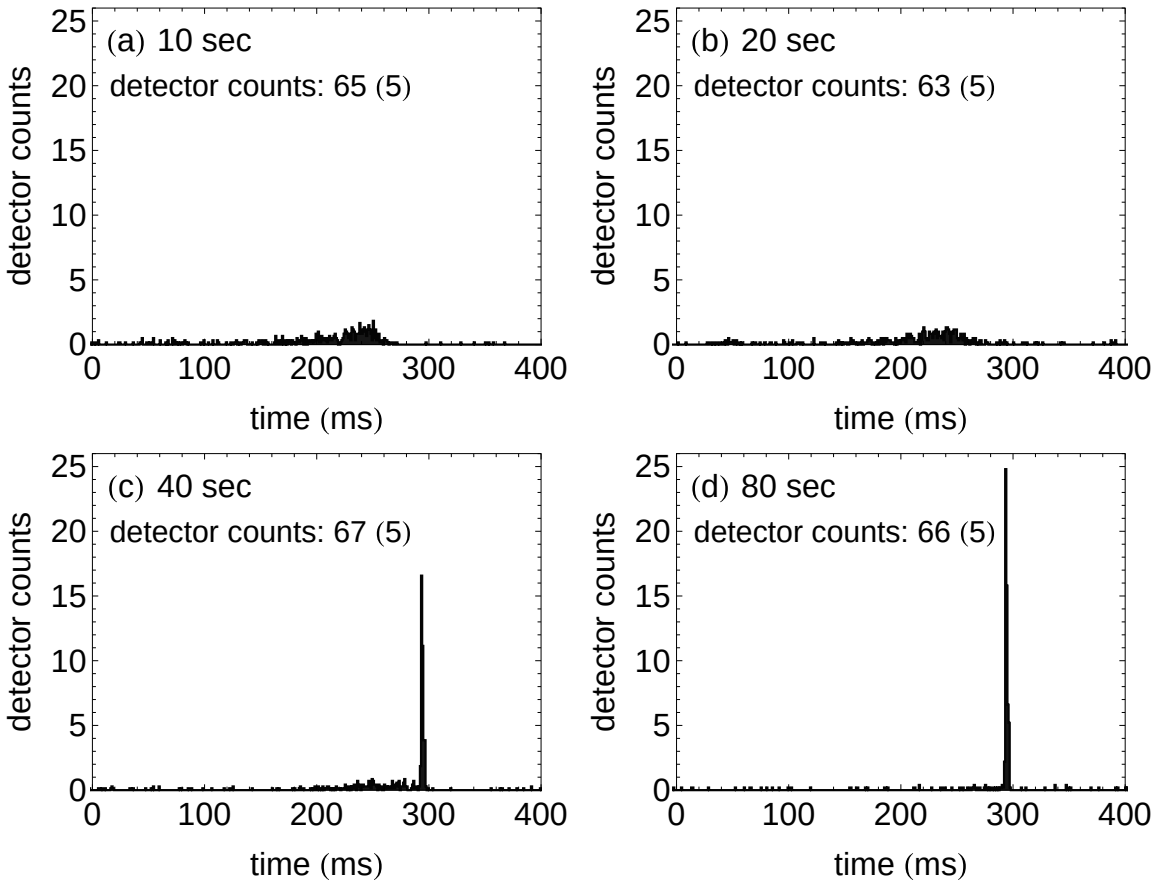

Figure 3.2: Antiproton annihilation counts during linear voltage ramp after different electron cooling times. Antiprotons leaving the trap later have lower energy; after 80 seconds of electron cooling, nearly all antiprotons are cooled to low energy. Figure from [35].

of those used for antiproton q/m measurements. We suspect that the difference in diameter may be responsible for the antiproton loss in the long HV well. While we used a similar length:diameter ratio as previous antiproton experiments, surface and noise interactions scaling with $1/\rho^2$ or higher could be more significant in our trap. For example, interaction of patch and surface potentials with the poorly-defined magnetron motion in the highly anharnomic HV well could cause particle loss, and may scale as fast as $1/\rho^4$ (see section 5.2.1). We therefore trapped antiprotons on an electrode adjacent to the degrader, and transferred them into the precision trap after electron cooling was complete (fig. 3.1e). Figure 3.3 shows the different HV wells used for loading antiprotons; the number of HV antiprotons increased with shorter wells, and only with the shortest well was electron cooling successful.

Once the antiproton cloud is cooled to thermal equilibrium, the electrons must be ejected from the trap. A fast pulse is applied to a neighboring electrode, which briefly opens the potential well to a ramp out of the trap (fig. 3.1f-g). If the pulse is shorter than the antiproton axial period but longer than that of the electron, electrons will leave the well while most antiprotons will remain trapped. The factor of $\sim$2000 difference in q/m ratios gives a large separation between these periods. 8 such pulses were used to ensure that no electrons remained trapped with the antiprotons. Finally, when only antiprotons are left in the precision trap, the antiproton number is reduced to one using techniques described in section 2.5.1. It can then be controlled and measured as a normal proton, using all the techniques described in chapter 2.

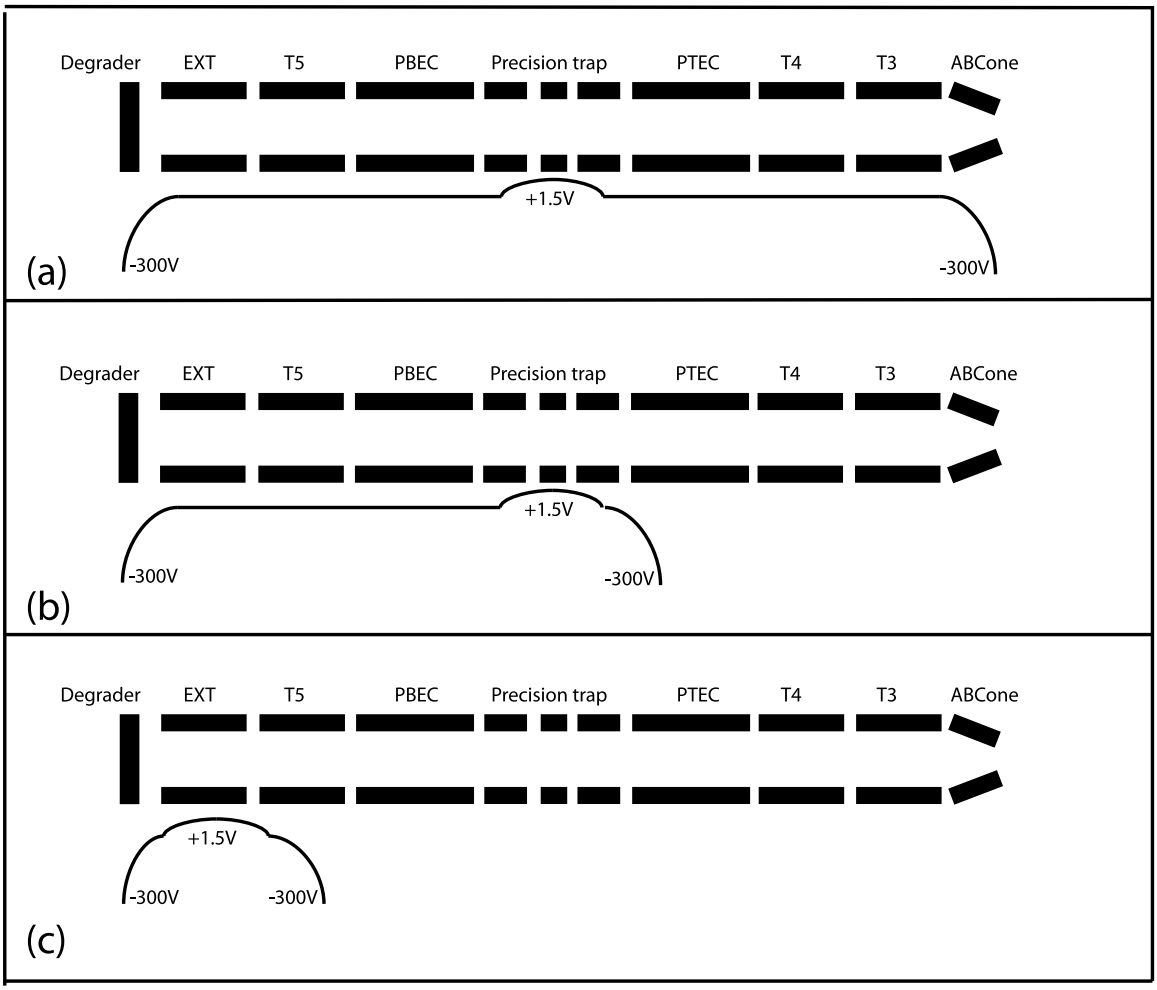

Figure 3.3: HV well configurations used for loading antiprotons. (Potentials not to scale.)
(a) Corresponds most closely to the configuration used in antiproton q/m measurements and antihydrogen experiments. 38±1 HV antiprotons were trapped per shot in this configuration.
(b) Intermediate configuration with a shorter well. 49±2 HV antiprotons per shot were trapped.
(c) Final configuration, which enabled successful antiproton cooling. 56±2 HV antiprotons were trapped per shot.

3.2 2013 measurement methodology

Once a single antiproton was obtained in the precision trap, it was transferred adiabatically[2] into the analysis trap. The magnetic gradient in the analysis trap couples the spin and cyclotron states to the axial frequency, as discussed in section 2.2, allowing us to detect changes in the spin and cyclotron states as changes in the axial frequency. Applying drives near the spin or cyclotron frequencies and measuring the response of the axial frequency thus gives the transition lineshapes. The frequency difference due to a spin flip is only 130 mHz out of ∼1 MHz, so measuring the response to a spin drive requires a stable axial frequency and high measurement resolution.

3.2.1 Measurement procedure in the analysis trap

Figure 3.4 outlines the measurement procedure. The axial frequency is measured before and after applying a near-resonant drive at a frequency under test, giving f_1 and f_2. A detuned drive with identical power is applied between f_2 and f_3, to control for systematics related to the drive strength.[3]

Axial frequency measurements for the spin-flip lineshape were performed using the SEO due to its improved SNR, while dips were used for the cyclotron frequency because f_z and f_+ should be measured under the same conditions to obtain f_c using equation 2.14. (The difference in axial frequencies measured with the SEO and dips

[2]The voltages applied to sequential electrodes are varied slowly compared to the axial and cyclotron frequencies, keeping the proton in the same quantum state as it moves between traps – see section 8.4 for details of the relative electrode and trap positions.

[3]In the proton magnetic moment measurement, the largest systematic effect observed (still well below the measurement precision) was a shift in the axial frequency depending on the applied spin-flip drive power. This systematic was suppressed here with a tuned-circuit drive line (see sec. 8.5). The detuning of the control drive in fig. 3.4 was kept large compared to the transition linewidth but small compared to the drive line bandwidth, to transmit the same power as the resonant drive.

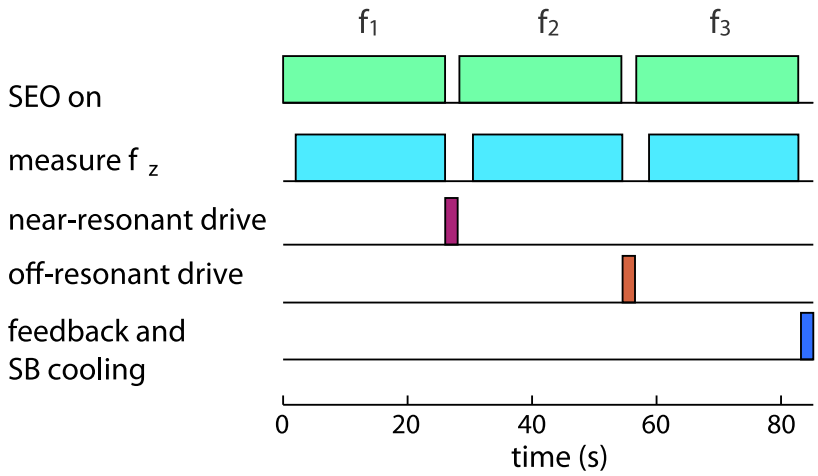

Figure 3.4: Measurement sequence for the 2013 antiproton magnetic moment measurement.

was too small to shift the measurement at ppm precision, but using dips eliminated the possibility of a systematic.)

After the third axial frequency measurement f_3, the proton is sideband cooled to prevent noise-driven growth over time in the magnetron state. Negative feedback applied during sideband cooling reduces the axial temperature and thus the distribution of magnetron states (see sec. 3.4). The entire sequence is then repeated many times at the same drive frequency, to build up a data set of frequency differences. Each frequency differences $f_{2,i} - f_{1,i}$ is recorded as Δ_i^s or Δ_i^c (for the spin or cyclotron), while each frequency differences $f_{3,i} - f_{2,i}$ is recorded as Δ_i^0.

3.2.2 Measuring transition probability via the Allen deviation

Once we have the two sets of frequency differences $\Delta_i^{s,c}$ and Δ_i^0, we compute an Allen deviation for each of them. This lets us time-average the frequency-difference signal and characterize the effect of the spin-flip or cyclotron drive. Used to characterize frequency stability over time, the Allen deviation of the control distribution and of the spin or cyclotron distributions at a drive frequency f is defined as

$$a_0^2 = \sum_{i=1}^{N} \frac{1}{2N} (\Delta_i^0)^2, \qquad a_f^2 = \sum_{i=1}^{N} \frac{1}{2N} (\Delta_i^{s,c})^2 \tag{3.1}$$

The Allen deviation in the presence of spin flips is approximately [35]:

$$a_f \approx \sqrt{a_0^2 + \frac{P}{2}\Delta_s^2} \tag{3.2}$$

where Δ_s is the axial frequency shift due to a spin flip. This equation can be inverted to give spin flip probability as a function of the deviations:

$$P_{sf} = \frac{2(a_f^2 - a_0^2)}{\Delta_s^2} \tag{3.3}$$

The difference in Allen deviations, $a^2 = a_f^2 - a_0^2$, is thus the signal of spin or cyclotron transitions. The uncertainty on this measurement can be estimated, for a Gaussian distribution of control frequency differences, as

$$\sigma(a) = \frac{a}{\sqrt{2N-2}} \tag{3.4}$$

Figure 3.5 shows two sets of axial frequency differences, one for a resonant and one for a non-resonant drive. The fits support the assumption of a Gaussian distribution; the increase in Allen deviation for a near-resonant drive is visible.

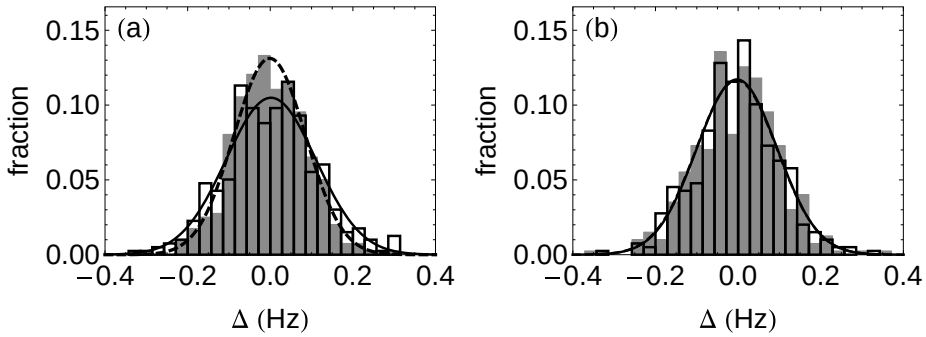

Figure 3.5: Histogram of axial frequency differences, with control drive in grey and spin-flip drive in outline. Dashed and solid lines are Gaussian fits to the control and spin-flip distributions, respectively. (a) shows the increase in standard deviation from a near-resonant drive; (b) shows the overlap between distributions for a non-resonant drive. Figure from [35].

3.2.3 Axial stability and selecting a cold cyclotron state

Equation 3.2 shows that for a saturated spin-flip drive (50% transition probability), the Allen deviation increases only by the addition of $\Delta_{sf}^2/4$ to the control Allen deviation a_0^2. Keeping the control Allen deviation small is thus essential to clearly resolving spin flips. Stability of the axial frequency is discussed in detail in [35] and [34]. Several elements contribute to a_0, including trapping potential instability, the noise floor of the FET, and 1/f noise in the detection chain. During the antiproton magnetic moment measurement, we observed significant axial instability due to ground loops in a neighboring experiment and temperature fluctuations in our electronics cabin.

However, the largest contributor is a changing cyclotron radius in the magnetic bottle. Like changes in the spin state, changes in the cyclotron quantum number n_c couple to the axial frequency via the magnetic gradient. In the presence of electric field noise near ν_c, the rate of cyclotron transitions is a function of n_c; a larger radius means a larger dipole moment interacting with the noise. Experimentally, we find

that a_0 increases proportionally to ρ_c or $\sqrt{n_c}$, as shown in figure 3.6.

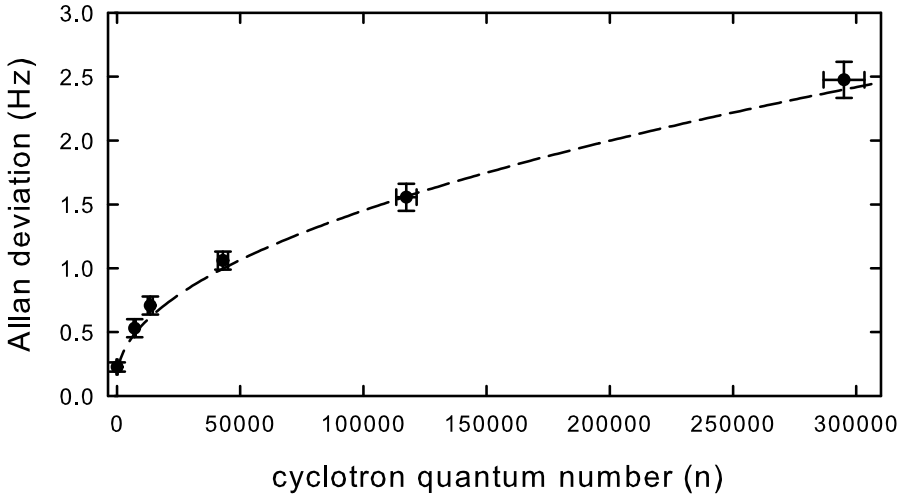

Figure 3.6: Measured control Allan deviation versus cyclotron quantum number, increasing as $\sqrt{n}$. Figure from [35].

The cyclotron motion in the analysis trap is decoupled from a resonator, and changes only in response to noise or drives. In the precision trap for this measurement, the cyclotron motion was in thermal equilibrium with a classical noise bath in its resonator. Therefore, every time the particle is transferred out of the precision trap, a different cyclotron state is selected from a Boltzmann distribution. This cyclotron state can then be measured via the axial frequency in the analysis trap. By re-selecting the state multiple times, we can obtain a proton with a low cyclotron energy (and thus high axial stability). We therefore set a threshold for the measured cyclotron energy, above which the proton is returned to the precision trap. In this measurement, the chosen threshold was $\Delta\nu_z \leq 100$ Hz, or a quantum number $n_c \sim 2000$.

3.3 Magnetic field stability

The axial stability demonstrates some of the issues with performing a precision measurement in an accelerator hall – it was affected by neighboring experiments and imperfect control over experimental conditions. While high axial stability is required to identify spin-flips, magnetic field stability is also important, since it affects the transition linewidth. The AD's basic principle of operation relies on changing magnetic fields – as the antiprotons slow down, the confining fields must decrease to keep them in the beam line. Using a cyclotron decay of a single antiproton, we identified the effect of the AD cooling cycle on our trapping field (figure 3.7). Fortunately, it shifted the field by significantly less than the measurement precision.

3.4 Lineshapes

The spin and cyclotron lineshapes originate from the magnetic gradient due to the iron ring. This gradient (equation 2.15) modifies the homogeneous background field with a term proportional to $z^2 - \frac{\rho^2}{2}$. Changes in either z^2 or ρ^2 therefore change the spin and cyclotron frequencies. The axial frequency is much larger than the transition rate and the linewidth, so the different magnetic fields sampled over the course of an axial oscillation average out; the average field throughout the oscillation determines the frequency. Changes in the axial oscillation amplitude do shift the transition frequency, as the particle experiences a different average field. Additionally, when the particle is sideband cooled the axial temperature is mapped onto the magnetron radius, which couples to the ρ^2 term in B_2.

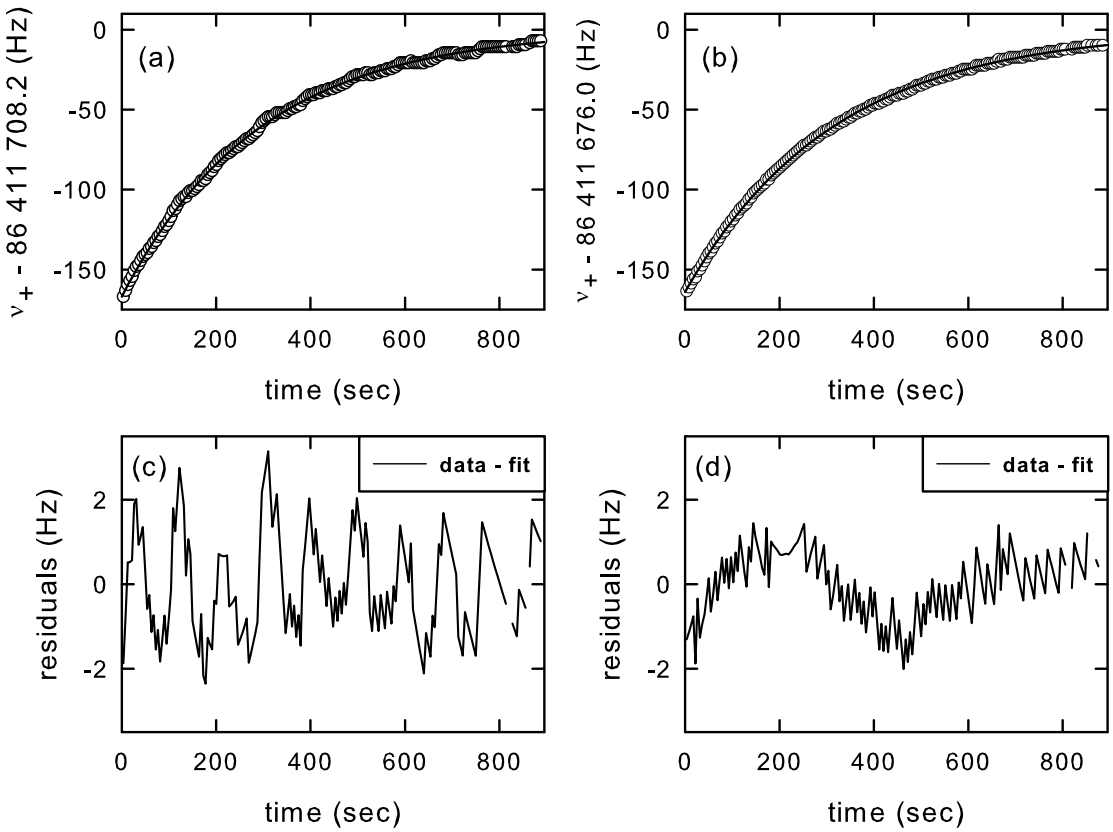

Figure 3.7: (a) and (b): Antiproton cyclotron decays with the AD cycle on and off, respectively. (c) and (d): Residuals for these decays. (c) shows a shift of 35 PPB in magnetic field due to the AD cycle. Figure from [35].

Considering the average axial amplitude $\langle z^2 \rangle$, the transition frequency is shifted from a zero-amplitude value by

$$\omega = \omega_0 \left(1 + \frac{B_2}{B} \langle z^2 \rangle \right) \tag{3.5}$$

Following [1], we identify a linewidth parameter $\Delta\omega_z$ due to the axial temperature distribution in the magnetic bottle:

$$\Delta\omega_z = \omega_0 \frac{B_2}{B} \langle z^2 \rangle = \omega_0 \frac{B_2}{B} \frac{k_b T_z}{m\omega_z^2} \tag{3.6}$$

The lineshape depends on the relative size of $\Delta\omega_z$ and other relevant timescales, and is evaluated for various conditions in [1, 52]. When $\Delta\omega_z \gg \gamma_z^{-1}$, the axial motion is loosely coupled to the thermal bath, and remains unchanged on the timescale of the transition linewidth $1/\Delta\omega_z$. The line profile is then a δ function for each possible amplitude, averaged over the Boltzmann distribution of axial amplitude states:

$$\chi(\omega) = \frac{\Theta(\omega - \omega_0)}{\Delta\omega_z} \exp\left[-\frac{\omega - \omega_0}{\Delta\omega_z} \right] \tag{3.7}$$

where Θ is the Heaviside step function. This distribution has a sharp lower edge, with no probability at negative amplitudes and the maximum probability at zero axial amplitude (as zero energy is the most probable state in the Boltzmann distribution). The axial Boltzmann distribution is also mapped onto the magnetron radius by sideband cooling, giving a similar lineshape, with a factor of $-1/2$ from the bottle term, and a reduced temperature $T_m = \frac{\omega_m}{\omega_z}$ from the sideband cooling limit. The sign change puts the exponential tail of the magnetron Boltzmann distribution below the edge of the axial distribution, while the other factors give the magnetron lineshape parameter

$$\Delta\omega_m = \frac{2\omega_m}{\omega_z} \Delta\omega_z \tag{3.8}$$

A convolution of the axial and magnetron distributions gives a combined lineshape

$$\chi(\omega) = \frac{1}{\Delta\omega_z + \Delta\omega_m}\left(\Theta(\omega_0 - \omega)\exp\left[-\frac{\omega_0 - \omega}{\Delta\omega_m}\right] + \Theta(\omega - \omega_0)\exp\left[-\frac{\omega - \omega_0}{\Delta\omega_z}\right]\right) \tag{3.9}$$

This is plotted in figure 3.8 for distributions with an 8K axial temperature and a 4K sideband-cooled magnetron temperature, as measured in 2013. (The magnetron temperature is lower because of negative feedback applied during sideband cooling.)

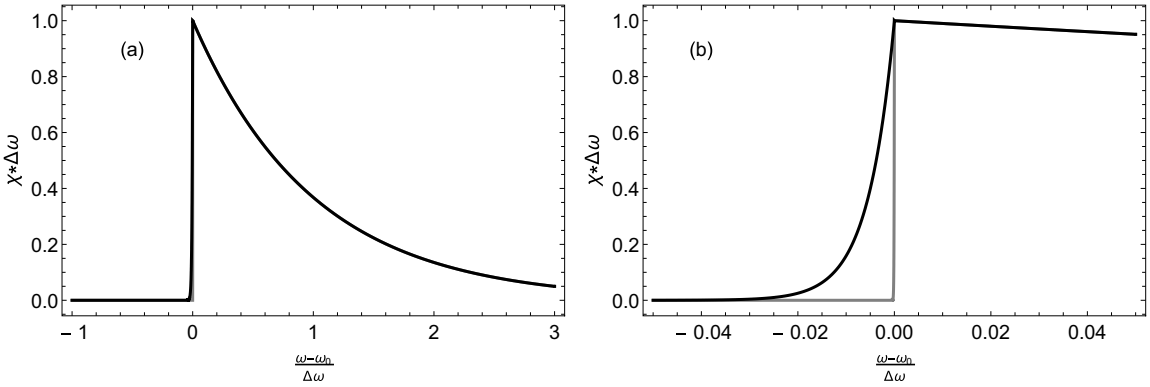

Figure 3.8: (a) Calculated lineshape for the antiproton magnetic moment measurement. (b) Detail highlighting the magnetron distribution.

3.5 Antiproton magnetic moment measurement

To measure the antiproton magnetic moment, the experimental sequence of figure 3.4 was repeated over the course of between 24 and 48 hours at each drive frequency for the spin lineshape, and 2-4 hours for the cyclotron lineshape. Figure 3.9 shows the measured lineshapes. The cyclotron lineshape can be resolved more easily because it is not a two-level system – multiple transitions can be driven during the period between the axial frequency measurements f_1 and f_2.

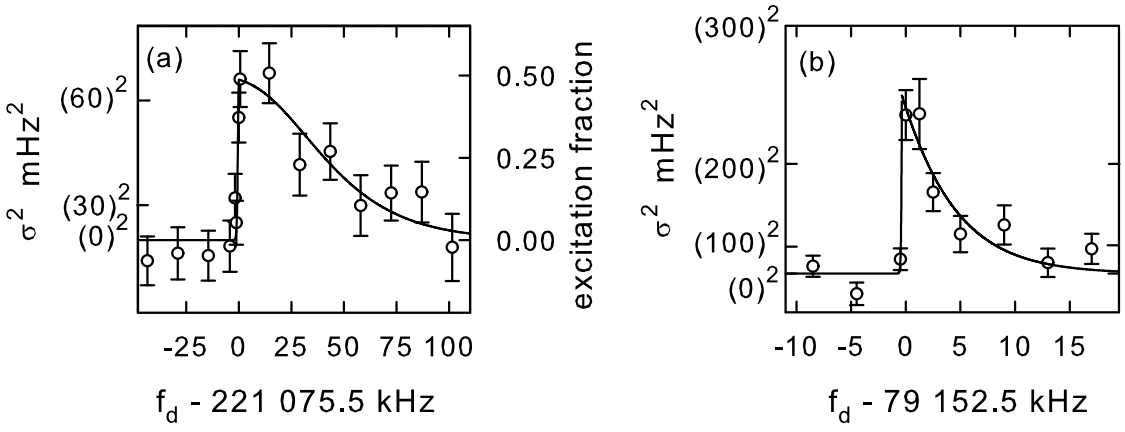

Figure 3.9: Data for the antiproton magnetic moment measurement.
(a) The spin-flip lineshape. (b) The cyclotron lineshape.

While the data show good agreement with the predicted lineshape (fig. 3.8), the reported result includes no fit; the solid lines in figure 3.9 are solely a guide. The sharp increase in transition rate was taken to be the zero-energy peak of the Boltzmann distribution at the resonant frequency, with an uncertainty equal to half the separation between high- and low-probability points. This is the leading uncertainty in the measurement. Broadening due to the distribution of magnetron states after cooling is added as an additional uncertainty. The free-space cyclotron frequency is determined from f_+, f_z, and f_m using equation 2.14; uncertainties in those measurements contribute below .1 ppm. Relevant uncertainties are given in table 3.1.

As a result of this work, we reported the first single-particle measurement of the antiproton magnetic moment [6]. The result was a measurement with 4.4 ppm fractional precision:

Table 3.1: Measurement uncertainties for the antiproton magnetic moment

Frequency	Uncertainty	ppm
Spin	Resonance frequency	1.7
Spin	Magnetron broadening	0.7
Cyclotron	Resonance frequency	1.6
Cyclotron	Magnetron broadening	0.7

$$\frac{\mu_{\bar{p}}}{\mu_N} \equiv \frac{g_{\bar{p}}}{2}\frac{q_{\bar{p}}/m_{\bar{p}}}{q_p/m_p} \approx -\frac{g_{\bar{p}}}{2} = -\frac{f_s}{f_c} = -2.792\ 845\ (12)^4 \tag{3.10}$$

This represents a factor of 680 improvement over the previous best measurement of $\mu_{\bar{p}}$ using exotic atoms. Compared to the earlier measurement [17] of μ_p, this gives

$$\frac{\mu_{\bar{p}}}{\mu_p} = -1.000\ 000\ \pm\ .000\ 005 \tag{3.11}$$

consistent with the prediction of CPT symmetry.

[4]The ratio of q/m values is necessary because the nuclear magneton μ_N is defined in terms of the proton charge and mass; the approximation is valid to better than 0.1 ppb.

Chapter 4

Observing Single-Proton Spin Flips

The precision of the antiproton magnetic moment measurement reported in chapter 3 was limited by the spin and cyclotron transition linewidths. The linewidth in the analysis trap is an unavoidable consequence of the magnetic bottle B_2, which is necessary to measure the spin state. The solution is clearly to split the measurement between a low-gradient precision trap and a high-gradient analysis trap. A measurement could be thus performed in two traps using quantum jump spectroscopy [5, 61]. A known spin state would be initialized in the analysis trap, then transferred to the precision trap; a drive applied at a test frequency; and the proton transferred back to the analysis trap, where the spin state is measured to determine if the test drive caused a spin flip. Over many trials at different drive frequencies, this gives the spin-flip transition lineshape.

Determining whether a spin flip drive in the precision trap caused a transition requires the ability to reliably identify the spin state in the analysis trap. Prior to our demonstration in 2013 [7] (and a simultaneous observation by a competing group

[8]), a spin flip of a single trapped proton or antiproton had never been observed. The magnetic moment of the proton is a factor of 660 smaller than that of the electron, while the difference in charge-to-mass ratios implies a nearly 2000-fold smaller motional amplitude at the same energy. Even with a large magnetic gradient, observing a single proton spin flip thus requires high SNR axial detection with excellent frequency stability.

Identifying single spin flips opens a path towards measurements of the magnetic moment at or below the part per billion level. The magnetic bottle linewidth for a measurement in our precision trap is orders of magnitude smaller than in the analysis trap. Two traps were installed in the first generation apparatus, although only the analysis trap was used in the 2013 antiproton measurement. In that apparatus, B_2 in the precision trap was estimated to be $\sim$1 T/m^2, compared to 290,000 T/m^2 in the analysis trap [35]. The new apparatus has been designed for a geometric cancellation of this residual bottle, and the remaining gradient should be of order 0.1 T/m^2 (see sections 8.1.1 and 10.2.4).

In this chapter, we will discuss spin-state identification, including the demonstration of single-spin-flip detection performed in the generation-one apparatus [7]. We then discuss prospects for improving spin-flip detection using similar techniques, and finally present a path to near-perfect spin-state identification using adiabatic fast passage.

4.1 Identifying single spin flips from frequency differences

The procedure for spin-state identification is essentially a single instance of the time-averaged procedure described in section 3.2.1 and figure 3.4. The spin state is identified using the change in axial frequency due to a spin flip Δ_{sf} – see equation 2.16. A spin flip from $|\downarrow\rangle$ to $|\uparrow\rangle$ causes the axial frequency to increase, as the effective z^2 potential well becomes deeper; the opposite happens for $|\uparrow\rangle$ to $|\downarrow\rangle$.

This is complicated by two effects. First, a saturated spin-flip drive causes a change in spin state at most 50% of the time, per equation 2.33. Second, the axial frequency in the absence of spin flips still shifts between measurements, as described by the control or background standard deviation[1] σ_0. If Δ_0 is a frequency difference chosen at random out of a Gaussian distribution with standard deviation σ_0, a spin-flip trial will give a frequency difference $\Delta = \Delta_0$ if no spin-flip took place, or $\Delta = \Delta_0 \pm \Delta_{sf}$ if a spin flip did occur. If the magnitude of σ_0 is comparable to Δ_{sf}, then background fluctuations can either mimic or mask the shift due to a spin flip. In the 2013 antiproton measurement, we averaged over control and resonant distributions to resolve the additional deviation due to spin flips. However, for quantum jump spectroscopy we cannot rely on this averaging, since we need to know the spin state at two specific points in time – before and after applying the precision-trap drive.

The size σ_0 of the background distribution strongly affects the capacity for spin

[1] We note here that in the context of single-spin-flip identification, the standard deviation, rather than the Allen deviation, is the significant quantity. This is because we are interested in individual frequency differences chosen at random out of a Gaussian distribution, not in characterizing the stability over time of the axial frequency. Standard deviations are represented as σ while Allen deviations are represented as a.

state identification. If σ_0 is much smaller than Δ_{sf}, every individual spin flip can be identified from the frequency difference after a drive attempt in the analysis trap; unsuccessful attempts will result in small frequency shifts. If on the other hand σ_0 is comparable to Δ_{sf}, individual spin flips can only be clearly identified when Δ_{sf} and the individual background fluctuation Δ_0 are in the same direction – otherwise, a small frequency difference could either be due to a small Δ_0 and no spin flip, or a large Δ_0 canceling some of Δ_{sf}. Figure 4.1 shows the distributions of possible frequency differences for different background distribution widths, demonstrating the difficulty in identifying a spin state from a spin-flip trial with a large background.

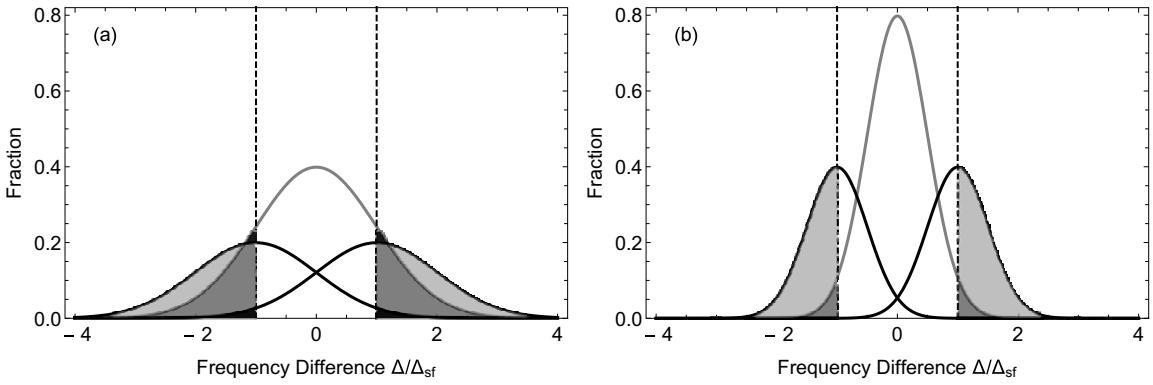

Figure 4.1: Distributions of frequency differences for two background widths, with a 50% probability of a spin flip and unknown initial spin state. Grey line is Δ when no spin flip takes place, black lines are Δ for a spin flip. Vertical dashed lines represent a threshold $\Delta = \Delta_{sf}$. Light gray regions are correctly identified spin states; dark gray regions contain correctly and incorrectly identified spins; black regions contain misidentified spins; and unshaded regions represent no spin assignment.
(a): $\sigma_0 = \Delta_{sf}$. (b): $\sigma_0 = \Delta_{sf}/2$.

With a background distribution σ_0 comparable to Δ_{sf}, spin states can be identified by setting a frequency threshold Δ_t. A frequency difference greater than the threshold,

$|\Delta| > \Delta_t$, is identified as a spin flip, with the initial and final spin states identified by the sign of Δ. In figure 4.1, dashed lines represent a threshold set to $\Delta = \Delta_{sf}$. In both cases, some spin-flips are not identified (unshaded regions), and some are misidentified (dark gray and black regions).

To quantify the ability to identify the spin state for given experimental values of σ_0 and Δ_{sf}, we define the efficiency E, inefficiency I, and fidelity F. We define the efficiency as the percentage of trials in which the initial spin is correctly identified. The inefficiency is the percentage of trials with an above-threshold event which leads to an incorrect assignment of the initial spin state. We define the fidelity as the percentage of identified initial spin states which are correctly identified. To calculate these, we write an expression for the probability of observing $\Delta > \Delta_t$ given a certain initial spin state, with a spin flip up, with no spin flip, or with a spin flip down:

$$\begin{aligned} P_{\downarrow\uparrow}(\Delta_t) &= P_{sf} \int_{\Delta_t}^{\infty} G(\Delta - \Delta_s, \sigma_0) d\Delta \\ P_{\uparrow\uparrow}(\Delta_t) = P_{\downarrow\downarrow}(\Delta_t) &= (1 - P_{sf}) \int_{\Delta_t}^{\infty} G(\Delta, \sigma_0) d\Delta \\ P_{\uparrow\downarrow}(\Delta_t) &= P_{sf} \int_{\Delta_t}^{\infty} G(\Delta + \Delta_s, \sigma_0) d\Delta \end{aligned} \tag{4.1}$$

where $G(\Delta, \sigma_0)$ represents a normalized Gaussian with mean 0 and deviation σ_0 at position Δ. In figure 4.1, $P_{\downarrow\uparrow}(\Delta_t)$ represents the shaded area under the black curve, while $P_{\uparrow\uparrow}(\Delta_t) + P_{\downarrow\downarrow}(\Delta_t)$ gives the shaded area under the gray curve. The efficiency E is then defined as the probability of correctly identifying the initial spin. We also define an inefficiency I as the probability of incorrectly identifying the initial spin[2].

[2]We focus on the identification of the initial spin using frequency differences because we are interested in determining the result of a test drive in the precision trap. Correctly identifying the final state is important for initializing the spin before the trial, but the initialization procedure can be repeated to attain a desired result without loss of information; after the precision-trap drive trial, failed spin identification risks losing information about the trial outcome.

For an identification that the spin was down at the start of the trial, this gives

$$\begin{aligned} E &= P_{\downarrow\uparrow}(\Delta_t) + P_{\downarrow\downarrow}(\Delta_t) \\ I &= P_{\uparrow\downarrow}(\Delta_t) + P_{\uparrow\uparrow}(\Delta_t) \end{aligned} \tag{4.2}$$

The fidelity F, or the probability that a spin assignment once made is correct, is then

$$F = \frac{E}{E+I} \tag{4.3}$$

We can calculate these parameters for the examples in figure 4.1. For the case $\sigma_0 = \Delta_{sf}$, we have an efficiency of 33% and a fidelity of 78% – we will make a spin state identification in 33% of trials, but it will only be correct 78% of the time. For the case $\sigma_0 = \Delta_{sf}/2$, we have a slightly lower efficiency, as fewer of the non-spin-flips are falsely identified; however, the fidelity is significantly improved. We have an efficiency of only 26%, but a fidelity of 96%. The fidelity can be improved at the cost of efficiency by increasing the threshold or by reducing σ_0; and with a smaller σ_0, a lower Δ_t can be chosen to improve efficiency without sacrificing fidelity. Figure 4.2 shows the evolution of these parameters with Δ_t and σ_0.

4.2 Experimental demonstration of single spin flip identification

In 2013, we experimentally demonstrated spin-state identification using this threshold method [7]. Figure 4.3 shows 15 hours of repeated axial frequency measurements, with a resonant spin-flip drive applied followed by a control drive.

The significantly increased scatter for the resonant drive in figure 4.3(b) indicates the presence of spin-flips. This is visible as a broadening in figure 4.4, which shows

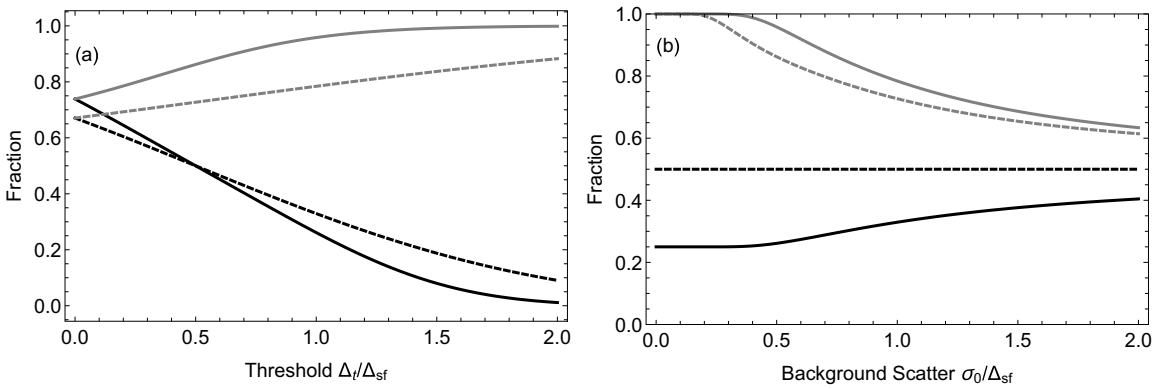

Figure 4.2: Efficiency and fidelity as a function of threshold and background scatter.
(a) Efficiency (black) and fidelity (gray) as a function of threshold Δ_t, with $\sigma_0 = \Delta_{sf}/2$ (solid lines) and $\sigma_0 = \Delta_{sf}$ (dashed lines).
(b) Efficiency (black) and fidelity (gray) as a function of background scatter σ_0, with $\Delta_t = \Delta_{sf}$ (solid lines) and $\Delta_t = \Delta_{sf}/2$ (dashed lines).

the data as a histogram. By defining a threshold we can identify individual frequency differences as spin flips. Because this data has a control standard deviation $\sigma_0 = 63$ mHz, close to half the spin-flip size $\Delta_{sf} \approx 130$ mHz, the efficiency and fidelity for a given threshold is very similar to the solid lines in figure 4.2 (a). For a threshold set at $\Delta_t = \Delta_{sf}$, we obtain an efficiency of 26% and a fidelity of 96%. Applying the threshold to individual frequency differences thus reliably identifies the spin state approximately one out of every four trials. Figure 4.5 shows a subset of the data with individual spin assignments for the $\Delta_t = \Delta_{sf}$ threshold.

To confirm that we are really seeing individual spin flips, we can examine correlations between adjacent spin states. Since the spin is a two-level system, two adjacent spin flips cannot be in the same direction. If above-threshold events indicate individual spin flips, then adjacent above-threshold events should have opposite sign (ie, $|\uparrow\rangle \rightarrow |\downarrow\rangle$ should be followed by $|\downarrow\rangle \rightarrow |\uparrow\rangle$ or vice versa). We therefore isolate

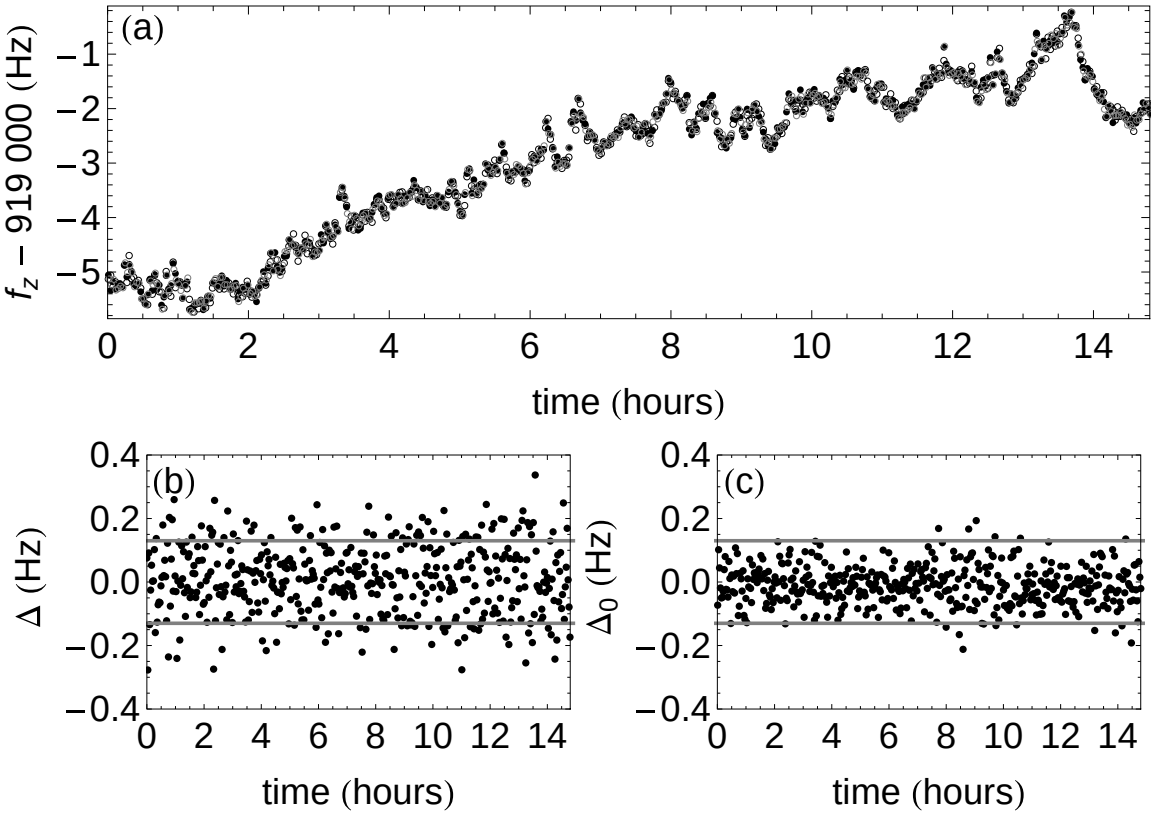

Figure 4.3: Axial frequency measurements used to demonstrate single spin flip detection. (a) Time series of frequency measurements. The overall drift is large compared to the size of a spin flip, but spin-state identification relies on frequency differences. (b) Time series of frequency differences Δ with a resonant spin-flip drive applied. (c) Time series of control frequency differences Δ_0 used to characterize background scatter σ_0. Gray lines show a threshold at the spin flip shift $\Delta_t = \Delta_{sf}$.

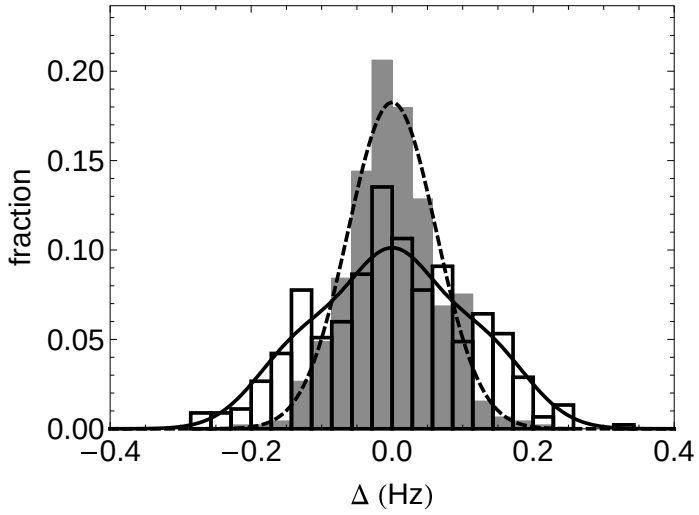

Figure 4.4: Histogram of frequency differences in the demonstration of single-spin-flip detection. The gray histogram shows control frequency differences Δ_0, with a fit (dashed line) to a Gaussian with standard deviation $\sigma_0 = 63$ mHz. The open histogram shows frequency differences Δ with a saturated spin-flip drive; the solid line is the sum of all three Gaussians of figure 4.1, calculated for standard deviation $\sigma_0 = 63$ mHz and $\Delta_{sf} = 130$ mHz.

only events where two adjacent differences are above the threshold, and examine the correlation $\Delta_2 - \Delta_1$. For real spin flips, this should give $\Delta_2 - \Delta_1 \approx \pm 2\Delta_{sf}$.

With an efficiency of 26%, adjacent spin states should be identified about $(.26)^2 \approx$ 7% of the time, or ∼30 pairs out of the 450 events included in figure 4.3. We in fact see 28, shown in histogram in figure 4.6(a). Because of the 96% fidelity, some of these will include misidentified spins – for each pair of identifications, there is a (.96)(.04)+(.04)(.96)≈8% chance that one of them is incorrect, giving a correlation $\Delta_2 - \Delta_1 \approx 0$ (as well as a (.04)(.04)≈ 0.2% chance that both are incorrect). Out of the thirty pairs of adjacent spin-state identifications, we thus expect $\sim$ 28 to have $\Delta_2 - \Delta_1 \approx \pm 2\Delta_{sf}$, and $\sim$ 2 to have $\Delta_2 - \Delta_1 \approx 0$. As figure 4.6(a) shows, we in fact have 25 on the wings, and 3 at the center. Error bars for the expected

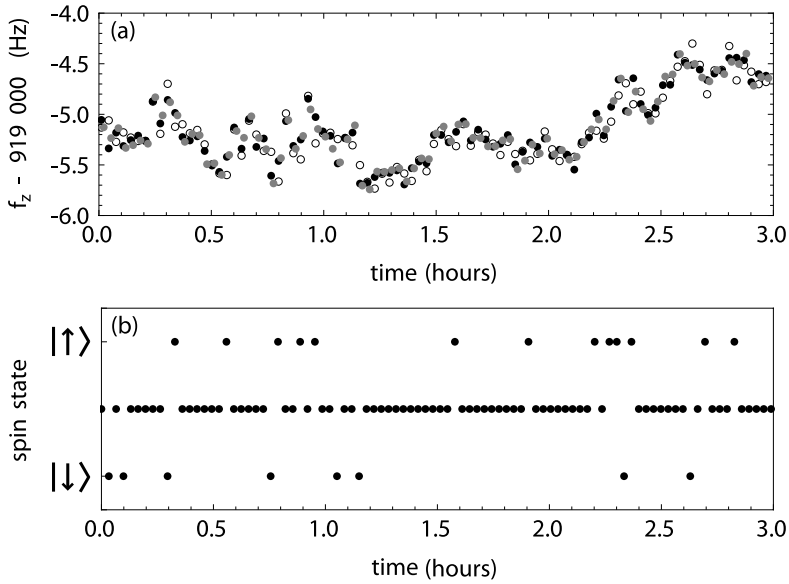

Figure 4.5: (a) Three-hour subset of frequency shifts from figure 4.3. Grey lines are the spin-state identification threshold, at $\Delta_t = \Delta_{sf}$. (b) Spin state assignments following that threshold. Points in the center are frequency differences $\Delta < \Delta_t$, for which no spin state was assigned.

correlations are generated by simulating sets of 450 differences many times, starting with a hidden initial spin state, with a 50% spin-flip probability at each step, and with each frequency difference chosen at random from the appropriate Gaussian in figure 4.1(b). One such simulation is shown in figure 4.6(b). The derived error bars give a range 30 ± 7 events with $\Delta_1 - \Delta_2 \approx \pm 2\Delta_{sf}$ and 2 ± 2 events with $\Delta_2 - \Delta_1 \approx 0$, consistent with our observations. We therefore conclude that we have observed individual spin-flips of a single proton.

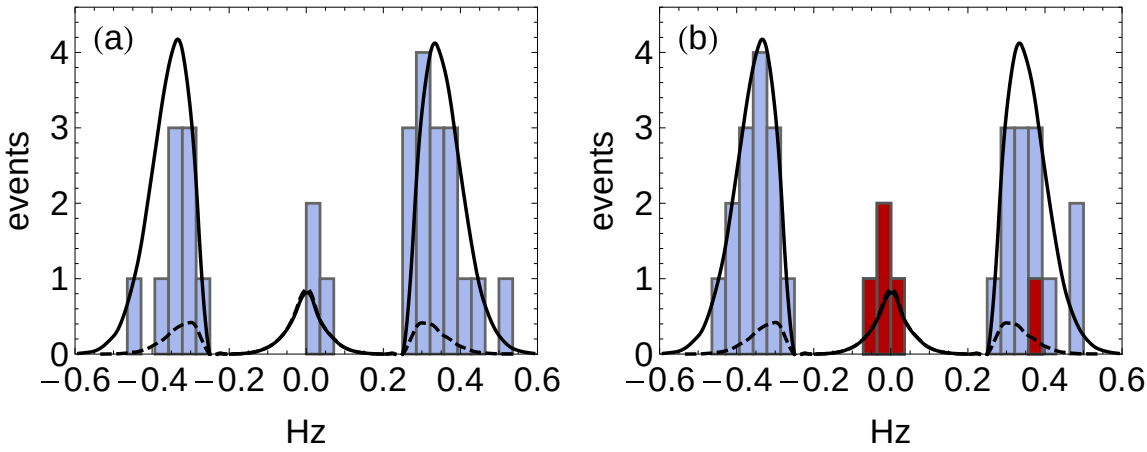

Figure 4.6: Correlations $\Delta_2 - \Delta_1$ for the single spin flip demonstration. The solid curve is the total predicted distribution for adjacent spin-state identifications, derived from simulating many sets of 450 frequency differences. The dashed curve is the prediction of correlations where one or more spin-state identifications is incorrect. (a) shows real correlations for the data set of figure 4.3. (b) shows one of many sets of simulated correlations, used to derive the predicted distribution and error bars. Accidental events are highlighted.

4.3 Improving spin-state detection for quantum jump spectroscopy

The demonstrated spin-state identification could be used to perform quantum jump spectroscopy and measure the g-factor in the precision trap. The spin state could be initialized using frequency differences in the analysis trap. The proton would then be transferred to the precision trap, where a spin-flip drive would be applied at a test frequency. The proton, now in an unknown spin state, would then be returned to the analysis trap, where its spin could be read out again using frequency differences.

However, the low efficiency poses a significant problem for this measurement. The initial spin could be determined reliably by setting a stricter threshold with near-unity

fidelity and low efficiency, and making several spin-flip attempts before one meets that threshold. An unidentified spin state after the precision trap drive, however, would render the trial useless (and a misidentified spin state would actively confound the experiment). The 26% efficiency required for a 96% fidelity implies that three out of every four trials would have to be discarded. Therefore, improving spin-state identification is a priority for single-spin-flip measurements of g.

4.3.1 Reduced background fluctuations and multiple trials

A straightforward way to improve spin-state identification is to reduce the background distribution σ_0. A saturated drive will flip the spin only 50% of the time, limiting the possible efficiency. However, a sufficiently low σ_0 would allow identification of trials when the spin does not flip; we could then make multiple trials to get a higher efficiency. For example, figure 4.7 shows the distributions for $\sigma_0 = \Delta_{sf}/4$, representing a factor-of-two improvement over the axial stability demonstrated in [7]. The three distinct histograms allow us not only to say that an above-threshold event is likely a spin flip, but also that a below-threshold event is likely not a spin flip.

We could then apply multiple consecutive spin-flip trials to increase the efficiency. After one spin-flip trial in the analysis trap with threshold Δ_t, we have E, I, and F as defined in equations 4.2 and 4.3. For multiple spin-flip attempts, we also have to account for spin-flips that have below-threshold frequency differences. The probability that a spin flip does take place and is not identified is N, the Nonidentification rate:

$$N = P_{sf} \int_{-\Delta_t}^{\Delta_t} G(\Delta - \Delta_s, \sigma_0) d\Delta = (P_{sf} - P_{\downarrow\uparrow}(\Delta_t) - P_{\uparrow\downarrow}(\Delta_t)) \tag{4.4}$$

The probability that a spin flip does not take place and no spin-flip is identified is

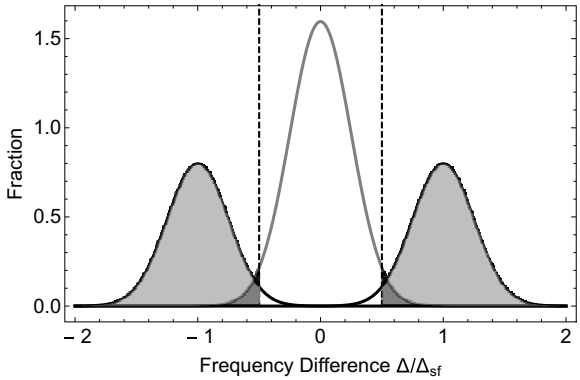

Figure 4.7: Distribution of axial frequency differences with hypothetical improved stability, $\sigma_0 = \Delta_{sf}/4$.

the Unflipped/Unidentified rate

$$\begin{aligned} U =& (1 - P_{sf}) \int_{-\Delta_t}^{\Delta_t} G(\Delta, \sigma_0) d\Delta \\ =& (1 - P_{sf}) \left(1 - \int_{-\infty}^{-\Delta_t} G(\Delta, \sigma_0) d\Delta - \int_{\Delta_t}^{\infty} G(\Delta, \sigma_0) d\Delta \right) \\ =& (1 - P_{sf} - P_{\downarrow\downarrow}(\Delta_t) - P_{\uparrow\uparrow}(\Delta_t)) \end{aligned} \tag{4.5}$$

Note that under these definitions we maintain the normalization E+I+N+U=1.

Given these definitions of N and U, if we apply a second trial after measuring a below-threshold frequency difference, the initial spin will be correctly assigned with a rate $U * E + N * I$; incorrectly assigned with a rate $N * E + U * I$; and remain unassigned with a rate $(U + N)^2$. In general, when the nth spin-flip attempt gives an above-threshold result, we will correctly identify the final spin with the same efficiency E and inefficiency I. A correctly identified final spin will give the correct initial spin when an even number of non-identified spin-flips took place, and the incorrect initial spin when there were an odd number of non-identified spin-flips. When we incorrectly identify the final spin, we still correctly identify the initial spin if there was an odd number of non-identified spin-flips. Repeating the process up to n times, and stopping

if an above-threshold event is measured, gives the n-trial efficiency and inefficiency

$$\begin{aligned} E_n &= \sum_{j=0}^{n-1} \left(E \sum_{k=0,\text{even}}^{j} \binom{j}{k} N^k U^{j-k} + I \sum_{k=1,\text{odd}}^{j} \binom{j}{k} N^k U^{j-k} \right) \\ I_n &= \sum_{j=0}^{n-1} \left(I \sum_{k=0,\text{even}}^{j} \binom{j}{k} N^k U^{j-k} + E \sum_{k=1,\text{odd}}^{j} \binom{j}{k} N^k U^{j-k} \right) \end{aligned} \tag{4.6}$$

When we incorporate the probability of not having any above-threshold events in n trials, $(U+N)_n = \sum_{j=0}^{n} \binom{n}{j} N^j U^{n-j}$, the probabilities obey the normalization $E_n + I_n + (U+N)_n = 1$. We can define the fidelity after up to n trials as $F_n = E_n/(E_n+I_n)$.

Figure 4.8 shows the efficiency and fidelity for multiple spin-flip trials. For the reported axial stability of $\sigma_0 = \Delta_{sf}/2$, the efficiency increases above 50%, but the fidelity falls with repeated trials, yielding little net gain. However, for an improved stability $\sigma_0 = \Delta_{sf}/4$, we have $E_n \to 96\%$ and $F_n \to 96\%$ with repeated spin-flip trials and a threshold $\Delta_t = \Delta_{sf}/2$. Two noteworthy properties of this calculation are demonstrated in figure 4.8: first, as the number of trials increases, the chances of having an above-threshold event at some point approach 1, meaning $E_n + I_n \to 1$ and $F_n \to E_n$ as $n \to \infty$. Second, as we increase the threshold further, the fidelity for multiple trials decreases even though the single-trial inefficiency I goes to zero – this is because the rate of a non-identified spin-flip followed by a correctly identified final spin state (and thus an incorrectly identified initial spin state) increases with the threshold, as more spin-flips are associated with below-threshold frequency differences.

A full factor-of-two improvement in axial stability may or may not be achievable. Changes to improve axial stability in the new apparatus include a new analysis trap design to reduce cyclotron transitions and a new cooling trap allowing improved cyclotron-radius selection (sec. 8.1.2 and [62]) and improved electrode surface treat-

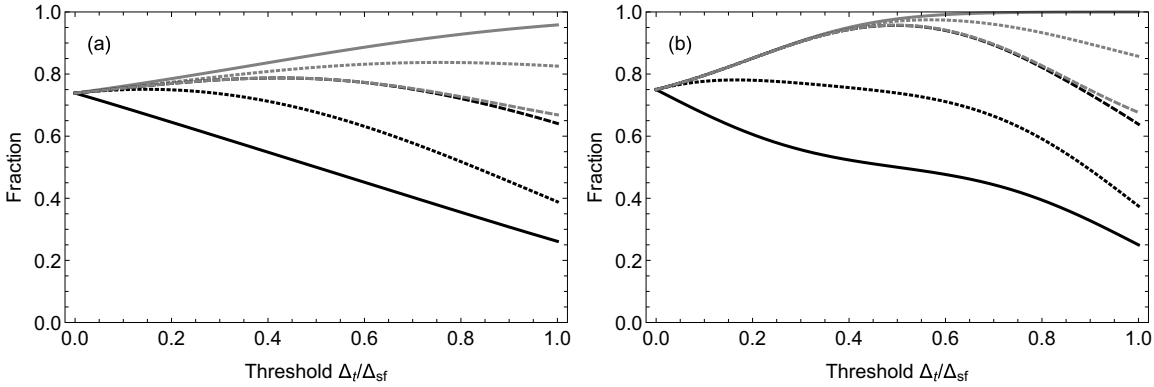

Figure 4.8: Correctly and incorrectly identified spin states versus threshold for repeated spin-flip trials. Black curves are the efficiency E_n, gray curves are the fidelity F_n, for one (solid), two (dotted), and ten (dashed) trials (stopping after measuring an above-threshold event). (a) shows this evolution for $\sigma_0 = \Delta_{sf}/2$, while (b) shows $\sigma_0 = \Delta_{sf}/4$.

ment (sec. 5.2.1 and 8.1.3), but the effect of these changes has not yet been measured.

Using more advanced analysis techniques, such as a Bayesian analysis of repeated spin-flip trials rather than a single threshold, it is possible to gain additional confidence in spin-state identification. However, without significant improvements in axial stability, the overlapping distributions of frequency differences will be a challenge for quantum jump spectroscopy.

4.3.2 Adiabatic fast passage

Improving axial frequency stability improves spin-state identification by better separating the spin-flip and no-spin-flip distributions, but sufficiently good stability could be difficult to achieve. If we could increase the spin-flip probability to nearly 100%, we would eliminate the no-flip distribution altogether. This can be achieved using adiabatic fast passage (AFP), a technique well established in NMR [11] and

previously demonstrated with electrons in a Penning trap [63, 64].

In an AFP experiment, the applied drive frequency ω_{sd} is slowly swept through resonance, at a rate $d\omega_{sd}/dt$. If the adiabatic condition $\Omega_r^2 \gg d\omega_{sd}/dt$ is met, the spin will adiabatically follow the drive, deterministically changing sign [65]. The distribution of axial frequency differences for AFP is shown in figure 4.9. This would allow a lower threshold and higher efficiency without reducing fidelity. For a background distribution $\sigma_0 = \Delta_{sf}/2$ as demonstrated in [7], a threshold at 0 with 100% spin-flip probability gives an efficiency of 100% and a fidelity of 98%.

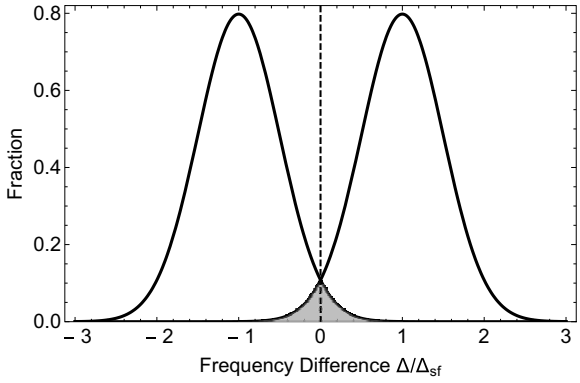

Figure 4.9: Distribution of frequency differences for adiabatic fast passage, with $\sigma_0 = \Delta_{sf}/2$ and $P_{sf} \to 100\%$. Dashed line is a threshold at $\Delta_t = 0$.

The magnetic gradient and the axial temperature distribution combine to challenge AFP in the analysis trap. As long as the sweep rate is much slower than the axial frequency, the spin follows the drive as though it experienced only the average field across an oscillation[3]. However, fluctuations in this average field seen by the particle at a timescale shorter than the sweep rate can interfere with AFP.

The spin lineshape averaged over many trials takes the form of equation 3.9, as dis-

[3]This was confirmed with robust numerical solutions to the equation of motion for an undamped proton in the magnetic bottle.

cussed in section 3.4, but on timescales shorter than $1/\gamma_z$ the lineshape is a Lorentzian whose center frequency depends on the axial amplitude [1]. This amplitude changes as Johnson noise in the axial resonator interacts with the proton. The Johnson noise is effectively uncorrelated, because the physical processes underlying it are much faster than the proton's motion. However, the proton has finite mass; its energy changes smoothly over time as driven by this random process. The rate of change is determined by the coupling between the proton and the noise bath, parameterized by the damping rate γ_z. During an AFP drive sweep, the spin-flip transition frequency will therefore select states from the axial Boltzmann linewidth Δ_z at a rate γ_z. The spin frequency changes at a rate $d\omega_s/dt \propto \gamma_z \Delta_z$. Like the drive sweep rate, $d\omega_s/dt$ must satisfy the adiabatic condition $d\omega_s/dt \ll \Omega_r^2$ to enable coherent population transfer. Meeting this condition for ω_s will be more difficult than for the drive frequency.

Using parameters from the first-generation apparatus, which should be similar in the new apparatus, we have $\Delta\omega_z \approx 2\pi * 24$ kHz and $\gamma_z \approx 2\pi * 2Hz$, giving $\Delta\omega_z * \gamma_z \approx 1.9 * 10^6$ (rad/s)^2. In section 8.5, a detailed calculation of the analysis trap spin-flip Rabi frequency from antiproton data gives $\Omega_{sf} = 410$ rad/s, which would give $\Omega_{sf}^2 \approx 1.7 * 10^5$ (rad/s)^2, violating the adiabatic condition.

Since $\Delta\omega_z$ is set by the bottle field, to implement AFP either γ_z must be reduced or Ω_{sf} increased. γ_z can be decreased to an extent through application of negative feedback (see [35] or sec. 7.2.8). It can be reduced more dramatically by detuning the particle from the amplifier, but this may be difficult to implement – methods to change the amplifier frequency risk reducing its Q, while changing the proton's frequency by shifting the trapping voltage risks increasing axial instability.

A more straightforward way to achieve AFP may be by increasing Ω_r with higher drive power. Results in section 8.5 show that we could increase the drive power by as much as 10dB without an analysis-trap axial frequency shift relevant at the sub-ppb level. An improved drive line promises a further increase in drive current; and we are less concerned with a power shift in the analysis trap since the measurement will take place in the precision trap. Combined with feedback cooling, we may thus be able to meet the adiabatic condition without detuning the particle from the resonator.

4.4 Conclusion

We have discussed a method for identifying the spin state of a single proton or antiproton using frequency differences in a magnetic gradient. In 2013, we reported the first observation of individual spin flips of a single proton using this method [7] (simultaneously reported with [8]). The efficiency and fidelity characterize the performance of spin-state identification. While the 96% fidelity and 26% efficiency demonstrated in our apparatus would be sufficient to perform a very slow magnetic moment measurement, improving spin-state detection beyond that demonstration would significantly improve a potential measurement using quantum jump spectroscopy. Spin-state identification could be improved with improved axial stability and multiple spin-flip trials; however, adiabatic fast passage in the magnetic bottle would allow near-perfect spin-state identification and is straightforwardly achievable with the large drive strength demonstrated in our analysis trap.

Chapter 5

Precision Measurement Methods with Single Spin Flips

The demonstration of spin-flip identification for a single trapped proton opened the road to a precision measurement at the ppb level using using quantum jump spectroscopy in two traps. It was at this point that the focus of my Ph.D. shifted to development of the second generation apparatus and methods. As a goal for the second-generation apparatus, we aimed for a measurement with a 0.2 ppb lineshape FWHM. This represented a further order-of-magnitude improvement over the estimated precision of a measurement in the first-generation apparatus using quantum jump spectroscopy with two traps [34, 35].

In this chapter, we analyze a method for achieving this target precision using single spin flips[1]. We first discuss previous measurements in two traps, and the limitations on their lineshapes and precision. We then present a method for measurement of the

[1]We are grateful to Professor Ed Myers of Florida State University for conversations and collaboration in the development of this method.

cyclotron frequency using separated oscillatory fields. Incorporating a simultaneous interrogation of the spin-flip transition frequency, also using separated oscillatory fields, provides a method for a full sub-ppb measurement of the magnetic moment. Finally, we estimate the amount of time required to make a measurement at our target precision using this method.

5.1 Historical two-trap measurement methods

The spin lineshape can in theory be measured using quantum jump spectroscopy and compared to a cyclotron lineshape, as was done (with the cyclotron and anomaly frequencies) for the electron g-2 measurements. However, a measurement at the sub-ppb level using separate lineshapes is subject to issues which were negligible at the few-ppm level of the 2013 antiproton measurement. For example, even with significant effort towards stabilization, the field in our magnet was observed to drift by 0.08 to 0.5 ppb per hour in a previous measurement of the electron magnetic moment [66]. Over the course of a proton magnetic moment measurement, the magnetic field drift would be much more than the target measurement precision. Rather than measure distinct lineshapes for the spin and cyclotron frequencies, a better solution is thus to associate each spin-flip trial with a full measurement of the cyclotron frequency taken under the same or similar conditions. With multiple trials at different spin-flip drive frequencies, a resonance for the g-factor can be measured. In considering ways to make this type of measurement, we can benefit from the history of g-factor experiments using two traps, especially recent proton and antiproton measurements. We will explore how some such measurements have combined spin and cyclotron measurements.

The two-trap method was originally demonstrated in measurements of the g-factor of the electron bound in highly charged ions [61]. Simultaneously to the application of a spin-flip drive, the cyclotron frequency was directly measured using the image charge in a resonator. This required a ∼3 eV excitation of the cyclotron motion, introducing a relativistic shift which did not limit the precision of [61] but would be over an order of magnitude larger for the proton.

More recently, a measurement of the antiproton magnetic moment by the BASE collaboration [18] used two separately trapped particles to perform interlaced cyclotron and spin-flip measurements. A measurement of the cyclotron frequency using axial-cyclotron sideband coupling [51] was performed before and after each application of the spin-flip drive in the precision trap. Separate particles were used so that the cyclotron radius of the spin-flip proton could be kept small, reducing σ_0, despite the sideband drive increasing the axial and cyclotron quantum numbers doing the cyclotron frequency measurements.[2]

This measurement achieved a 1.5 ppb precision, with a lineshape having a full width at half maximum (FWHM) of ∼32 ppb[3]. This FWHM included a 30 ppb contribution from drive saturation, and a 9 ppb contribution from changes in the magnetic field between the time when ν_c was measured and the application of the spin-flip drive. The contribution from magnetic field fluctuations included short-term changes in the magnetic field and fluctuations in ν_c arising from the temperature distribution during the split-dip measurement. The measurement uncertainty includes

[2]The analysis of the cyclotron-axial sideband drive follows that for magnetron-axial sideband coupling in section 2.6, with the coupled cyclotron temperature being $\frac{\omega_c}{\omega_z}T_z$.

[3]We include the FWHM rather than the Gaussian linewidth σ because we expect data from the methods we propose in the next two chapters will not be normally distributed; the FWHM provides a well-defined way to compare lineshapes.

a 1.1 ppb statistical uncertainty and a 1.0 ppb systematic uncertainty, dominated by a possible change in the axial temperature and frequency during application of the spin-flip drive (see sec. 8.5 for discussion of a similar effect in our 2011 proton magnetic moment measurement and its reduction in our 2013 antiproton measurement).

Subsequently, the BASE collaboration reported a measurement of the proton magnetic moment [19] with an uncertainty of 0.3 ppb and FWHM of ~ 3 ppb. For this result, the cyclotron frequency was measured with sideband coupling while the spin-flip drive was applied in the precision trap. This required re-selection of a cold cyclotron state after every spin-flip trial, leading to a slower experiment cycle. However, this method also eliminated the dominant systematic uncertainty in [18], since the cyclotron and spin frequencies were measured in the same magnetic field.

Simultaneous measurements reduced the linewidth from magnetic field fluctuations, while a weaker drive eliminated saturation. The remaining linewidth arose from instability of the axial frequency measurement using dips, which maps to instability on ν_c. The reported uncertainty was dominated by statistical fitting uncertainty of 0.28 ppb, with .08 ppb systematic uncertainty due to detuning between the particle and axial resonator during the measurement of ν_c, and .03 ppb uncertainty on the relativistic shift due to the high cyclotron energy from the sideband coupling.

5.2 Separated oscillatory fields measurement of ν_c

Simultaneously interrogating the spin and cyclotron frequencies avoids significant systematic (and statistical) uncertainties. However, using axial-cyclotron sideband coupling introduces uncertainty to the cyclotron frequency measurement, which

broadens the linewidth of the g-factor resonance. For a high-precision measurement of the magnetic moment, we propose using a different method to measure ν_c – that of separated oscillatory fields (SOF).

First proposed and implemented by Ramsey [67] with spatially separated fields and an atomic beam, SOF measurements using trapped particles apply two phase-coherent drive pulses separated by some evolution time T_e. A particle's motion or quantum state is initialized with a known phase given by the first drive pulse. During the evolution time, the particle accumulates phase free from the influence of the drive. The second pulse then interferes constructively or destructively with the evolved phase. The interference pattern can be measured as a function of evolution time or drive detuning, and the natural frequency determined with high precision.

The use of separated oscillatory fields for cyclotron frequency measurements in Penning traps has a multi-decade history. In the "pulse and phase" technique [68, 69], the cyclotron motion is excited to a known phase and amplitude with a coherent drive; after a free evolution time, a cyclotron-axial coupling pi-pulse transfers the phase of the cyclotron motion to the axial motion, which is read out using phase-sensitive detection. A variation on this technique, the "pulse and amplify" method [70, 71], uses an amplifying mode-coupling pulse rather than a pi-pulse to reduce the effect of relativistic shifts and other systematics on light particles.

For determining ν_c in a two-trap g-factor measurement, we propose instead following a method described in [72, 73]. The proton is excited from a thermal state to a well-defined cyclotron amplitude and phase with a drive pulse. The cyclotron motion then accumulates phase during the evolution time T_e. A second pulse, phase-coherent

with the first, will interfere with the accumulated phase and either increase or decrease the cyclotron radius. This maps the detuning between the drive and cyclotron frequencies onto the final cyclotron radius ρ_{cf}. The proton can then be transferred to the analysis trap, and ρ_{cf} precisely measured using the axial frequency shift in the magnetic bottle. Figure 5.1 demonstrates a sequence for this measurement.

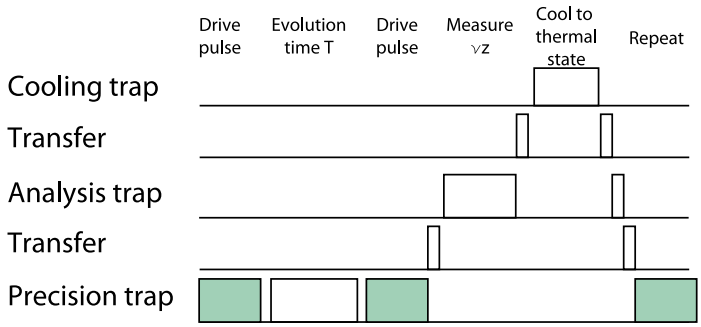

Figure 5.1: Procedure for SOF measurement of cyclotron frequency (time axis not to scale).

To enable free phase evolution during T_e, the particle must be decoupled from the thermal bath in the cyclotron resonator. An FET switch with this function is demonstrated in section 7.4. When cyclotron damping is necessary – for example, returning to a cold state after measuring ρ_{cf} in the analysis trap – the proton can be transferred to the cooling trap, discussed in section 8.1.2, or the FET switch can be opened, re-coupling the particle to the precision trap cyclotron resonator.

Our evaluation of the SOF cyclotron frequency measurement follows reference [73]. The first drive pulse is applied for a time τ, during which the cyclotron motion is excited to a complex amplitude

$$\boldsymbol{\rho}_c = Ae^{i\phi_d(\tau)} + \rho_0 e^{i\phi_0} \approx Ae^{i\phi_d(\tau)} \tag{5.1}$$

where A is the excitation imparted by the drive, $\phi_d(\tau)$ is the phase of the drive field at

the end of the excitation, and ρ_0 and ϕ_0 are respectively the initial thermal radius and phase of the particle. Because ϕ_0 and ρ_0 are randomly selected from the Boltzmann distribution in the cooling trap, to have a well-defined initial phase at the start of T_e the pulsed amplitude A must be much larger than ρ_c. The phase of the drive at the start of the second pulse is $\omega_d T_e + \phi_d(\tau)$, while the cyclotron phase has evolved to $\omega_+ T_e + \phi_d(\tau)$. After the second pulse, the complex amplitude of the cyclotron motion is the vector sum of the proton's evolved phase, the initial thermal radius, and the phase imparted by the drive:

$$\begin{aligned}\boldsymbol{\rho}_{cf} &= A\exp[i(\omega_+ T_e + \phi_d(\tau))] + A\exp[i(\omega_d T_e + \phi_d(\tau))] + \rho_0\exp[i(\phi_0 + \omega_+ T_e)] \\ &= Ae^{i\phi_d}\left(\exp[i\omega_+ T_e] + \exp[i\omega_d T_e] + \frac{\rho_0}{A}\exp[i(\phi_0 - \phi_d + \omega_+ T_e)]\right) \\ &= 2Ae^{i\phi_d}\left(\cos\left[\frac{\omega_+ - \omega_d}{2}T_e\right]\exp\left[i\frac{\omega_+ + \omega_d}{2}T_e\right] + \frac{\rho_0}{2A}\exp[i(\phi_0 - \phi_d + \omega_+ T_e)]\right)\end{aligned} \tag{5.2}$$

If we then approximate $\omega_+ - \omega_d \ll \omega_+$ this simplifies to

$$\boldsymbol{\rho}_{cf} \approx 2Ae^{i(\phi_d + \omega_+ T_e)}\left(\cos\left[\frac{\omega_+ - \omega_d}{2}T_e\right] + \frac{\rho_0}{2A}\left(\cos[\phi_0 - \phi_d] + i\sin[\phi_0 - \phi_d]\right)\right) \tag{5.3}$$

The final, real cyclotron amplitude thus encodes the detuning. The thermal component changes this real amplitude by the thermal initial radius times the cosine of the random thermal phase $(\phi_0 - \phi_d + \omega_+ T_e)$. The cosine or sine of a random phase will be between $\pm.87$ 67% of the time, so we can approximate this as an uncertainty:

$$\rho_{cf} \approx 2A\left|\cos\left[\frac{\omega_+ - \omega_d}{2}T_e\right]\right| \pm .87\rho_0 \tag{5.4}$$

Figure 5.2 shows the relative cyclotron amplitude after the second pi-pulse, as a function of the evolution time. We can invert this equation to measure the frequency:

$$\omega_+ = \omega_d + \frac{2}{T_e}\left(\cos^{-1}\left[\frac{\rho_{cf}}{2A}\right] + N_\pi\right) \tag{5.5}$$

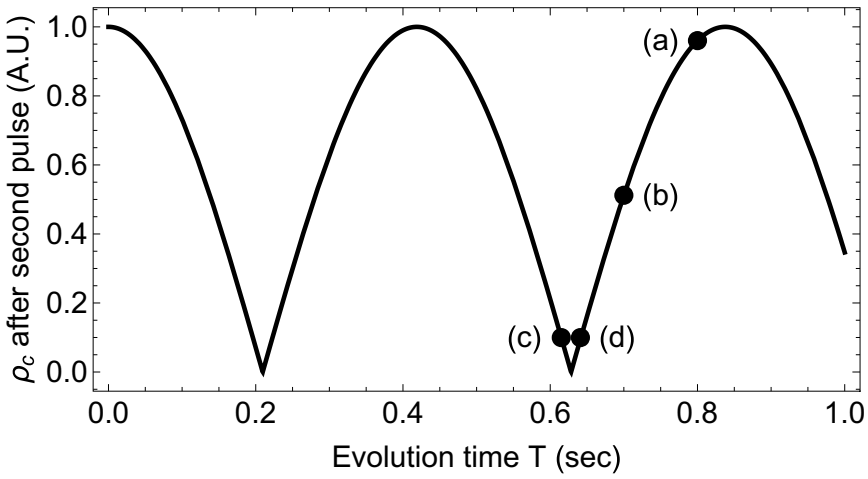

Figure 5.2: Fringe pattern for SOF cyclotron measurement: ρ_c after second pulse for 15 Hz drive frequency detuning. See text for labeled points.

where we have introduced N_π, representing the number of times the phase accumulates a full π during T_e. (This is π rather than 2π because we lose some phase information by measuring an amplitude – positive and negative ρ_{cf} are indistinguishable.)

In order to accurately measure the cyclotron frequency, N_π must be known. We therefore would employ a series of measurements with increasing evolution times, as used in [72]. Additionally, two points would be measured near each evolution time. This is necessary to account for lack of a priori knowledge of where the points will be on the fringe due to magnetic field drift [73], and the increased information from points with a larger slope – for example, in figure 5.2, point (b) determines ω_+ with more precision than point (a), while points (c) and (d) may lead to confusion about the number of fringes counted. (These two points do not necessarily need to be measured with the same evolution times, as long as the first evolution determines the cyclotron frequency with enough precision to put the second evolution at a high-information point like point (b) of figure 5.2. The second evolution then gives the high-precision

determination of the cyclotron frequency.)

To estimate the precision achievable with this method, we propagate the uncertainties through equation 5.5. Using standard error propagation we obtain

$$\sigma_{\omega_+}^2 = \left(\frac{\partial \omega_c}{\partial \rho_c}\right)^2 \sigma_{\rho_c}^2 + \left(\frac{\partial \omega_c}{\partial A}\right)^2 \sigma_A^2 \tag{5.6}$$

We assume that N_π, T_e and ω_d are known exactly. N_π should always be known exactly from earlier measurements with shorter evolution times. If uncertainty in T_e and ω_d becomes significant – for example, due to jitter in a timing source – then this analysis should be appropriately modified. Applying equation 5.6 to equation 5.5 gives

$$\sigma_{\omega_+}^2 = \frac{1}{A^2 T_e^2 \left(1 - \frac{\rho_{cf}^2}{4A^2}\right)} \sigma_{\rho_{cf}}^2 + \frac{\rho_{cf}^2}{A^4 T_e^2 \left(1 - \frac{\rho_{cf}^2}{4A^2}\right)} \sigma_A^2 \tag{5.7}$$

5.2.1 Measuring ρ_c in the analysis trap

Under this proposal, ρ_{cf} and A are both measured using the magnetic bottle in the analysis trap. The axial frequency in the analysis trap encodes the cyclotron quantum number n_c through equation 2.16, and thus the cyclotron radius. The axial frequency would be measured as a sequence of four steps: transfer the proton adiabatically to the analysis trap; employ a fast voltage scan to find its frequency using an axial or sideband drive; sideband cool it into the center of the trap; and precisely measure the axial frequency using the SEO.

We would like to estimate the achievable precision and time required per measurement with this method. The precision with which we can determine ω_+ for a given T_e depends on how accurately this procedure measures n_c, and thus ρ_{cf}. Three effects limit our measurement of n_c: uncertainty in the measurement of the axial frequency;

voltage instability, which changes the axial frequency for a given cyclotron radius; and changes in the total magnetic moment after the second SOF pulse[4]. We will estimate the magnitude of these effects in terms of uncertainty in the measured n_c.

Short-term axial frequency resolution

For this estimate, we will assume that we replicate the frequency resolution of the previous-generation apparatus [35]. The axial stability was found to be proportional to the square root of the cyclotron quantum number n_c. For low n_c, a standard deviation of as little as 70 mHz was achieved, corresponding to an uncertainty $\sim$1.4 cyclotron quanta. However, the actual axial frequency precision with the SEO will depend on n_c after the second SOF pulse. We numerically estimate this based on the measured scaling by multiplying the low-energy stability by $\sqrt{n_c}$.

We isolate short-term effects from the long-term drift and instability measured in the analysis trap of the first-generation experiment[5], because the rest of this section attempts to estimate the size of all longer-term effects in the context of the new procedure. Long-term drift in the first-generation apparatus was relatively well characterized by a cyclotron random walk; see [35] and later in this section.

Voltage instability

We separate voltage instability into short- and long-term effects, with short-term being at or below the time required to measure the axial frequency. Short-term voltage

[4]These same effects limit the precision of a single measurement of A, but because A should be identical for identical drive parameters we can average multiple measurements, and σ_A should be much smaller than $\sigma_{\rho_{cf}}$.

[5]For example, fig. 4.3(a) features a drift of $\sim$4 Hz over 14 hours

instability is accounted for in the achievable axial frequency resolution. On the other hand, some effects could cause the trapping potential to shift between measurements, distorting our estimate of ρ_{cf} from the measured axial frequency. These include drifts in trapping potential due to power supply or thermal effects, as well as and hysteresis or voltage shifts from the application of large voltages during transfers.

We can estimate the overall voltage drift using measured data in the new apparatus. With a single particle in the precision trap, we observed a drift of the dip frequency <0.8 Hz over 12 hours (see chapter 9), or on average .07 Hz per hour. This voltage stability would imply a drift equivalent to 1.4 cyclotron quanta per hour.

Particle transfers could also generate changing voltage offsets, either in the power supply or in the large cryogenic capacitors, which would shift the axial frequency. Preliminary tests in the precision trap of the new apparatus were hindered due to temperature and pressure instability – these tests were performed just prior to the Northwestern move, as the nitrogen dewars were already warming up. We did observe a longer-than-expected time constant in the particle's frequency as the trap voltage returned to its pre-transfer state, shown in figure 5.3, which may require ringing in the voltages by setting them temporarily below the desired trap potential. [6]

The best way to deal with this issue may be to keep the voltages constant on the ring and compensation electrodes, and change their relative potential by changing the voltages on the endcaps and transfer electrodes. This is electrostatically equivalent to increasing the trap voltage, and would keep the precision power supplies and large capacitors stable. This procedure should be investigated to determine what its effect is on transfer success rate and frequency settling time.

[6] Similar effects were observed in the electron experiment; see [74] for details.

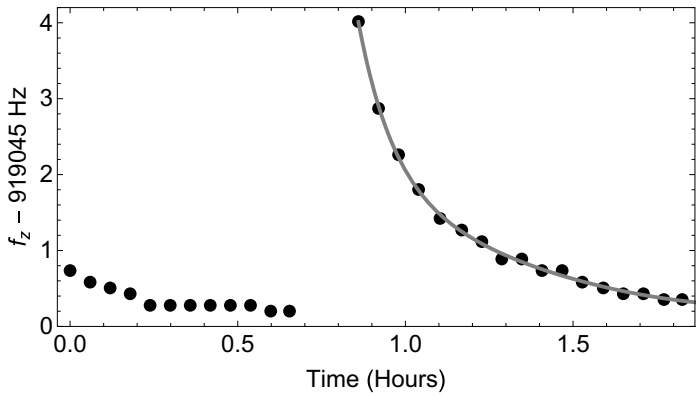

Figure 5.3: Axial frequency before and after a transfer. Grey line is a fit to an exponential decay plus a polynomial drift.

For this estimate, we will put a bound of 0.15 Hz, or 3 cyclotron quanta, per transfer on transfer effects – this roughly corresponds to the average frequency difference from before to after a transfer, in a series of five tests similar to figure 5.3, after subtracting the increased linear drift due to warmup-related thermal effects. (This was a high-uncertainty measurement; if the transfer effect is determined to be significantly larger, we can switch to a transfer scheme that keeps the trap potentials constant.)

Changes in n_c during the measurement procedure

Next, we account for changes in n_c after the second drive pulse, which shift the measured ρ_c away from ρ_{cf}. n_c was observed to drift in the previous apparatus as though it followed a classical random walk. For cold particles, this exhibited a step rate of $\sim$0.02 transitions per second.

The mechanism inducing these transitions is unclear, but two hypotheses have been proposed. The first is spurious RF fields near the cyclotron frequency, either from Johnson noise or coupled in from ambient noise. Figure 4.5 of [35] demonstrates

that even when applied to radially symmetric electrodes, near-resonant noise can cause cyclotron transitions; and the electrodes used for magnetron sideband cooling must have the same radial asymmetry needed to drive cyclotron transitions. All drive and DC lines going in to the analysis trap are heavily filtered, so little cyclotron noise should enter the trap – see section 8.4 as well as section 4.2.2 of [35]. However, further investigation of noise sources and couplings should follow if background transitions prove a significant source of noise on the measurement of n_c. Reference [75] proposes a method for determining the extent to which motional heating is dominated by noise on trap electrodes, which could be adapted to use in our Penning trap. The BASE collaboration reported a dependence of background cyclotron transition rate on magnetron radius [76]; this should be investigated in our apparatus, in case a focus on improving feedback cooling could reduce this transition rate.

Besides RF noise, a second hypothesis is "anomalous heating" as measured in quantum information studies using trapped ions [77, 78, 79]. Anomalous heating rates have been shown to depend on trap surfaces [80] – showing separate dependence on both adsorbed molecule layers and surface composition or condition [81]. This heating was measured in [82, 83] to scale proportionally to d^{-4}, for trap size d measured from 30 to 1000 μm, and to $f^{-2.5}$ for frequencies from 0.5 to 5 MHz. Predicting the heating rate in our trap requires extrapolating to 86 MHz, but gives a heating rate an order of magnitude smaller than observed in [6, 35]. However, in addition to the uncertainty from this extrapolation, our trap could be more susceptible to anomalous heating than the ion traps in these studies. These traps operate from room temperature to 4K [84], and thus feature a high room-temperature vacuum. In our sealed trapcan,

where XHV pressures are created by cryopumping of gas onto cold surfaces, a thicker layer of adsorbed atoms could be frozen onto the trap surface, increasing the heating rate, in an opposite analogue to the effect of plasma or ion cleaning [80, 83].

The transition rate in the new analysis trap should be equal or less than that observed in [6, 35]. If background transitions are due to anomalous heating from an adsorbed molecule layer, the rate should be unchanged. However, an improved surface-finishing procedure (section 8.1.3) and a new design which locates radially asymmetric electrodes further from the particle (section 8.1.2) should reduce the transition rate due to surface inhomogeneity and RF noise, respectively. That said, as a conservative approximation we will assume that the transition rate stays unchanged from the previous apparatus, at .02/sec for a cold proton.

Whatever the source of the perturbing electric field, the cyclotron transition matrix element $\langle n|\, qx\, |n+1\rangle$ will be proportional to $\sqrt{n+1}$. The transition rate therefore grows as n, and the random-walk deviation as $\sqrt{n}$. This determines the scaling of both the short-term axial frequency uncertainty and the long-term changes in ρ_c – the Allen deviation represents the effect of cyclotron transitions during the axial frequency averaging time, while the random walk represents the number of transitions taking place between the second pi-pulse and the final axial frequency measurement.

The transition rate in the precision trap should scale due to the larger electrode radius. If it is due to anomalous heating from surface effects, we would expect the precision trap rate to be smaller by a factor of 2^4, putting it at $\approx .001$ transitions per second. Alternatively, the interaction strength for RF noise near the cyclotron frequency would scale like the cyclotron damping rate, $\gamma_c \propto \frac{1}{\rho_0^2}$, putting our estimate

at .005 transitions per second for a cold proton (both cases increasing with n).

To translate the transition rate into an uncertainty, we note that for a one-dimensional classical random walk with N steps, the variance of the number of steps in a given direction is $\sigma_\uparrow^2 = \langle n_\uparrow^2 \rangle - \langle n_\uparrow \rangle^2 = NP_\uparrow P_\downarrow \approx N/4$. The total distance traveled (change in n_c) is $d = n_\uparrow - n_\downarrow = 2n_\uparrow - N$, so the variance will be $\sigma_d^2 = 4\sigma_\uparrow^2 \approx N$. This will add in quadrature with the other sources of error.

Sideband cooling – magnetron distribution

The magnetic moment due to the magnetron state also shifts the axial frequency in the bottle. After each transfer, the magnetron motion must be sideband cooled – the magnetron radius typically increases during transfers, either because of field misalignments or noise processes in the anharmonic transfer-electrode potential. Sideband cooling thermalizes the magnetron motion with the axial Boltzmann distribution. With feedback cooling active during sideband cooling, a distribution of magnetron frequency differences was observed in our single-spin-flip demonstration with a spread of 122 mHz [7], corresponding to ~ 3 cyclotron quanta.

5.2.2 Estimated cyclotron frequency measurement precision using SOF

Many of these noise sources scale with the cyclotron quantum number, so we will estimate the ρ_{cf} to be measured. A condition for a well-defined initial phase is for the pulse to increase the thermal energy by a factor of 50-100 [85]. We will consider a 2K sub-thermal proton selected using multiple trips to the analysis trap (as described

in section 3.2.3) and increased either by a factor of 50 (to 100K) or by a factor of 100 (to 200K). The post-pulse amplitudes would be $\rho_c \sim 2.3$ μm ($n_c \sim 24,000$) or $\rho_c \sim 3.4$ μm ($n_c \sim 50,000$). Table 5.1 collects noise sources for both cases.

Table 5.1: Sources of uncertainty in determining the cyclotron quantum number after SOF measurement

Source	$\sigma_n(2K)$	σ_n (100K)	σ_n (200K)
Trapping voltage drift	.0004/sec	.0004/sec	.0004/sec
Transition rate in analysis trap	.02/sec	1.3/sec	2.6/sec
Transfer hysteresis or voltage shift	3	3	3
Frequency determination uncertainty	1.4	14	20
Magnetron cooling distribution	3	3	3

Finally, we need to estimate the time required for the analysis-trap axial frequency measurement. This will determine the expected number of background cyclotron transitions, as well as the time required to measure ω_+ using SOF. To measure the axial frequency we will need to scan a range of possible trapping voltages determined by the magnetic bottle shift and the range of possible final cyclotron radii ρ_{cf}. After the second pulse, ρ_{cf} will be in a range between 0 and 2A. On average, we will have to search half of that range before we find the particle, because its position on the fringe after the second pulse is unknown a priori. For the 100K case we have $A = 2.3$ μm, and for the 200K case we have $A = 3.4$ μm. From equation 2.16, this implies a frequency range of 1.2 or 2.4 kHz.

The axial frequency need only be measured with enough precision that we are limited by the thermal initial radius $\boldsymbol{\rho}_0$. To optimize this process, we propose a dual

sweep technique. First, we sweep the trap voltage with a strong sideband heating or cooling drive to roughly locate the resonance. With a strong heating drive in the new apparatus, we have demonstrated resonance widths of up to 4 mV or ~ 430 Hz (figure 2.14). Introducing a safety factor of 2 to guarantee repeatability, we estimate a maximum step size of 215 Hz. Covering a range of 2.4 kHz would thus require 11 steps. Once we have localized the response, we will sweep a cooling drive through the 200 Hz range in approximately 2 minutes to center the proton and ensure that we have found the correct frequency. An axial frequency uncertainty of 215 Hz will on average resolve ρ_{cf} to better than the limit posed by the thermal initial radius. With a standard step time in the voltage sweep of 45 seconds, the entire process of finding the axial frequency would take on average 620 seconds. For the 100K particle, we could cover the 1.2 kHz range in 6 steps; this would reduce the average time to ~ 400 seconds. With that time estimate, we can return to table 5.1 and calculate the uncertainty on n_c:

$$\begin{aligned} \sigma_n\ (200K) &= \sqrt{(.0004 * 620)^2 + 2.8 * 620 + 3^2 + 20^2 + 3^2} = 46 \\ \sigma_n\ (100K) &= \sqrt{(.0004 * 400)^2 + 1.3 * 400 + 3^2 + 14^2 + 3^2} = 31 \end{aligned} \tag{5.8}$$

We can use equation 2.10 to propagate this through to uncertainty in ρ_{cf}:

$$\begin{aligned} \sigma_{\rho_{cf}}\ (200K) &\approx 1.2\,\text{nm} \\ \sigma_{\rho_{cf}}\ (100K) &\approx 1.1\,\text{nm} \end{aligned} \tag{5.9}$$

This uncertainty is much smaller than that due to the initial thermal state, which from equation 5.4 was $\sigma(\rho_{cf}) \approx .87\rho_0$, or $\sim .28\mu$m for a 2K initial state. (However, the detailed estimate was necessary to prove this; and we will use this framework

again in chapter 6.) Using equation 5.7, we calculate the uncertainty in ω_+:

$$\begin{aligned}\sigma_{\omega_+}\ (2K \to 100K) &\approx \frac{.14\ rad/s}{T_e} = \frac{250}{T_e}\ \text{ppt} \\ \sigma_{\omega_+}\ (2K \to 200K) &\approx \frac{.09\ rad/s}{T_e} = \frac{160}{T_e}\ \text{ppt}\end{aligned} \tag{5.10}$$

We are able to get a sub-ppb measurement of ω_+ from a few-percent measurement of ρ_{cf} because we know N_π from a series of less precise measurements – for $T_e = 1s$, knowing N_π is already a measurement of ω_+ with an effective 2π rad/s uncertainty. Similarly, reference [73] reports a sub-ppb measurement of ω_+ from a phase measurement with 10° uncertainty or ∼3% precision.

We note that because of the phase-coherent measurement, the precision increases linearly with T_e – for example, with a ten-second evolution time we can measure ω_+ with an uncertainty less than 10 mrad/s. We can also increase the precision by increasing the amplitude during the evolution time or choosing a colder initial cyclotron state; precision improves linearly with A/ρ_0. All of these require additional time – increasing A comes at the cost of a longer initial voltage scan required to find the axial frequency in the magnetic bottle, choosing a smaller ρ_0 requires more trips to the cooling trap, and increasing T_e requires more initial measurements at shorter evolution times to find the location on the fringe. Depending on the precision required, these competing timescales can be optimized differently. For example, the same precision can be achieved with a 4K initial particle energy and an excitation to 400K; this would increase the time required for the initial voltage scan by on average 8 minutes, but avoid the need to select a sub-thermal cyclotron state. Additionally, as we noted above each ω_+ measurement requires two iterations of the process of figure 5.1, making two points on the fringe of figure 5.2, because of the a priori

unknown position on the fringe. One of these points could be performed with a 4K initial state and a larger uncertainty, because its main function is to let us locate the other point on the fringe pattern. Systematic uncertainties are expected to be smaller than the estimated statistical uncertainty – for example, the relativistic shift for a 200K proton is only ~ 20 ppt, with an uncertainty at or below the ppt level.

5.3 Separated oscillatory fields measurement of ν_s

We can use the same principle of separated oscillatory fields to drive spin flips. A $\pi/2$ pulse puts the spin state into an equal superposition of spin-up and down. A second $\pi/2$ pulse after an evolution time T_e will interfere either constructively or destructively with the evolved phase, depending on the detuning between the drive and the Larmor frequency. Following [67], the expression for spin-flip probability after the second pulse is

$$P_{SF} = 4\left(\frac{\Omega_s}{\Omega_s'}\right)^2 \sin^2\left[\frac{\Omega_s'\tau}{2}\right] \times \left(\cos\left[\frac{\Delta_s T_e}{2}\right]\cos\left[\frac{\Omega_s'\tau}{2}\right] - \frac{\Delta_s}{\Omega_s'}\sin\left[\frac{\Delta_s T_e}{2}\right]\sin\left[\frac{\Omega_s'\tau}{2}\right]\right)^2 \quad (5.11)$$

where Ω_s is the spin-flip Rabi frequency (equation 2.32) applied during the $\pi/2$-pulses, $\Delta_s = \omega_s - \omega_{sd}$ is the detuning between the drive frequency ω_{sd} and the Larmor frequency ω_s, $\Omega_s' = \sqrt{\Delta_s^2 + \Omega_s^2}$ is the generalized Rabi frequency for the detuned drive, τ is the $\pi/2$-pulse duration, and T_e is the evolution time between pulses. This gives a characteristic Ramsey fringe pattern, shown in figure 5.4 for T_e=10 s and τ = 1 s. By making multiple spin-flip trials at slightly different drive frequencies, we can measure the excitation fraction and sweep out the central fringe to determine ω_s.

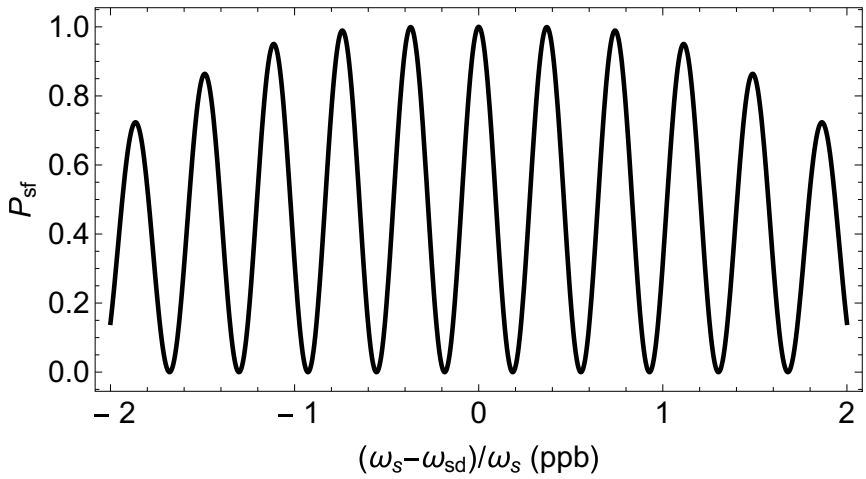

Figure 5.4: Ramsey fringe pattern of equation 5.11 for $\tau = 1$ second and $T = 10$ seconds.

For strong $\pi/2$ pulses or small detunings $\Omega_{sf} >> \Delta$, equation 5.11 reduces to

$$P_{sf} \approx \cos^2\left[\frac{\Delta T}{2}\right] \tag{5.12}$$

The fractional width (FWHM) of the central fringe is therefore approximately

$$\frac{\Delta_{\omega sf}}{\omega_{sf}} \approx \frac{\pi}{T_e \omega_{sf}} = \frac{2.3}{T_e} \text{ ppb} \tag{5.13}$$

Our precision goal, a 0.2 ppb FWHM, is thus achieved with $T_e \approx 10$ seconds.

As discussed in section 5.1, measuring separate spin and cyclotron resonances is impractical. Instead, each single-spin-flip trial should be associated with a ν_c measurement, with the ratio of applied spin-flip drive to ν_c giving the tested value of g/2. If the $\pi/2$ spin pulses are simultaneous with the cyclotron SOF pulses, the cyclotron and spin phases evolve under identical conditions – the SOF cyclotron measurement will give the integrated field which determines the spin precession. Figure 5.5 demonstrates a proposed sequence for this g-factor measurement.

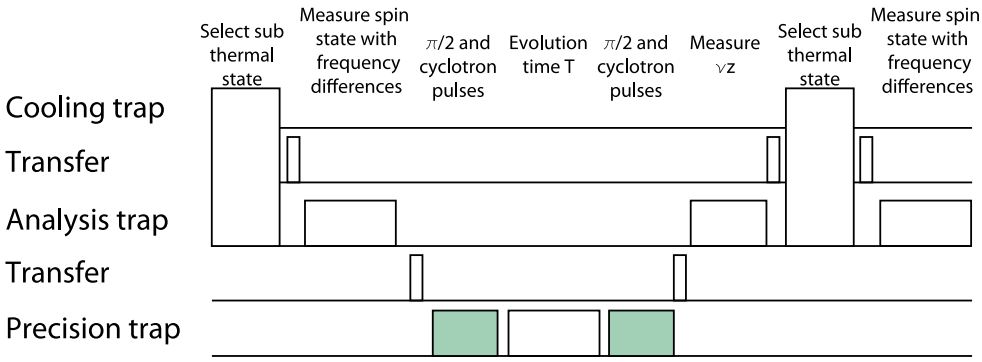

Figure 5.5: Procedure for SOF measurement of the g-factor (time axis not to scale; also note, preliminary cyclotron frequency measurement sequence omitted for clarity).

Following the calculations in the preceding section, each cyclotron measurement will have a higher precision than the g/2 linewidth. The statistical precision of a g/2 measurement using this simultaneous SOF method will be determined by the width of the central Ramsey fringe and the uncertainty on P_{sf} at each point. The width of the fringe is determined by T_e, and the uncertainty on P_{sf} by the number of spin-flip trials performed. Therefore, as a way to evaluate the method, we will estimate the time required to achieve a target level of statistical precision.

To do this, we need a criterion to evaluate the data acquisition rate. The point with the greatest slope gives the most information for a fit to the center frequency; we therefore focus on the P_{sf}=50% points at the wings of the central fringe. We will calculate the number of trials and amount of time required to achieve a signal-to-noise ratio of 3 at this point[7].

[7]In an actual measurement, the tested g-factor for each trial will be determined by the measured cyclotron frequency as well as the applied drive. Magnetic field drift means that the actual g-factor under test will differ from what is intended; rather than being individually binned into data points, the entire data set will be analyzed using statistical methods. However, that analysis is outside the scope of this thesis, and we will make our estimates for data binned at individual test values of g.

Measuring a spin-flip probability is equivalent to measuring the success fraction from a binomial distribution. Choosing a confidence interval estimator for the binomial probability is difficult – for a small number of trials, standard estimators tend to either over- or under-estimate the real confidence interval. Reference [86] provides an overview of different estimators. To compare to a standard deviation of normally distributed data, we will evaluate the 70% confidence interval for the binomial probability. For a 50% success rate, the Wilson score (which tends to underestimate the uncertainty) gives an interval $50 \pm 16\%$ after 10 trials, while the exact or Clopper-Pearson confidence interval (which tends to overestimate the uncertainty) gives $50 \pm 16\%$ after 15 trials. We will split the difference and estimate that 12 spin-flip trials will be required at a given g-factor value to satisfy our statistical criterion.

In addition to measuring and fitting the central Ramsey fringe, a complete g-factor measurement must also conclusively identify which fringe is the central one. Given prior knowledge of g, we could simply measure the correct fringe; however, this assumes the existence of a previous measurement at our target precision. Additionally, doing this would bias what should be an independent result. Instead, our measurement should identify the central fringe from the data.

We could do this by measuring the peak amplitude of adjacent fringes. The $\sin^2\left[\frac{\Omega'_{sf}\tau}{2}\right]$ term in equation 5.11 scales the height of the fringes depending on the detuning and duration of the $\pi/2$-pulse. However, for the case depicted in figure 5.4, the amplitude of the adjacent fringes is still 99.9%. Distinguishing the central from adjacent fringes would require over 100 spin-flip trials, and would be essentially impossible with even a small possibility of misidentified spin states. Increasing the

$\pi/2$-pulse amplitude and decreasing τ reduces the height of neighboring fringes, at the cost of a narrower power-broadened linewidth. To rotate the spin by $\pi/2$ in the presence of magnetic field drift, the pulse should apply roughly equal rotation across at least a 1 ppb range of field values, which gives a minimum Ω_s of $\sim \pi/2$ rad/sec and leaves us in the case pictured in fig. 5.4.

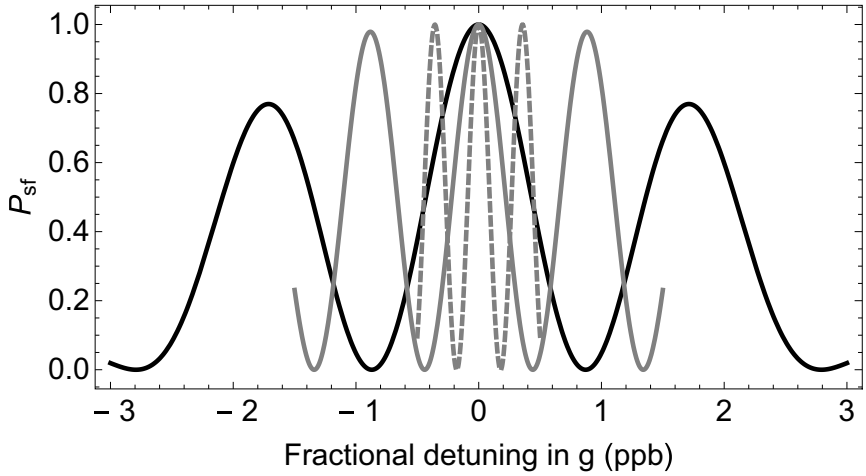

Figure 5.6: Ramsey fringes for evolution times of 1.2 sec (black), 3.5 sec (gray), and 10.4 sec (gray dashed), with 1 second $\pi/2$-pulse time. The amplitudes of the central and neighboring fringes are distinguishable for the 1.2-second fringe; knowledge of each previous central fringe then identifies the central fringe for the next evolution time.

A better way to identify the central fringe is to make a series of measurements with increasing evolution times, similarly to the proposed SOF measurement of ω_+. Keeping τ at 1 second but reducing T_e to $\sim$1.2 seconds would let us distinguish the central and adjacent fringe amplitude in a reasonable number of trials. We would then perform an intermediate measurement with $T_e\sim$3.5 seconds, before going to a final T_e of 10 seconds. As demonstrated in figure 5.6, increasing T_e by a factor of 3 at each step guarantees that each central fringe is within the FWHM of the previous

measurement, while the peaks of all neighboring fringes are outside that range. This increases the amount of time required by a factor of 3 over measuring only the central fringe, but it preserves the epistemological validity of our independent measurement.

We therefore conclude that between the three evolution times, we will have to make ∼36 spin-flip trials at each ratio to meet our statistical criterion (an SNR of 3 at the half-maximum points). To turn this into an estimate of time required, we need to evaluate the duration of each spin-flip attempt.

We make several plausible assumptions about technical improvements to the apparatus to evaluate the time per spin-flip trial. We assume transfer times between traps are negligible, using the system of diodes described in [62] to bypass the long filter time constants; a similar system was demonstrated in [87]. We assume the cooling trap (described in section 8.1.2) reduces the cyclotron cooling time to the predicted ∼5 seconds. We further assume, as discussed in section 5.2.1, we will be able to axial-frequency distortions from transfers. With these distortions minimized, we should be able to determine whether a proton is cold enough for spin-state detection by transferring it into the analysis trap, applying known voltage settings, and activating the SEO. Including SB cooling after the transfer and multiple SEO averages to reject transfer-related voltage drift, we estimate approximately 3 minutes required to find the axial frequency in the analysis trap, in addition to the ∼5 second cooling trap thermalization time. Finding a 1-in-5 proton, with a 2K cyclotron energy for spin-state detection, should therefore take on average 15.5 minutes

With these assumptions, we can sum up the total time required per data point for a 0.2 ppb measurement of the g-factor. Including 20 minutes for the cyclotron

frequency measurement, 15.5 minutes to find a 2K proton, and 3.5 minutes (two SEO averaging times on either side of the AFP sweep) to find the spin state, we estimate a total measurement time of 39 minutes per spin-flip attempt.

In addition to the simultaneous SOF spin-flip trials, we will occasionally need to measure ω_c using a shorter T_e, to ensure we always know N_π for the final evolution time. Magnetic field drift will shift ω_c between measurements.[8] A cyclotron frequency measurement with a shorter evolution time should be repeated when the magnetic field drifts enough to shift the cyclotron frequency by a significant fraction of the fringe width. After that time, we would need to repeat a cyclotron frequency measurement with a shorter evolution time to locate the fringe.

For a 10-second evolution time, the fringe width is $\pi/T \approx .3$ rad/sec out of the cyclotron frequency $\omega_c = 2\pi * 87$ MHz. A previous measurement of the electron magnetic moment in the same magnet [66] measured a typical field drift of 1 ppb per 12 hours, as well as an occasional faster drift (for unknown reasons) of 0.5 ppb per hour. For a 1/12 ppb per hour drift rate, the cyclotron frequency will change by a fringe width after around 6.6 hours; for 0.5 ppb per hour, it will change by a fringe width after around 1.1 hours. This sets how often we will need to insert an additional cyclotron frequency measurement. In total, this gives an estimate of 25 hours of data acquisition per point on the lineshape in the low-drift case, and 31 hours in the high-drift case.[9]

[8] Magnetic field drift during T_e is accounted for by the simultaneous SOF method, because the spin and cyclotron phases evolve under the same field.

[9] This estimate assumes that AFP can be successfully demonstrated with near-unity efficiency and fidelity in the analysis trap. Without AFP, data acquisition times will increase by a factor >4 (including extended spin-state preparation times as well as un- and misidentified final spin states) unless axial stability is improved significantly beyond the values demonstrated in the previous apparatus.

Chapter 6

Quantum Walk Measurement of the Magnetic Moment

Methods which measure the spin and cyclotron frequencies simultaneously, under the same combination of electric and magnetic fields, reduce the range of possible systematic shifts and uncertainties in a g-factor measurement. This advantage is being pursued for electron and positron measurements in our lab [88] and was used in the recent proton measurement discussed in section 5.1 [19], and is an integral part of the proposed dual SOF method described in chapter 5. In this chapter, we propose a new method for such a simultaneous measurement which improves the acquisition of statistics. This is accomplished by storing the information from multiple spin transitions as diffusion in the cyclotron energy.

Under the proposed scheme, drives are simultaneously applied at frequencies ν_{ad} near the anomaly frequency $\nu_a = \nu_s - \nu_+$ and ν_{sd} near the spin frequency ν_s, with the ratio ν_{ad}/ν_{sd} between them fixed. A signal, consisting of a broadening in the range of

cyclotron radii ρ_c, takes place only if the drives are simultaneously resonant[1]. This only occurs when the frequency ratio of the applied drives matches the ratio of the proton's true transition frequencies, $\nu_{ad}/\nu_{sd} \approx \nu_a/\nu_s$.

After describing the proposed procedure, we derive quantum mechanical equations of motion for the proton wavefunction addressed by both drives. We elucidate its behavior with a comparison to a quantum walk. Finally, we calculate the time required to measure the g-factor for a number of possible experimental scenarios. We use this to illustrate different ways the method can be adapted to conditions in the new environment at Northwestern University, where the experiment was recently moved. We conclude by comparing the quantum walk method with the separated oscillatory fields method described in chapter 5.

6.1 Measurement scheme

An outline of the measurement scheme is as follows. We begin by measuring the cyclotron quantum number n_c in the analysis trap. As described in section 5.2.1, we do this by measuring the axial frequency in the bottle field. This measurement couples the cyclotron and spin magnetic moments to the classical thermal bath in the

[1]We note that a proposal was put forward by Gabrielse and Quint in 1993 [89] for a parts-per-million measurement using this drive scheme with a different detection method.

axial amplifier, collapsing the wavefunction and giving a well-defined initial state.[2]

Once n_c is known, we transfer the particle to the precision trap. We then apply simultaneous drives at ν_{ad} and ν_{sd} – with the field symmetries required to make transitions – fixing the ratio ν_{ad}/ν_{sd}, for some time T. Figure 6.1 shows the spin and cyclotron energy levels addressed by the drives in this scheme. If neither drive is resonant, the particle remains in its initial state. If only one drive is resonant, the particle undergoes Rabi oscillations between two levels. However, if both drives are simultaneously resonant, then their combined action will couple up and down the ladder of cyclotron energy levels. With effectively equal[3] probability to climb up and down, n_c undergoes a quantum walk (the quantum mechanical analogy of a classical random walk). At the end of the drive period, the cyclotron state will be in a superposition of energy levels, with the width of the superposition depending on the Rabi frequencies and detunings of the applied drives, as well as the time spent on resonance during the drive period.

Finally, we transfer the particle back to the analysis trap and measure n_c. This

[2]We refer to n_c as a single quantum number after a measurement in the analysis trap, rather than a superposition, because the magnetic bottle couples the cyclotron motion to the dissipative, classical thermal bath in the axial amplifier. If the axial frequency resolution is narrower than the shift between cyclotron states, this dissipation collapses the cyclotron wavefunction down to a single state. This is achievable using the self-excited oscillator for ∼8 seconds, or using a dip for a longer time. Classical technical noise, including the FET noise floor, trapping voltage instability, etc, may prevent us from reliably identifying the value of n_c at the single-quantum level, but the strength of the coupling to the classical thermal bath is determined by the resolution of a single measurement, not the stability of many measurements over time. This is the same principle which causes the spin state to collapse from a measurement of the axial frequency after applying a spin-flip drive, in spite of sub-unity efficiency and fidelity.

However, collapsing the wavefunction to a single state is not necessary for the method; the equations we will derive apply to any single-particle wavefunction composed of a superposition of cyclotron states with an average quantum number $\langle n_c \rangle$. We refer to a single quantum number n_c both because we believe that to be a more correct analysis, and because it simplifies the conceptual discussion of the method.

[3] The ratio of probabilities to make transitions up and down is $P_\uparrow/P_\downarrow = \sqrt{n+1}/\sqrt{n} \approx 1$ for a 2-4K proton.

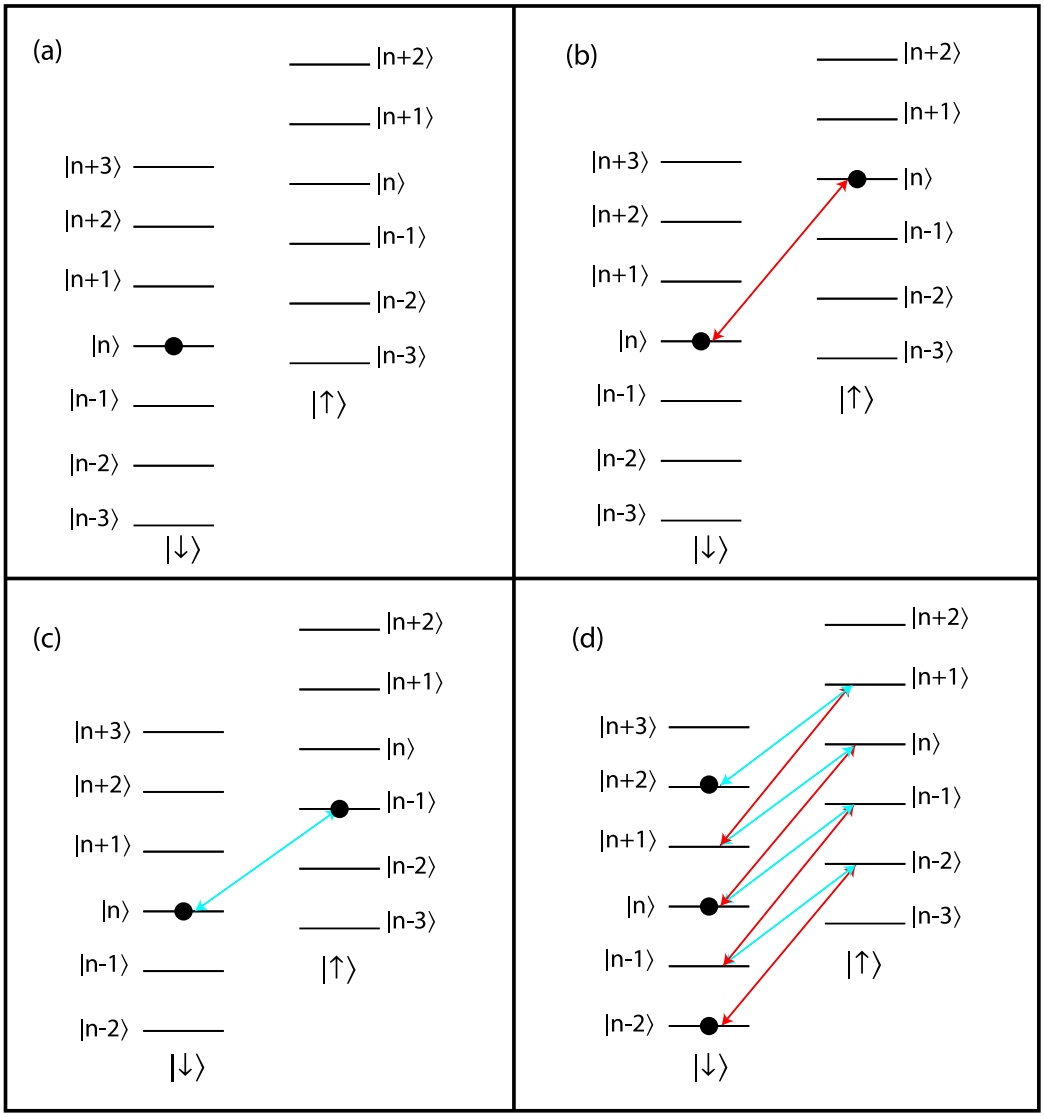

Figure 6.1: Energy level diagram for the simultaneous drive method. (a) Initial state with cyclotron energy ladders for spin up and spin down. (b) Spin-flip drive induces transitions between spin states. (c) Anomaly drive induces one simultaneous spin and cyclotron transition. (d) When both drives are resonant, the proton climbs up and down the cyclotron energy ladder in a diffusion of quantum states.

returns an n_c value, depending on the probabilities in the superposition after the quantum walk. The change in cyclotron between measurements is recorded as Δn_c. This procedure is repeated several times to find the standard deviation of differences, $\sigma(\Delta n_c)$, at a given frequency ratio. Measuring at different ratios gives a resonance profile for the ratio ν_a/ν_s.

A block diagram of the proposed measurement scheme is presented in figure 6.2. In addition to the steps outlined above, we have added blocks for interlaced measurements of ν_c using pulsed SOF (similar to the procedure described in sec. 5.2, with minor modifications described in sec. 6.5). These measurements of ν_c are required to account for magnetic field drift[4] (see sec. 6.4.3). Under the optimal version of this scheme, a second "magnetometer" proton is trapped and kept spatially separated from the original "measurement" proton, as demonstrated in a similar experiment in [18]. The "measurement" proton then remains at an appropriate cyclotron radius without needing frequent re-thermalization and selection. The time between ν_c measurements will depend on how well the magnetic field drift can be controlled and predicted in the system at Northwestern.

During each trial, the ratio of drive frequencies ν_{ad}/ν_{sd} will be held constant. The specific values of each frequency will depend on the known magnetic field value during the trial, as determined from the most recent cyclotron frequency measurement and our overall knowledge of the magnetic field drift rate (ie, if the magnetic field follows a predictable linear or polynomial drift). Once the ratio is chosen, the frequencies

[4]While we measure the response to the ratio ν_{ad}/ν_{sd}, which is independent of the magnetic field value, in order for transitions to take place both ν_{ad} and ν_{sd} must be resonant simultaneously during the trial. Choosing specific values or ranges for the drive frequencies requires a magnetic field estimate, but that estimate does not affect the final measurement of the g-factor.

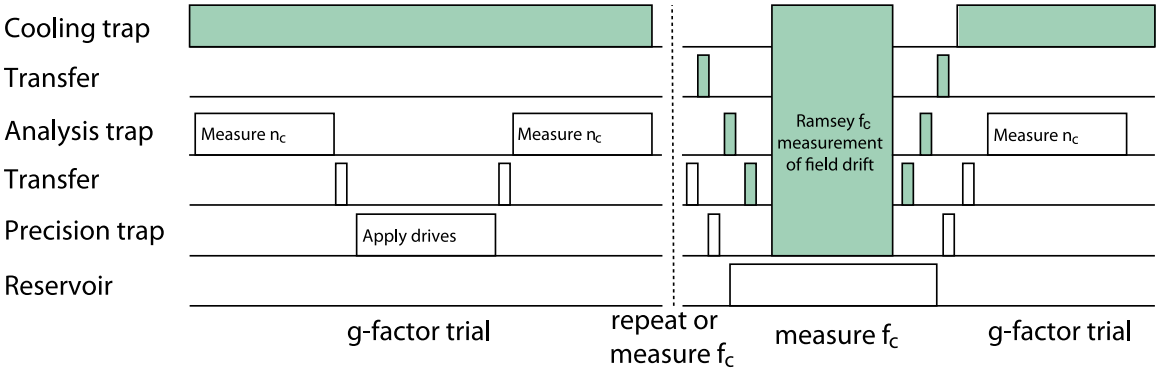

Figure 6.2: Measurement sequence for the simultaneous drive method. The interval between cyclotron frequency measurements with the magnetometer proton will depend on measured field drift. Empty and filled blocks represent the location of the measurement and magnetometer protons, respectively. Time axis not to scale.

can either be held constant as the magnetic field drifts through resonance, or chirped through some range.[5] Some specific scenarios for this decision are evaluated in sec. 6.6.

Not included in this basic sequence is the cooling and re-selection of the measurement proton's cyclotron radius (by repeatedly thermalizing in the cooling trap followed by axial frequency measurement in the analysis trap – see sec. 3.2.3). We expect this to occur very infrequently (estimated at fewer than four times per day) once an initial ρ_c has been selected. The acceptable range of cyclotron radii will depend on measured experimental parameters – a large ρ_c will increase radius-dependant noise (sec. 5.2.1 and 6.4.2) while an extremely small ρ_c will limit the anomaly-drive transition rate (sec. 6.2.2). If the proton's cyclotron energy grows too large or small, it will

[5]Whether or not the drives should be chirped will depend on the uncertainty in our ν_c measurement, the measured long-term magnetic field drift rate, and the short-term inconsistency in the field drift. For example, if the uncertainty in ν_c is larger than the drift rate, we will chirp the drives through the range of uncertainty to guarantee that they pass through resonance during the trial. If the short-term inconsistency is large compared to the long-term drift, then we will chirp the drives to ensure that the time spent on resonance is similar for each trial.

have to be re-thermalized.[6]

Cyclotron radius growth over time will arise from background transitions and transitions driven during the measurement cycle. Background transitions can be kept low with appropriate filtering and trap geometry.[7] The net growth in cyclotron radius during the drive cycle should also be small, as the relative transition-rate scaling $P_\uparrow/P_\downarrow = \sqrt{n+1}/\sqrt{n} \approx 1$ implies an essentially undirected random walk. Based on [76] and cyclotron radius growth rates during the 2013 measurement, we expect cyclotron radius re-selection to be necessary less than once per 24 hours.

6.2 Single-drive Rabi frequencies and lineshapes

Later in this chapter, we will evaluate the transition rate and lineshape using the quantum mechanical equations of motion for the cyclotron and spin states under the action of both drives, both of which are themselves classical fields. To ensure we correctly derive those equations, we must first consider the transition probabilities and lineshapes for the two individual drives.

In our 2013 magnetic moment measurement, the linewidth was dominated by the axial thermal distribution in the magnetic bottle gradient. In the electron magnetic moment measurements, the relativistic shift between cyclotron states was large compared to the measurement precision. In the simultaneous anomaly-spin drive method

[6] The calculations in this chapter are performed for a proton with 2K of cyclotron energy.

[7] The BASE experiment recently reported an observed background transition rate of 6 quanta per hour [76] in a similar Penning trap. Improvements to our new apparatus in filtering, surface treatment (sec. 8.1.3), and trap geometry (sec: 8.1.2 and ref [62]), should dramatically reduce the rate of background transitions compared to that measured in the first-generation apparatus (which was already small compared to the change in cyclotron radius due to the simultaneous spin and anomaly drives).

proposed here, we also consider the effect of power broadening, since we need enough power to drive multiple transitions. In this section, we will show that the magnetic-bottle lineshape approximates a sharp δ function, the relativistic shift is negligible, and that power broadening dominates the lineshape.

6.2.1 Residual magnetic bottle linewidth and relativistic shift

In the weak-drive limit, the transition probability per time would equal $\frac{\pi}{2}\Omega_A^2\chi(\omega)$, with χ being the magnetic bottle lineshape parameter [90]. (We discussed the magnetic bottle linewidth in the analysis trap in chapter 3; the lineshape parameter χ in the precision trap arises from the same Boltzmann distribution of axial states, but with a different form due to the smaller magnetic gradient.) However, we intend to drive many transitions during each trial, and with the new trap design, the residual magnetic bottle linewidth should be quite small. Therefore, power broadening from the drive, rather than the magnetic bottle, becomes the dominant lineshape source.

To demonstrate this, we will plot the linewidth parameter χ [52]:

$$\chi(\omega) = \frac{4}{\pi} Re\left[\frac{\gamma'\gamma}{(\gamma'+\gamma)^2}\right] \sum_{n=0}^{\infty}\left[\frac{(\gamma'-\gamma)^{2n}(\gamma'+\gamma)^{2n}}{(n+\frac{1}{2})\gamma' - \frac{1}{2}\gamma - i(\omega-\omega_0)}\right] \tag{6.1}$$

where γ is the axial damping width, ω the drive frequency, ω_0 the resonant frequency of the transition, and γ' is given by

$$\gamma' = \sqrt{\gamma^2 + 4i\Delta\omega_z} \tag{6.2}$$

with $\Delta\omega_z = \omega_0 \frac{B_2}{B}\frac{k_b T_z}{m\omega_z^2}$ defined as in equation 3.6.

We will drive on the anomaly frequency $\omega_a = \omega_s - \omega_+ \approx 2\pi{*}156$ MHz. B_0 is the homogeneous field, approximately 5.7 Tesla in the precision trap. For γ we use the

measured single-particle damping width of $2\pi * 0.7$ Hz. B_2 should be small thanks to the new trap electrode geometry discussed in sec. 8.1.1, and can be made smaller with techniques discussed in sec. 10.2.4. Estimates for the new trap geometry give a residual bottle field $\sim$0.03 T/m^2. For this chapter, we will use a more conservative 0.1T/m^2 to demonstrate that even with a safety factor the B_2 contribution is negligible. An axial temperature of 6K and frequency of 1 MHz give $\Delta\omega_z \approx 22$ mrad/sec. Figure 6.3 (a) shows a numerical calculation of the linewidth parameter. It is sharply peaked around $\omega = \omega_0 + \Delta\omega_z$, with a width $\approx \frac{\Delta\omega_z^2}{\gamma}$ =0.1 mrad/s (equation 6.27 of ref. [1]). The overall lineshape including both power broadening and the magnetic bottle will be a convolution of the two effects; since the magnetic bottle linewidth is approximately a delta function compared to the power-broadened linewidth, we can safely neglect it.[8]

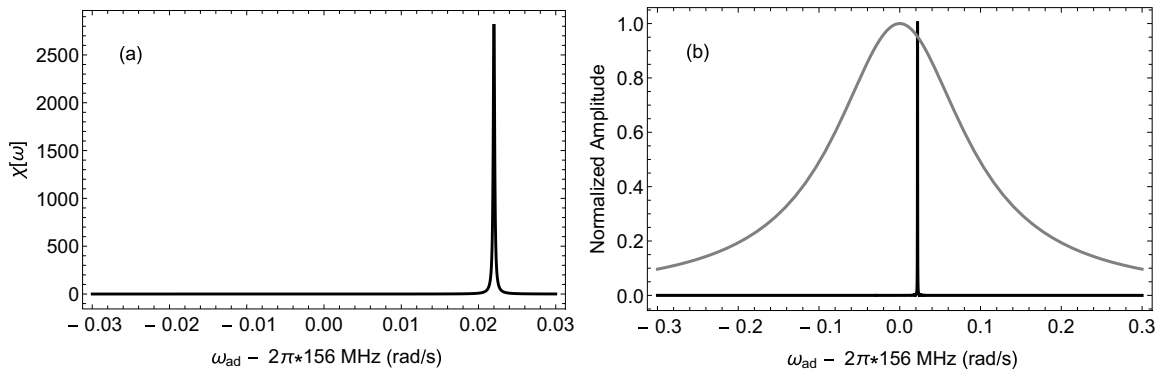

Figure 6.3: (a) Lineshape parameter $\chi(\omega)$ for estimated experimental values. (b) Normalized lineshape parameter (black) and transition rate for a 0.1 ppb power-broadened linewidth (gray) for the anomaly transition.

[8]The shifted center frequency, $\omega = \omega_0 + \Delta\omega_z$, may be significant depending on the value of the magnetic bottle and the amount of line-splitting desired. Using techniques described in [45] we expect to measure the magnetic bottle with 0.01 T/m^2 precision, giving a .002 ppb uncertainty from the shift to the anomaly frequency.

The relativistic shift is another possible contributor to the lineshape. The proton's mass shifts proportionally to its energy, as discussed in section 2.5, by an amount $\frac{\Delta\nu_+}{\nu_+} = -\frac{E_c}{E_c+m_pc^2} \approx -\frac{E_c}{m_0c^2}$. The cyclotron energy is given by $E_c = \hbar\omega_c\left(n+\frac{1}{2}\right)$. A change of one quantum thus shifts the frequency by a fractional amount $\frac{\hbar\omega_c}{m_0c^2} \approx 4*10^{-16}$, or 0.2 μrad/sec. For an expected 10-50 transitions, the shift during a drive cycle would be only two parts in 10^{14} of the cyclotron frequency, much less than either the power-broadened or magnetic bottle linewidth.[9]

6.2.2 Power-broadened linewidth and drive strength

To determine the linewidth due to power broadening, and compare it to the residual magnetic bottle linewidth, we consider only the spin or anomaly transition and take the rotating wave approximation for a driven two-level system. The interaction with a linear oscillating field at frequency ω_d, with strength such as to induce resonant Rabi oscillations at a Rabi frequency Ω, with zero energy chosen to correspond to the ground state, can be written as

$$H = -\hbar\frac{\Omega}{2}\left(e^{-i\omega_d t}+e^{i\omega_d t}\right)\left(|1\rangle\langle 2| + |2\rangle\langle 1|\right) + \omega_0|2\rangle\langle 2| \tag{6.3}$$

After applying the standard rotating wave approximation, the Schrodinger equation in the rotating frame gives, for a drive with detuning $\omega_0-\omega_d=\Delta$, a wavefunction $|\psi\rangle = C_1|1\rangle + C_2|2\rangle$, and initial conditions $C_1=1,\ C_2=0$:

$$\begin{aligned} i\frac{\partial|\psi\rangle}{\partial t} &= -\left(\frac{\Omega}{2}\left(|1\rangle\langle 2| + |2\rangle\langle 1| + \Delta|2\rangle\langle 2|\right)\right)|\psi\rangle, \\ |C_2|^2 &= \frac{1}{2}\frac{\Omega^2}{\Delta^2+\Omega^2}\left(1-\cos\left(t\sqrt{\Delta^2+\Omega^2}\right)\right) \end{aligned} \tag{6.4}$$

[9] A larger change in cyclotron radius will occur when the measurement proton needs to be rethermalized. However, the attendant relativistic shift is still much smaller than the measurement precision – a change in energy as large as 20K still corresponds only to a shift of 2 ppt.

This gives probability oscillations at the generalized Rabi frequency $\sqrt{\Delta^2+\Omega^2}$. The amplitude $\frac{\Omega^2}{\Delta^2+\Omega^2}$ of these oscillations is a Lorentzian with width Ω, giving the power broadening (plotted in figure 6.3(b)). On resonance this gives a transition probability of $\frac{1-\cos(\Omega t)}{2}$, which for small Ωt gives $\frac{t^2\Omega^2}{4}$.

We apply the anomaly drive current using a tuned-circuit transmission line, as described in sections 2.7 and 8.5. Connected to both compensation electrodes, we get two loops of current travelling in opposite directions, creating the xz gradient required to drive anomaly transitions. The effective Rabi frequency for this drive scheme is [1]

$$\Omega_A = \frac{qg}{2m}\frac{b_1\rho_c}{2} \tag{6.5}$$

where ρ_c is the proton's cyclotron radius and b_1 is the strength of the linear, oscillating drive magnetic field, given by

$$\mathbf{b} = b_1\boldsymbol{\rho}(t)\cos(\omega_{ad}t) \tag{6.6}$$

with ω_{ad} being the angular frequency of the anomaly drive. For two current loops of radius a, the xz-gradient component of the field near trap center is

$$b_1 = \frac{3a^2\mu_0 Id}{4(a^2+\frac{d^2}{4})^{5/2}} \tag{6.7}$$

where d is the distance between the two current loops and I is the drive current through each electrode.

Suppose our goal is power broadening of order 0.1 ppb. Since the Lorentzian width of equation 6.4 is Ω, we would choose a Rabi frequency such that $\frac{\Omega_A}{2\pi 156 MHz} = 1*10^{-10} \rightarrow \Omega_A \approx 100$ mrad/s. The bottle lineshape and relativistic shift are both at least 1000 times smaller, and can be neglected compared to the power broadening.[10]

[10] A Bloch-Siegert shift is also imposed by the counter-rotating component of the linear oscillating field, but with magnitude $\sim \Omega_A^2/\omega_a^2 \approx 10^{-20}$ this is negligible [91].

We briefly digress to evaluate whether this Rabi frequency is achievable. We use the geometry of the precision trap to calculate the gradient b_1 due to the drive. For current in the compensation electrodes, equation 6.7 gives a field strength $\sim .06\ \frac{T}{m}$ per ampere. The Rabi frequency is then given by equation 6.6 to be $\Omega_A \approx 8 * I * \rho_c$, where I is again the drive current in amperes and ρ_c is the cyclotron radius in μm. The cyclotron radius of a proton at temperature T can be found from equipartition:

$$\rho_c = \sqrt{\frac{2k_bT}{\omega_c^2 m}} \tag{6.8}$$

For a 2K proton, $\rho_c = 0.3\ \mu$m. To achieve a $2\pi * 15$ mHz ($\Omega_A = 100$ mrad/s) Rabi frequency we would need ~40 mA of current at a drive frequency of 156 MHz in the electrodes. In the 2013 antiproton measurement, a current of ~31mA was coupled through the spin-flip drive line with no measurable power shift. Data in figure 8.19 indicates a 25% increase in drive current would still apply a power shift to the axial frequency of less than .7 Hz, shifting the cyclotron frequency by less than 0.1 ppb. This should be carefully measured as a possible systematic shift, but the axial frequency precision means uncertainty on this shift will not limit the overall measurement precision.

The analysis of the spin-flip drive proceeds similarly, with two caveats. First, 0.1 ppb of $2\pi * 240$ MHz implies a Rabi frequency of 150 mrad/s. Second, the Rabi frequency is given by the oscillating magnetic field B_s from two parallel half-loops of current, with radius a and at an axial distance z from the particle:

$$\Omega_S = \frac{|B_s|\,\mu_p}{\hbar}, \qquad B_s = \frac{\mu_0 I}{\pi}\frac{az}{(a^2+z^2)^{3/2}} \tag{6.9}$$

For the precision trap, with current flowing through both sides of an endcap, we have B_S of $1.5 * 10^{-5}$ Tesla per ampere of driving current. Only 75 μA is required for a

Rabi frequency of 150 mrad/s. (The current required to drive spin flips is so much less than for anomaly transitions because the spin-flip Rabi frequency does not depend on the small cyclotron radius ρ_c.)

6.3 Quantum state evolution under the simultaneous anomaly and spin drives

To drive simultaneous spin and anomaly transitions, we apply RF fields near the transition frequencies. Geometries for these drives are given in sec. 2.7. We will generate the interaction Hamiltonian terms for these drives in order to apply the rotating wave approximation and obtain equations of motion for the proton wavefunction.

The spin-flip drive is applied via parallel currents, creating an oscillating linear magnetic field along the $\hat{x}$ axis. This field acts on the magnetic moment of the proton, adding a term to the Hamiltonian

$$\boldsymbol{\mu}\cdot\boldsymbol{B} = \frac{\hbar}{2}\frac{g}{2}\frac{e}{m_p}B_S\cos(\omega_{sd}t)\sigma_x = \frac{\hbar}{2}\Omega_s\left(e^{i\omega_{sd}t}+e^{-i\omega_{sd}t}\right)\left(|\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow|\right) \tag{6.10}$$

where we have written $\sigma_x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} = \left(|\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow|\right)$ (for a spin quantized along the $\hat{z}$ axis) and $\cos(\omega_{sd}t) = \frac{1}{2}\left(e^{i\omega_{sd}t}+e^{-i\omega_{sd}t}\right)$.

The anomaly drive is applied using opposed current loops formed using the compensation electrodes. These create a magnetic field gradient at the proton's position of the form $\mathbf{b} = b_1\boldsymbol{\rho}(t)\cos(\omega_{ad}t)$. The interaction energy $\boldsymbol{\mu}\cdot\boldsymbol{b}$ from this field is proportional both to the spatial coordinate $\boldsymbol{\rho}$ and to the spin operators. We can then

write the interaction Hamiltonian term as

$$\boldsymbol{\mu}\cdot\boldsymbol{b} = \frac{\hbar}{2}\frac{g}{2}\frac{e}{m_p}b_1\cos(\omega_{ad}t)(\sigma_x\hat{x}+\sigma_y\hat{y}) = \frac{\hbar}{2}\frac{g}{2}\frac{e}{m_p}b_1\cos(\omega_{ad}t)\begin{bmatrix} 0 & \hat{x}-i\hat{y} \\ \hat{x}+i\hat{y} & 0 \end{bmatrix} \tag{6.11}$$

The position operators $\hat{x}$ and $\hat{y}$ can be defined in terms of cyclotron angular momentum raising and lowering operators, $\hat{a} = \sqrt{\frac{\hbar}{2m_p\omega_+}}(\hat{x}+i\hat{y})$, $\hat{a^\dagger} = \sqrt{\frac{\hbar}{2m_p\omega_+}}(\hat{x}-i\hat{y})$

$$\boldsymbol{\mu}\cdot\boldsymbol{b} = \frac{\hbar}{2}\frac{g}{2}\frac{e}{m_p}b_1\cos(\omega_{ad}t)\sqrt{\frac{2m_p\omega_+}{\hbar}}\begin{bmatrix} 0 & \hat{a^\dagger} \\ \hat{a} & 0 \end{bmatrix} \tag{6.12}$$

or equivalently

$$\begin{aligned}\boldsymbol{\mu}\cdot\boldsymbol{b} =& \frac{\hbar}{2}\frac{g}{2}\frac{e}{m_p}b_1\left(e^{i\omega_{ad}t}+e^{-i\omega_{ad}t}\right)\sqrt{\frac{2m_p\omega_+}{\hbar}}\times \\ & \left(|\uparrow\rangle\langle\downarrow|\sqrt{n+1}\,|n\rangle\langle n+1| + |\downarrow\rangle\langle\uparrow|\sqrt{n+1}\,|n\rangle\langle n+1|\right)\end{aligned} \tag{6.13}$$

Making the approximation $n >> 1$ for a thermal particle, we can write this as

$$\boldsymbol{\mu}\cdot\boldsymbol{b} = \frac{\hbar}{2}\frac{g}{2}\frac{e}{m_p}b_1\left(e^{i\omega_{ad}t}+e^{-i\omega_{ad}t}\right)\rho(|n,\uparrow\rangle\langle n+1,\downarrow| + |n,\downarrow\rangle\langle n-1,\uparrow|) \tag{6.14}$$

Finally, using the definition of Ω_A from equation 6.5, we have

$$\boldsymbol{\mu}\cdot\boldsymbol{b} = \frac{\hbar}{2}\Omega_A\left(e^{i\omega_{ad}t}+e^{-i\omega_{ad}t}\right)(|n,\uparrow\rangle\langle n+1,\downarrow| + |n,\downarrow\rangle\langle n-1,\uparrow|) \tag{6.15}$$

6.3.1 Rotating wave approximation for simultaneous drives

We now move to evaluate quantum state evolution under both drives. The interaction Hamiltonian terms obtained above are added to the non-interacting Hamiltonian, which is just the energy of the two-level spin and the cyclotron harmonic oscillator:

$$H_0 = \sum_n \hbar\left[-\frac{\omega_s}{2}+\left(n+\frac{1}{2}\right)\omega_c\right]|n,\downarrow\rangle\langle n,\downarrow| + \hbar\left[\frac{\omega_s}{2}+\left(n+\frac{1}{2}\right)\omega_c\right]|n,\uparrow\rangle\langle n,\uparrow|$$

Combining the interaction and non-interacting terms we have the Hamiltonian

$$\begin{aligned}\mathrm{H}=\sum_n \Bigg[& -\frac{\hbar\Omega_A}{2}\left(e^{-i\omega_{ad}t}+e^{i\omega_{ad}t}\right)\left(|n,\uparrow\rangle\langle n-1,\downarrow|+|n,\downarrow\rangle\langle n+1,\uparrow|\right) \\ & -\frac{\hbar\Omega_S}{2}\left(e^{-i\omega_{sd}t}+e^{i\omega_{sd}t}\right)\left(|n,\uparrow\rangle\langle n,\downarrow|+|n,\downarrow\rangle\langle n,\uparrow|\right) \\ & +\hbar\left(-\frac{\omega_s}{2}+\left(n+\frac{1}{2}\right)\omega_c\right)|n,\downarrow\rangle\langle n,\downarrow|+\hbar\left(\frac{\omega_s}{2}+\left(n+\frac{1}{2}\right)\omega_c\right)|n,\uparrow\rangle\langle n,\uparrow|\Bigg]\end{aligned} \tag{6.17}$$

In Appendix A, we apply a pair of unitary transformations and the rotating wave approximation to derive a Hamiltonian in the rotating frame:

$$\begin{aligned}\tilde{H}_{RWA}\approx & \sum_n\left[-\frac{\hbar\Omega_A}{2}(|n,\uparrow\rangle\langle n-1,\downarrow|+|n,\downarrow\rangle\langle n+1,\uparrow|)-\frac{\hbar\Omega_S}{2}(|\uparrow\rangle\langle\downarrow|+|\downarrow\rangle\langle\uparrow|)\right] \\ & +\sum_n\left[\hbar(\Delta_{sf}+(n-n_0)\Delta_{cd})\,|n,\uparrow\rangle\langle n,\uparrow|+[\hbar(n-n_0)\Delta_{cd}\,|n,\downarrow\rangle\langle n,\downarrow|\right]\end{aligned} \tag{6.18}$$

where $\Delta_{sf}=\omega_{sd}-\omega_s$ represents the detuning between the spin-flip drive and transition frequency, and $\Delta_{cd}=\omega_{sd}-\omega_{ad}-\omega_c$ represents the combined detuning of both drives from the true cyclotron frequency. Similarly, the detuning of the anomaly drive from that transition frequency is defined as $\Delta_{ad}=\Delta_{sd}-\Delta_{cd}$. To evaluate the dynamics of this system, we apply a general wavefunction in the rotating frame:

$$\left|\tilde{\psi}\right\rangle=\sum_n\left(\tilde{C}_{n\uparrow}\,|n,\uparrow\rangle+\tilde{C}_{n\downarrow}\,|n,\downarrow\rangle\right) \tag{6.19}$$

Note that the transformation to the rotating frame applies a time-dependant phase shift to each state; the wavefunctions before and after the unitary transformation are related by $\tilde{C}_{n\uparrow}=e^{i(\omega_s+(n-n_0)\omega_c)t}C_{n\uparrow}$ and $\tilde{C}_{n\downarrow}=e^{i(n-n_0)\omega_c t}C_{n\downarrow}$. As expected, the unitary transformations preserve the probability amplitude, $\tilde{C}^{\star}_{n\uparrow,\downarrow}\tilde{C}_{n\uparrow,\downarrow}=C^{\star}_{n\uparrow,\downarrow}C_{n\uparrow,\downarrow}$, and equations 6.18 and 6.19 together obey the Schrodinger equation in the rotating

frame. By orthogonality, the Schrodinger equation in the rotating frame gives

$$i\hbar\frac{\partial}{\partial t}\sum_n(\tilde{C}_{n\uparrow}|n,\uparrow\rangle+\tilde{C}_{n\downarrow}|n,\downarrow\rangle)=\tilde{H}_{RWA}\left|\tilde{\psi}\right\rangle \tag{6.20}$$

$$\begin{aligned}=&\sum_n-\frac{\Omega_A}{2}(\tilde{C}_{n-1,\downarrow}|n,\uparrow\rangle+\tilde{C}_{n+1,\uparrow}|n,\downarrow\rangle)-\frac{\Omega_S}{2}(\tilde{C}_{n,\downarrow}|n,\uparrow\rangle+\tilde{C}_{n,\uparrow}|n,\downarrow\rangle)\\&+(n-n_0)\Delta_{cd}\tilde{C}_{n+1,\downarrow}|n+1,\downarrow\rangle+(\Delta_{sf}+(n-n_0)\Delta_{cd})\tilde{C}_{n+1,\uparrow}|n+1,\uparrow\rangle\end{aligned} \tag{6.21}$$

Multiplying by $\left\langle\tilde{\psi}\right|$ we obtain equations of motion for the wavefunction coefficients:

$$\begin{aligned}i\frac{\partial}{\partial t}\tilde{C}_{n,\uparrow}&=-\frac{\Omega_A}{2}\tilde{C}_{n-1,\downarrow}-\frac{\Omega_S}{2}\tilde{C}_{n,\downarrow}+(\Delta_{sf}+(n-n_0)\Delta_{cd})\tilde{C}_{n,\uparrow}\\i\frac{\partial}{\partial t}\tilde{C}_{n,\downarrow}&=-\frac{\Omega_A}{2}\tilde{C}_{n+1,\uparrow}-\frac{\Omega_S}{2}\tilde{C}_{n,\uparrow}+(n-n_0)\Delta_{cd}\tilde{C}_{n,\downarrow}\end{aligned} \tag{6.22}$$

6.3.2 Quantum walk overview

Under this drive scheme, the cyclotron state will undergo a quantum walk along the ladder of energy levels. A quantum walk is the quantum mechanical analogy of the classical random walk; quantum walks have been extensively studied in the context of quantum computing and search algorithms [92, 93].

In a classical random walk, a walker takes multiple steps, choosing a direction randomly before each step. In a quantum walk, a particle's position in some Hilbert space evolves with a direction determined by a random quantum variable. The thought experiment first used to describe the quantum walk [94] is relevant to our system. The spin of a spin-1/2 particle is put into a superposition at each step. The spin state then determines which direction the particle's position or momentum evolves – i.e., if the particle spin is up, it moves left, while if its spin is down it moves right. As this

process is repeated, the wavefunction evolves into a superposition of states, which is collapsed when the position or momentum is measured.

Returning to the proton, each measurement of the quantum walk will select a state at random from this superposition, depending on the probability amplitude. The distribution of multiple measurements reveals the properties of the underlying process. In the procedure of figure 6.2, an individual measurement gives the change in cyclotron quantum number Δn_c during the drive time. Over multiple measurements, we will determine the standard deviation $\sigma(\Delta n_c)$, defined as

$$\sigma(\Delta n_c) = \sqrt{\langle \Delta n_c^2 \rangle - \langle \Delta n_c \rangle^2} = \sqrt{\left[\sum_n (n - n_0)^2 \left(C^*_{n,\uparrow}(t) C_{n,\uparrow}(t) + C^*_{n,\downarrow}(t) C_{n,\downarrow}(t) \right) - \left(\sum_n (n - n_0) \left(C^*_{n,\uparrow}(t) C_{n,\uparrow}(t) + C^*_{n,\downarrow}(t) C_{n,\downarrow}(t) \right) \right)^2 \right]} \tag{6.23}$$

This gives us the standard deviation of the superposition of states created during the random walk, and thus the combined transition rate as a function of the drive frequency ratios.

The standard deviation of a quantum walk has been studied, as a quantum mechanical property of the wavefunction [95, 96]. While the standard deviation of a one-dimensional classical random walk grows as $\sqrt{T}$ [97], the standard deviation of a one-dimensional quantum walk grows linearly with T [95].

6.3.3 Quantum walk of the cyclotron state

Equation 6.22 gives an infinite set of coupled differential equations. These equations could be solved using adiabatic elimination – since we begin in a well-defined

initial state, we can neglect distant states with large $(n - n_0)$. However, the analytic solutions quickly lose utility as we increase the number of terms. We will first consider an analytically solved case from the literature, which we will compare to a numerical solution of equation 6.22. We will then use numerical solutions of equation 6.22 with different possible parameters of the experiment at Northwestern, such as magnetic field drift and cyclotron uncertainty, to characterize the measurement method.

A difference between equation 6.22 and the standard quantum random walk (e.g. of [94]) is the cyclotron detuning term $(n-n_0)\Delta_{cd}$. This represents a phase difference accumulated as the cyclotron state steps away from the origin, due to the different phase rotation rates of the energy levels. Quantum walks of this form are studied in reference [98]. The phase difference per step causes interference in the superposition, with the interference increasing as the superposition evolves away from the origin and as the evolution time increases. Reference [98] calls this the "harmonic case" of a quantum walk, because the phase difference term behaves like a restoring force towards the original state. For a quantum walk involving with this sort of phase difference per step, in the case where the "harmonic" term $(n_0 - n)\Delta_{cd} << 2\pi$, reference [98] derives the coupled equations

$$\frac{d}{dt}C_{c,n} \approx in\Phi C_{c,n} + i\frac{d}{2}\left(C_{c,n-1} + C_{c,n+1}\right) \tag{6.24}$$

where the Φ term represents the phase difference or effective harmonic term, and the d term represents the step probability. These equations have the exact solutions in terms of Bessel functions [99]

$$C_n(t) = J_{-n}\left(\frac{2d}{\Phi}\sin\left[\frac{\Phi}{2}t\right]\exp\left[i\frac{n}{2}(\Phi t - \pi)\right]\right) \tag{6.25}$$

This can be thought of as a toy model for our system, ignoring the existence of two spin states and dealing only with the cyclotron ladder. If we remove the spin-flip terms and add the two parts of equation 6.22 together, we get equation 6.24. This lets us identify the parameters in equation 6.25 with our system: d=$\Omega/2$ and $\Phi = \Delta_{cd}/2$.

The upper panels of figure 6.4 show the probability distribution derived from numerical solution of equation 6.22, while the lower panels compare properties of this solution with equation 6.24. Besides reinforcing the conceptual identification of our system with a quantum walk, we identify two important parameters from the analogy. First, at small times or detunings such that $T << 1/\Delta_{cd}$, $\sigma(\Delta n_c)$ grows linearly with time, as expected for the quantum walk. By taking the limit as $\Delta_{cd} \to 0$ of $d\sigma/dt$ using equations 6.25 and 6.23, we can identify the slope of this growth as $\Omega/2\sqrt{2}$. Second, we identify a timescale at which σ starts to decrease, as the effective harmonic term causes the probability to return towards the origin. In quantum walk literature, this is known as a recurrence. The timescale can be found from equation 6.25 – taking the n=0 term gives probability of a return to the origin, and we find that this probability begins to increase (and σ begins to decrease) after time $T = 2\pi/\Delta_{cd}$.

Leaving the toy model, we will relax the restriction that $\Omega_A = \Omega_{sf}$ and investigate numerical solutions for equation 6.22 with detunings set to zero. Figure 6.5 demonstrates that the growth of $\sigma(\Delta n_c)$ depends on the lower of the two Rabi frequencies.

Plotting the slope versus the lower of the two Rabi frequencies (fig. 6.6) reveals a linear relationship. This has the same slope $\Omega/2\sqrt{2}$ derived from 6.25, but with the generic Ω replaced by $\min[\Omega_A, \Omega_{sf}]$.

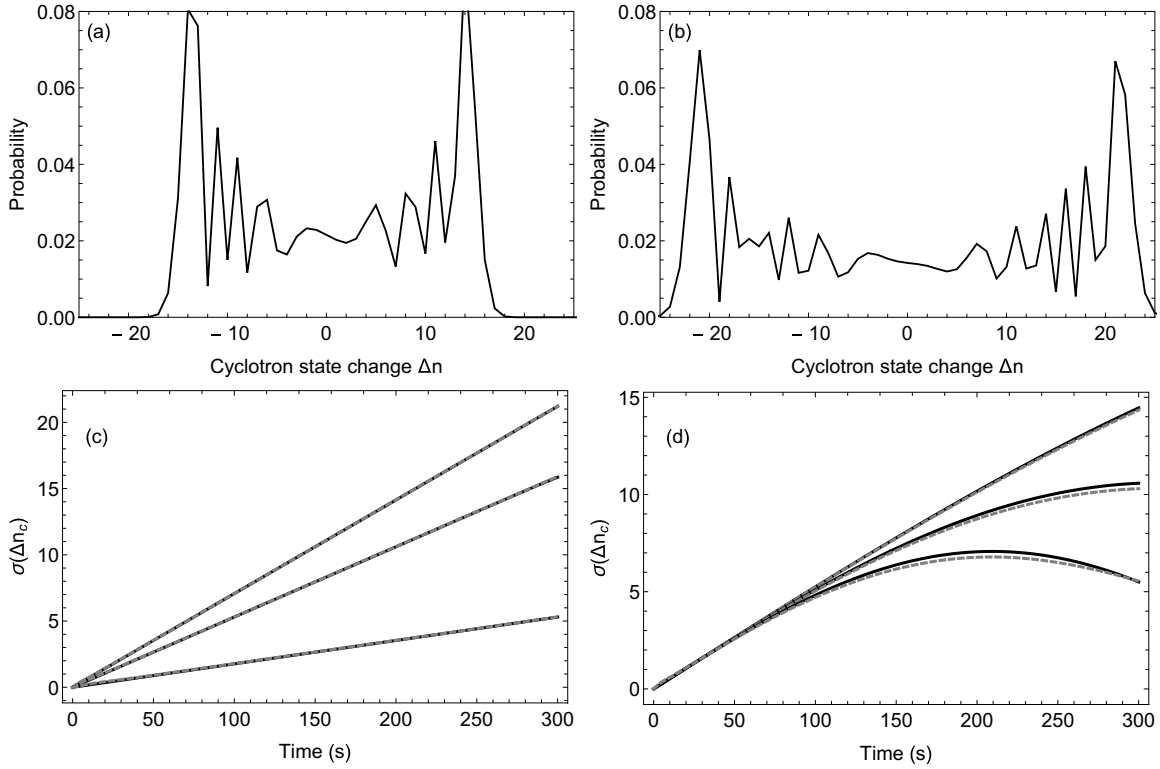

Figure 6.4: Quantum walk in the cyclotron state under simultaneous drives.
(a) Probability of measuring Δn_c after 200 second evolution time with $\Omega_A = \Omega_{sf} = 150$ mrad/s, $\Delta_{sf} = \Delta_{cd} = 0$, demonstrating characteristic behavior of a quantum walk (see e.g. [93])
(b) Probability of measuring Δn_c after 300 second evolution time.
(c) Standard deviation vs time for 200 mrad/s (top), 150 mrad/s (middle) and 50 mrad/s (bottom) Rabi frequencies, $\Delta_{sf} = \Delta_{cd} = 0$. Black lines are numerical solutions to equation 6.22; dashed grey lines are solutions to equation 6.25.
(d) Standard deviation vs time for 150 mrad/s Rabi frequencies, with $\Delta_{sf} = 0$ and $\Delta_{cd} = 10$ (top), 20 (middle), or 30 mrad/s (bottom).

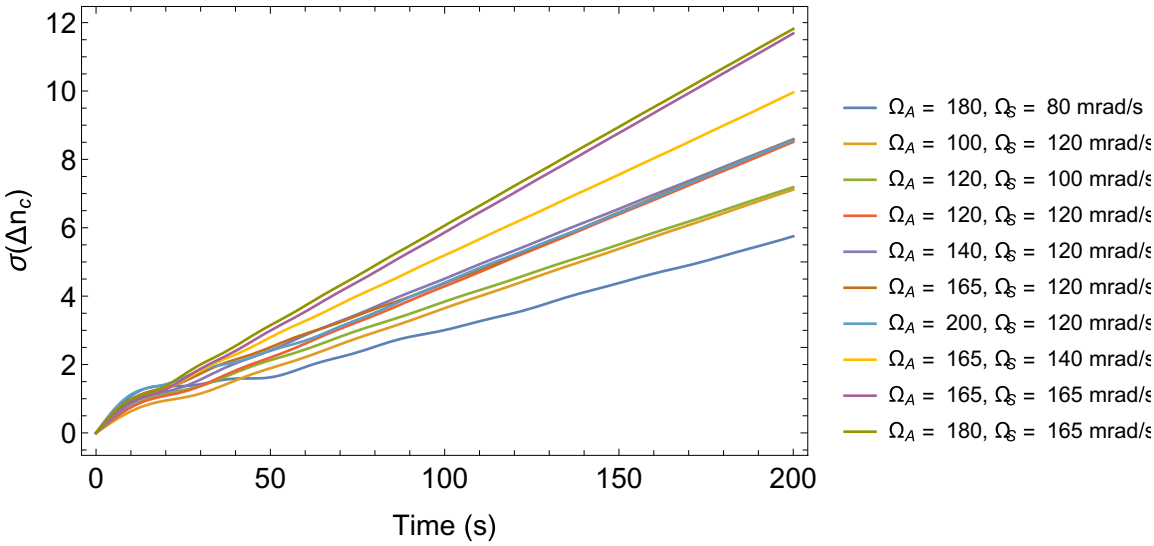

Figure 6.5: $\sigma(\Delta n_c)$ vs time, for different combinations of Rabi frequencies.

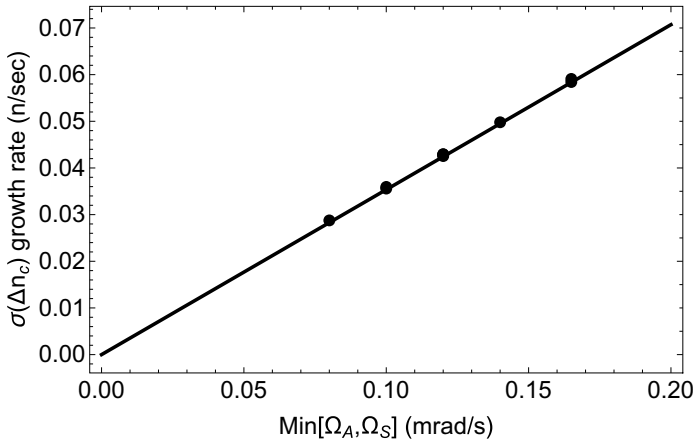

Figure 6.6: Growth in $\sigma(\Delta n_c)$ vs smaller of the two Rabi frequencies. Solid line is $y = \frac{1}{2\sqrt{2}}x$, matching the prediction of our analytically solved toy model (equation 6.25).

6.4 Lineshape and Precision

6.4.1 Detuning and Power Broadened Linewidth

We can find the linewidth by examining the signal for different detunings. In figure 6.4 we identified the effect of nonzero Δ_{cd} – at short times, the signal accumulates linearly, but at times approaching $2\pi/\Delta_{cd}$ the signal growth rate decreases. In practice, we will hold Δ_{cd} close to zero, adjusting the drive frequencies as indicated by pulsed SOF measurements of ω_c using the magnetometer proton. We will measure the ratio ω_{ad}/ω_{sd} by varying the "ratio detuning" $\Delta_r = \frac{\omega_{ad}}{\omega_{sd}} - \frac{\omega_a}{\omega_s} = \frac{\omega_a+\Delta_{ad}}{\omega_s+\Delta_{sf}} - \frac{\omega_a}{\omega_s}$. Our equations of motion were in terms of Δ_{sd} and Δ_{cd}, so we find an expression for Δ_r:

$$\Delta_{sf} = \frac{\omega_s\Delta_{cd} + \omega_s^2\Delta_r}{\omega_s - \omega_a + \omega_s\Delta_r} \approx \frac{\omega_s\Delta_{cd} + \omega_s^2\Delta_r}{\omega_c} = \frac{g}{2}\left(\Delta_{cd} + \omega_s\Delta_r\right) \tag{6.26}$$

(since $\Delta_r \sim 10^{-10} << 1$), or:

$$\frac{\Delta_{sf}}{\omega_s} = \frac{g}{2}\left(\frac{\Delta_{cd}}{\omega_s} - \Delta_r\right) \rightarrow \frac{\Delta_r}{r} = \frac{1}{\omega_a/\omega_s}\left(\frac{\Delta_{cd}}{\omega_s} - \frac{2}{g}\frac{\Delta_{sf}}{\omega_s}\right) \tag{6.27}$$

$$\frac{\Delta_r}{r} = \left(\frac{\Delta_{cd}}{\omega_a} - \frac{2}{g}\frac{\Delta_{sf}}{\omega_a}\right) \tag{6.28}$$

In figure 6.7 we plot the lineshape with $\Delta_{cd} = 0$ for different values of Δ_{sf}. Each point is the slope of $\sigma(\Delta n_c)$ obtained from a numerical solution of equation 6.22 with the appropriate parameters. Because it depends on the excursion of the quantum random walk, the lineshape we obtain from these equations of motion is not a pure Gaussian or Lorentzian. As figure 6.7 shows, it approximately matches the square root of a Lorentzian, but is more sharply peaked.

Calculating the standard deviation lineshape with different Rabi frequencies shows (fig. 6.8) that while the signal on resonance depends on the weaker of the two Rabi

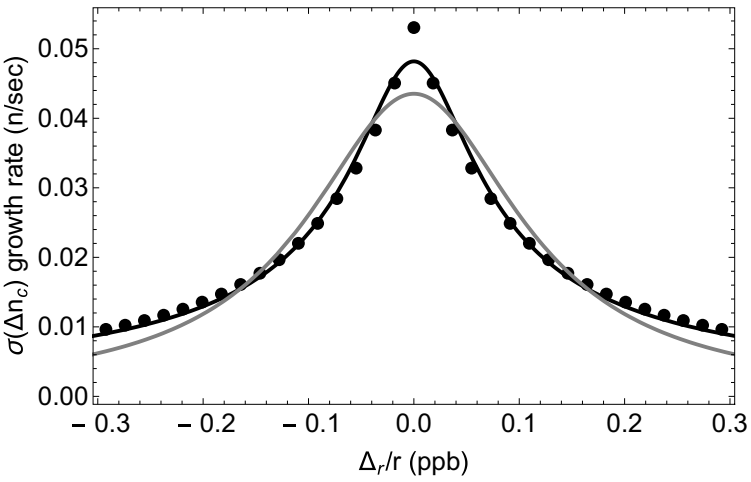

Figure 6.7: Growth in $\sigma(\Delta n_c)$ vs detuning, for 150 mrad/s Rabi frequencies. Gray line: best fit parameters to $A\frac{1}{\pi\gamma\left(1+\left(\frac{x}{\gamma}\right)^2\right)}$. Black line: best fit parameters to $A\sqrt{\frac{1}{\pi\gamma\left(1+\left(\frac{x}{\gamma}\right)^2\right)}}$

frequencies, the linewidth depends on power broadening due to both. We conclude that using equal Rabi frequencies is the optimal configuration – increasing either one relative to the other only broadens the lineshape without increasing the signal rate.

Figure 6.9 demonstrates how this dependence on both drive strengths increases with the detuning. Lacking an analytical expression for the lineshape that replicates this behavior, for the following estimates we will use the full width at half maximum (FWHM) of the lineshape to quantitatively evaluate the linewidth or precision of a measurement using different parameters.

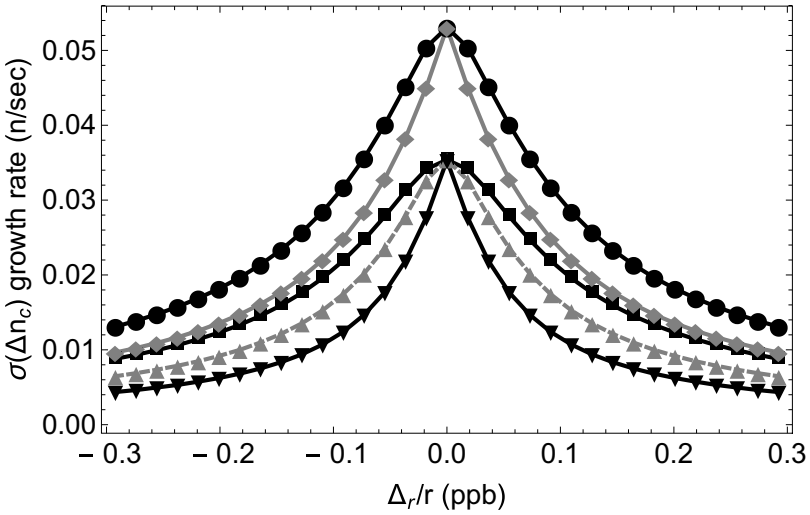

Figure 6.8: Growth rate of $\sigma(\Delta n_c)$ vs detuning for different Rabi frequencies.
Black circles: $\Omega_A = 207$ mrad/s, $\Omega_s = 150$ mrad/s, FWHM $\approx$.3 ppb.
Gray diamonds: $\Omega_A = 150$ mrad/s, $\Omega_s = 150$ mrad/s, FWHM $\approx$.2 ppb.
Black squares: $\Omega_A = 207$ mrad/s, $\Omega_s = 100$ mrad/s, FWHM $\approx$.3 ppb.
Gray triangles: $\Omega_A = 100$ mrad/s, $\Omega_s = 150$ mrad/s, FWHM $\approx$.2 ppb.
Black triangles: $\Omega_A = 100$ mrad/s, $\Omega_s = 100$ mrad/s, FWHM $\approx$.15 ppb.

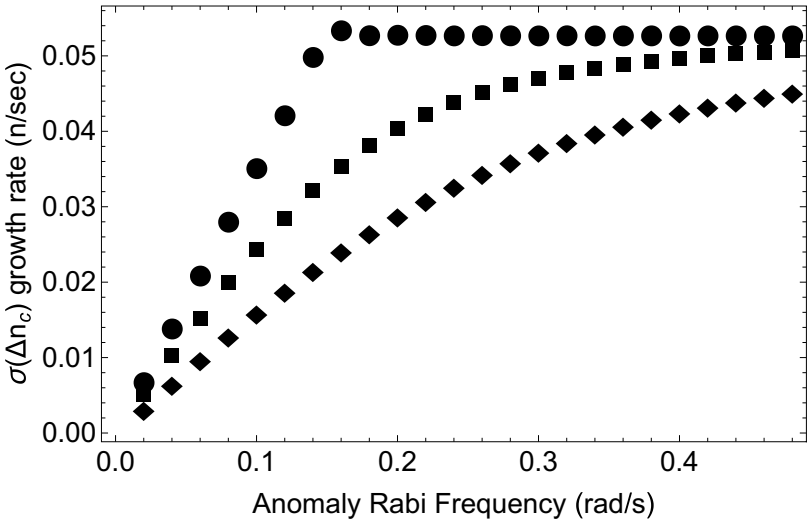

Figure 6.9: Rate of change of $\sigma(\Delta n_c)$ vs anomaly Rabi frequency for different detunings, with spin Rabi frequency held at 150 mrad/s. This demonstrates the change in lineshape with Rabi frequencies seen in figure 6.8. Circles: $\Delta_r/r = 0$. Squares: $\Delta_r/r = .05$ ppb. Diamonds: $\Delta_r/r = .1$ ppb.

6.4.2 Estimating Background - Noise and Cyclotron State Change

We now have an idea of the signal amplitude and linewidth for given drive parameters. There are two more factors which we need to consider before we can make our final estimates to evaluate this method – the accuracy with which we can identify n_c, and the effect of magnetic field drift.

The cyclotron quantum number n_c should be measured using the SEO in the analysis trap. Typical changes in the quantum number between measurements should be on the order of 10 to 100, corresponding to frequency shifts in the magnetic bottle of between 0.5 and 5 Hz. We therefore do not anticipate requiring any broad scans to find the axial frequency, as for the ν_+ measurement proposed in section 5.2; instead, we should be able to return the proton to consistent voltage settings and activate the SEO. Including sideband cooling, we estimate approximately 3 minutes required to find n_c in the analysis trap as discussed in section 5.3.

Since our signal is the change in n_c after applying drives in the precision trap, the accuracy with which we can measure this change using the analysis trap determines the level of signal which we can resolve.

The analysis of noise sources on this measurement largely follows the discussion in section 5.2.1. However, we are concerned with smaller changes in quantum number; and we are using a much smaller radius, so many of the background sources scale favorably. There are a few additional relevant background sources that did not apply to the Ramsey cyclotron frequency measurement discussion.

We treat the magnetic moment shift due to Rabi oscillations from a single resonant

drive (i.e., if one drive is resonant but the other is not, meaning the ratio under test is incorrect) as a contributor to noise in Δn_c. A single spin flip yields an axial frequency change equivalent to three cyclotron quanta, and an anomaly transition yields four cyclotron quanta. The direction of this shift is random since we do not characterize the initial spin state, so we add it as ± 2 cyclotron quanta in quadrature with other noise sources. Additionally, we incorporate background transitions in the precision trap, at the rate calculated in section 5.2.1 for a cold proton.

Table 6.1: Sources of uncertainty in determining n_c for anomaly drive measurements

Source	Effective Δ_n
Trapping voltage drift in either trap	.0004/sec
Background transitions in precision trap	.005/sec
Background transitions in analysis trap	.02/sec
Transfer hysteresis or voltage shift	3
Frequency determination uncertainty	1
Spin or anomaly Rabi oscillations	2
Magnetron cooling distribution	3

We can look at each of these as independent random sources of noise affecting our measurement of the number of n_c changes excited by the drive field. The noise sources are summarized in table 6.1. They would then add in quadrature, and we would have an effective "background signal," in units of effective cyclotron state changes:

$$\sigma_{noise} = \sqrt{\left(.0004(t_P + t_A)\right)^2 + (.005t_P)^2 + (.02t_A)^2 + 3^2 + 1 + 2^2 + 3^2} \qquad (6.29)$$

where t_P and t_A are the total time between n_c measurements spent in the precision and analysis traps, respectively.

Say that we drive in the precision trap for 11 minutes and that we take 3 minutes in the analysis trap to determine the axial frequency. We would then have, again in units of effective cyclotron state changes:

$$\sigma_{noise} = \sqrt{(.0004 * 840)^2 + (.005 * 660)^2 + (.02 * 180)^2 + 23} \approx 7 \tag{6.30}$$

6.4.3 Magnetic Field Drift and Signal-To-Noise

The standard deviation $\sigma(\Delta n_c)$, on resonance and with 150 mrad/s Rabi frequencies, is approximately .053 transitions per second. Over our 11 minutes of drive time, this would give a standard deviation of 35 cyclotron levels. However, we still need to account for the effect of magnetic field drift. In the scenario where the magnetic field drifts at 0.5 ppb per hour, during the drive time we would expect it to drift by .09 ppb, slightly reducing the average transition rate. Solving equation 6.22 with Δ_r drifting at this rate gives $\sigma(\Delta n_c)$ ~32 after 11 minutes.

Since the signal is a standard deviation, it adds in quadrature with the standard deviation due to noise sources. Figure 6.10 shows the approximate total standard deviation in measured cyclotron quantum number, including the noise estimate as well as the effect of magnetic field drift.

Because we would be measuring a standard deviation, we will need to make multiple measurements at each ratio value. The standard deviation of a measured sample standard deviation σ after N measurements is

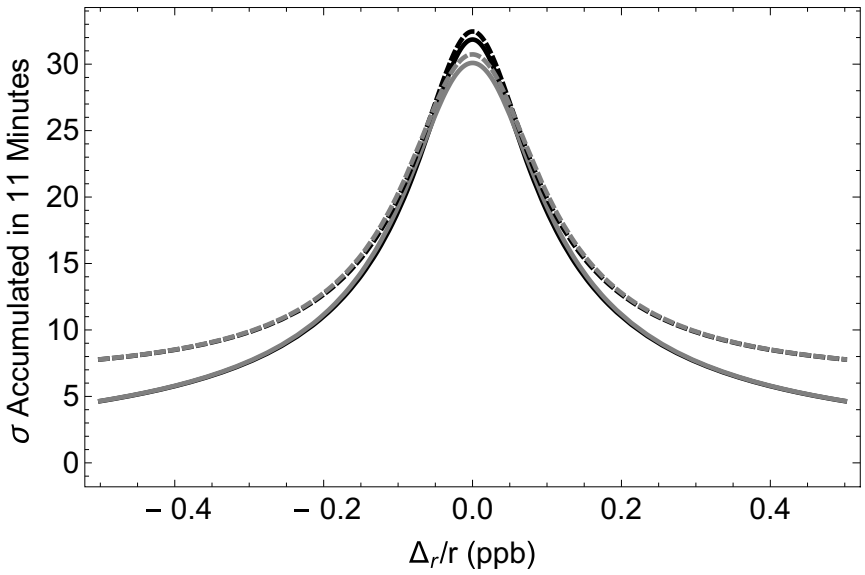

Figure 6.10: Cyclotron standard deviation after 11 minutes with 150 mrad/s Rabi frequencies. Black lines without drift, grey lines include 0.5 ppb/hour magnetic field drift. Solid lines without noise, dashed lines include estimated noise added to signal in quadrature.

$$SD(\sigma) = \sigma\sqrt{1 - \frac{2}{N-1} * \left(\frac{\Gamma(\frac{N}{2})}{\Gamma(\frac{N-1}{2})}\right)^2} \tag{6.31}$$

Equating this standard deviation to an error bar, after ten trials on resonance we should have a result of $\sigma(\Delta_n) = 32 \pm 8$.

We can use equation 6.29 to calculate the "noise" transitions accumulated during each measurement (σ_{noise}), which we will compare to the accumulated signal obtained from solving equations 6.22. The signal-to-noise ratio on this measurement is then

$$SNR = \frac{\sqrt{\sigma_{signal}^2 + \sigma_{noise}^2} - \sigma_{noise}}{\sqrt{\sigma_{signal}^2 + \sigma_{noise}^2}\sqrt{1 - \frac{2}{n-1}\left(\frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})}\right)^2}} \tag{6.32}$$

After ten trials, on resonance and with drive time 11 minutes each, this gives an SNR $\frac{\sqrt{32^2+7^2}-7}{\sqrt{32^2+7^2}(.23)} \approx 3.3$.

6.5 Cyclotron detuning Δ_{cd}

Our ultimate goal is to carefully calculate the rate of data acquisition and linewidth using this proposed measurement method. To do this, one final element remains to be analyzed – we must develop a framework to account for the effects of nonzero Δ_{cd}.

In the preceding sections we have assumed $\omega_{sf} - \omega_{ad} = \omega_{cd}$ is held constant and exactly at the proton's cyclotron frequency ω_c. In an actual experiment, magnetic field drift means that there will be some nonzero detuning Δ_{cd} during part of the trial. The effect of this was shown in figure 6.4; effectively, the timescale of recurrence towards the origin sets a limit on how close to zero Δ_{cd} must be to acquire signal. After a time $\frac{2\pi}{\Delta_{cd}}$ the standard deviation starts to decrease if the detuning remains constant.

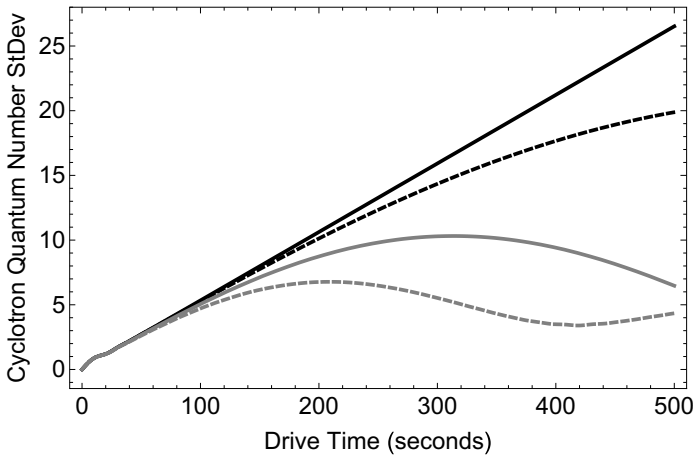

Figure 6.11: Numerical calculation of $\sigma(\Delta n_c)$ vs time, for $\Omega_A = \Omega_{sf}$ =150 mrad/s.
Black solid line: $\Delta_{cd} = 0$ mrad/s; Black dashed line: $\Delta_{cd} = 10$ mrad/s
Gray solid line: $\Delta_{cd} = 20$ mrad/s; Gray dashed line: $\Delta_{cd} = 30$ mrad/s

This sets a stricter limitation than power broadening on the acceptable Δ_{cd}. Acquiring a measurable signal thus implies all three frequencies ω_{cd}, ω_{sf} and ω_a are

resonant. In making our final estimates for precision and data acquisition rates, we have to account for the fact that Δn_c will only be significant while Δ_{cd} is close to zero. Essentially, ω_+ must be measured to determine the magnetic field, and the drive frequencies adjusted to match the measured frequency.

Three experimental parameters restrict how well we can hold Δ_{cd} close to zero during a drive sweep. First, we can only minimize Δ_{cd} as well as we can measure the cyclotron frequency. Second, Δ_{cd} will vary during the drive time according to the magnetic field drift. Finally, the predictability of the magnetic field drift – how closely it adheres to some linear or polynomial model – will set how frequently we have to interlace our drive trials with new measurements of ω_+ to keep Δ_{cd} small. In the following section, as we estimate data acquisition time and linewidth, we will evaluate different possible scenarios for each of these parameters.

We will predict ω_+ during the drive time from interleaved measurements, as per figure 6.2. We propose using the SOF method described in section 5.2, where two drive pulses separated by a time T interfere constructively or destructively with the free phase evolution during that time, mapping the cyclotron frequency onto the radius, which can be measured in the analysis trap. As in chapter 5, we will use two sets of pulses with slightly different detunings, to locate the position on the fringe.

The application of this method differs in two small ways from the proposal in chapter 5. Because we aren't trying to match a spin-flip Ramsey sequence, we don't need the full 10 second evolution time. A statistical uncertainty <10 mrad/s is achievable with T=1 second. We therefore can reduce the frequency of the shorter-T measurements (to make sure we know which fringe we are on) from once per hour to once or

twice per day. Second, we need to know ω_+ when the spin and anomaly drives are applied, rather than during the SOF evolution time. We therefore must account for possible magnetic field drift during the (on average) 13 minutes of axial frequency search in the analysis trap. We can do this with a third measurement, with the same detuning and evolution time as the second pair. This lets us measure the cyclotron frequency change which took place during the axial frequency search. A 0.1 ppb drift in the cyclotron frequency (13 minutes of drift at 0.5 ppb/hour) corresponds to a ∼50 mrad/s shift in ω_c. Assuming a 1.75 Hz detuning between the cyclotron frequency and the SOF pulses (as in [72]), equation 5.4 gives a search range of ∼45 Hz to identify the 0.1 ppb drift. This range is small enough that we can just activate the SEO in the analysis trap after the second pulse, rather than needing to search with drive scans.

For our estimates of linewidth and data collection time, we will consider three different scenarios for uncertainty on ω_+. The first is the same scenario used in chapter 5 – an uncertainty range of order a few mrad/s, with an appropriately chosen post-pulse energy and evolution time. In this scenario, our ability to minimize Δ_{cd} is mostly limited by the magnetic field drift. We also consider two scenarios where as-yet-unidentified systematic effects limit our ω_+ precision, with a range of 20 mrad/s or 40 mrad/s. In these scenarios, we ensure that our drives pass through resonance by covering the entire range, either by sweeping our drives while keeping their ratio constant, or by keeping the drives constant and allowing the magnetic field drift to carry Δ_{cd} through the range.

To account for magnetic field drift, we will solve equations 6.22 with time-dependent detunings. We will consider two scenarios, the "typical" 1 ppb/12 hour drift of the

previous electron magnetic moment measurement performed in this magnet [66], as well as the .5 ppb/hour drift observed for a minority of that measurement.

A final consideration is the predictability of the field drift. If it can be characterized by a linear or polynomial fit between measurements of ω_+, then we can interpolate and adjust our drives between trials to "follow" the drift. We will investigate the data collection time with different allowable intervals between measurements of ω_+.

6.6 Numerical estimates of statistical precision for different experimental parameters

Equation 6.32 gives the signal-to-noise for a certain time on resonance per trial and number of trials at a particular ratio. In the following estimates, σ_{signal} is calculated by numerically solving equations 6.22. We account for magnetic field drift by including time-dependent Δ_{sf} and Δ_{cd}.

To match the signal-to-noise criterion set in chapter 5, we evaluate the signal-to-noise at the Δ_{sf} value where $\sqrt{\sigma_{signal}^2 + \sigma_{noise}^2} - \sigma_{noise}$ is half that for $\Delta = 0$. Together with Δ_{cd}, this determines $\frac{\Delta_r}{r}$ and thus the width of the lineshape. With 150 mrad/s Rabi frequencies, we find this at approximately $\Delta_{sf,1/2} \approx .165$ rad/s, corresponding to $\Delta_r/r \approx .06$ ppb, for a FWHM of .12 ppb. The contribution from Δ_{cd} in equation 6.28 is to broaden this lineshape slightly. Signal will accumulate as magnetic field drift (or an applied chirp of the drive frequency) sweeps Δ_{cd} through 0, with the accumulation limited by the recurrence time $\frac{2\pi}{\Delta_{cd}}$. Recurrences during the drive time do not occur when $t < \frac{2\pi}{\Delta_{cd}}$. Incorporating the drift rate, this becomes

$t < \frac{2\pi}{t*d\omega/dt}$. Figure 6.12 demonstrates how the signal accumulates as Δ_{cd} drifts or is swept through resonance – at large detunings, we observe oscillating recurrences, while net signal only accumulates for small Δ_{cd}.

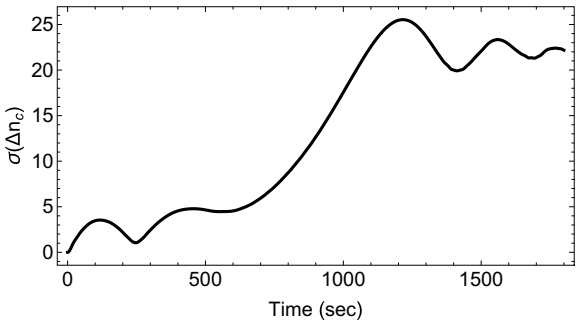

Figure 6.12: Cyclotron standard deviation accumulated as Δ_{cd} drifts through zero at a rate .5 ppb per hour. Calculated for $\Delta_{sf} = 0$.

To determine the true FWHM of the measured lineshape, we must incorporate a line broadening due to the range of Δ_{cd} values covered. We add double the calculated Δ_{cd} into equation 6.28, to account for a drift range above or below resonance.

For dB/dt = 1 ppb/12 hours, we have $d\omega/dt \sim$ 45 mrad/s/hour, giving t∼11.8 minutes, during which time Δ_{cd} covers a range of $\sim$ 9 mrad/s. This gives an FWHM of 0.15 ppb in Δ_r/r.

For dB/dt = 0.5 ppb/hour, we have $d\omega/dt \sim$ 270 mrad/s/hour, giving t∼5 minutes, during which time Δ_{cd} covers a range of 22 mrad/s. This gives an FWHM of 0.2 ppb in Δ_r/r.

The time-dependent detunings used to solve equations 6.22 are given by:

$$\begin{aligned} \Delta_{sf} &= \left(\frac{d\omega_s}{dt}\left(\frac{-T}{2}+t\right)+\Delta_{sf,1/2}\right) \\ \Delta_{cd} &= \left(\frac{d\omega_c}{dt}\left(\frac{-T}{2}+t\right)\right) \end{aligned} \tag{6.33}$$

where T is the total time that the drives are applied, $d\omega_s/dt$ and $d\omega_c/dt$ are determined by the magnetic field drift or drive chirp rate, and $\Delta_{sf,1/2}$ is the center value where $\sqrt{\sigma_{signal}^2 + \sigma_{noise}^2} - \sigma_{noise}$ is half that for $\Delta_{sf} = 0$.

We will use equation 6.29 to calculate σ_{noise}, with t_P and t_A being the time between n_c measurements spent in the precision and analysis traps respectively. Finally, for the number of measurements n in equation 6.32, we take the total cycle time T_{tot}, including ν_+ measurements performed on the magnetometer proton, drive time t_P, and time required for n_c measurements in the analysis trap before and after the drive time. The drift rate and ν_+ measurement precision will determine what fraction of the drive time is spent on resonance and how much time is spent on cyclotron-frequency measurements versus g-factor measurements, giving us an estimate of SNR versus time for each case.

Case I: 1 ppb / 12 hour drift rate, precise ω_+ measurements

With the low "typical" drift rate previously observed in this magnet, we use a drive time of $\sim$ 11 minutes to optimize the number of trials. We evaluate different scenarios for how well we can interpolate between ω_+ measurements. If we can make three drive trials between each SOF ω_+ procedure, we obtain a SNR of 3 at $\Delta_{sf,1/2}$ after 3.8 hours. If we can only interpolate every other measurement, this would take 4.4 hours; and if we required an ω_+ measurement after every drive trial, we would meet our SNR criterion after 6.4 hours.

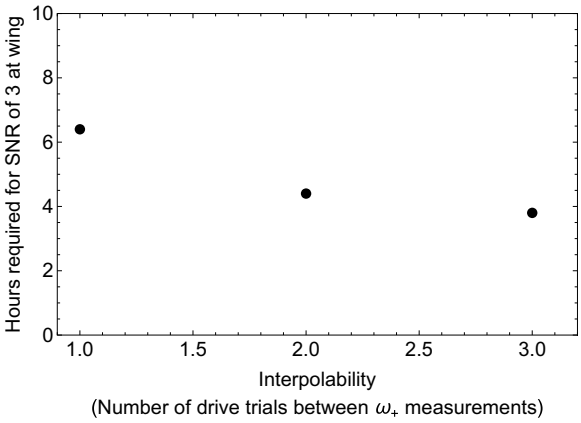

Figure 6.13: Time to meet SNR criterion vs cyclotron-frequency interpolability.

Case II: 1 ppb / 12 hour drift rate, 20 mrad/s ω_+ sweep range

If our ability to measure ω_+ using the SOF procedure is limited by as-yet-undetermined systematics, we may have to sweep over a larger range of detunings to acquire a signal. To cover a 20 mrad/s range at 1 ppb/12 hours requires a 25 minute drive time. As above, we also consider ability to interpolate between ω_+ measurements – for example, if the center of the 20 mrad/s range can be predicted by a polynomial.

In addition to increasing the time required for each trial, the longer drive time slightly increases the estimated σ_{noise} in equation 6.32, as background transitions and voltage drift slowly accumulate while the proton is in the precision trap. Incorporating this, we find that obtaining an SNR of 3 at the half-maximum point would require 6.7 hours with a ω_+ measurement every third trial, 7.4 hours with every second trial, and 9.6 hours with an ω_+ measurement after every drive trial.

Case III: 1 ppb / 12 hour drift rate, 40 mrad/s ω_+ sweep range

A larger range of uncertainty to cover requires a larger drive time. At 1 ppb/12 hours, a 40 mrad/s range would require 50 minutes of drive time each trial. Considering the increased time per trial, we find that meeting our SNR criterion would require 19 hours with a ω_+ measurement every third trial, 20.3 hours with every second trial, and 24 hours with an ω_+ measurement after every drive trial.

Case IV: .5 ppb / hour drift rate, precise or 20 mrad/s ω_+ measurement

The higher drift rate means we spend less time per trial accumulating signal, requiring a larger number of trials to meet our SNR criterion. The time required is therefore somewhat more sensitive to interpolation between ω_+ measurements. The background signal is slightly lower, and the time per trial is also smaller, but this does not compensate for the reduced signal accumulated per trial. (Recall that the signal per trial increases linearly with time, whereas the statistical uncertainty reduces with the square root of the number of trials). Balancing signal accumulation per trial with the number of trials, we find that the optimal drive time is approximately 7 minutes.

We find that obtaining an SNR of 3 at the half-maximum point would require 5.1 hours with a ω_+ measurement every third trial, 6.2 hours with every second trial, and 9.6 hours with an ω_+ measurement after every drive trial.

As mentioned above, during the 5 minute drive time with 0.5 ppb/hour drift or sweep rate, we cover a range of ~ 22 mrad/s. The results for a 20 mrad/s cyclotron uncertainty range are therefore the same as for arbitrarily precise ω_+ measurement. We can compare the time required to that for the above case of 1PPB/12 hour drift

and a 20 mrad/s range; if we are able to interpolate for multiple measurements, sweeping at .5 ppb/hour achieves our SNR criterion faster than the slower 1 ppb/12 hour drift. Once experimental parameters are measured, determining the proper sweep rate and drive time will be an optimization problem, of which each case in this section is a specific example.

Case V: 0.5 ppb / hour drift rate, 40 mrad/s ω_+ sweep range

Here we have a more dramatic example of the optimization choices between different sweep rates. Sweeping through a 40 mrad/s range at 0.5 ppb / hour requires only $\sim$9 minutes, compared to 50 at the lower drift rate. This is only slightly longer than the optimal drive time determined for the case of precise ω_+ measurements. We find that meeting our SNR criterion would require 5.4 hours with a ω_+ measurement every third trial, 6.5 hours with every second trial, and 9.7 hours with an ω_+ measurement after every drive trial. This indicates that in the presence of larger systematic uncertainty in ω_+, we would significantly benefit from intentionally sweeping our drives rather than following the field drift rate.

6.7 Comparison between Anomaly and Single-Spin-Flip methods

Both the simultaneous SOF and the anomaly-drive methods measure the spin, cyclotron or anomaly, and axial frequencies at the same time and under the same trapping fields, and so neither should have a significant advantage over the other in

terms of systematic uncertainties. In sections 5.3 and 6.6, we have attempted to make detailed numerical estimates of the data-acquisition time required for a measurement with a 0.2 ppb FWHM lineshape. We have also attempted to include guiding principles for adapting those estimates to achievable experimental conditions. In this section, we will first summarize both estimates under the optimistic but achievable conditions used for our numerical estimates. We will then briefly summarize the most relevant experimental parameters, to identify how the comparison between sub-PPB methods changes if the achievable level is better or worse than estimated here.

We have already adapted our estimates to two different conditions for magnetic field stability – the 1 ppb / 12 hour "typical" rate and the 0.5 ppb/hour "atypical" rate observed during the previous measurement of the electron magnetic moment in this magnet [66]. Under the lower drift rate, the simultaneous anomaly-spin flip drive measurement has a significant advantage over the separated oscillatory fields method. If the magnetic field drift can be interpolated reliably between cyclotron-frequency measurements, for example using a polynomial fit, then the anomaly-drive method is a factor of 6 faster. If the magnetic field cannot be interpolated and the cyclotron frequency must be re-measured after every drive trial, this advantage is reduced to a still-significant factor of 4.

Under the higher drift rate, this advantage is reduced, although it still depends on whether we can interpolate between ω_+ measurements. With no interpolation, the anomaly-drive is faster only by a factor of order 3. However, if we can interpolate and make two trials before measuring ω_+, this improves to better than a factor of 7.

So, under the estimated parameter set, the anomaly-drive method outperforms

the SOF method, especially at the lower predicted field drift rate. Without making further detailed numerical estimates, we will now consider how divergence from our assumptions affects the comparison between the two methods.

First, we consider if the magnetic field stability in this magnet at Northwestern cannot match even the worse stability observed in the same magnet during the electron measurement. As field drift gets worse, the time required for the SOF method scales more favorably than the anomaly method. With higher drift rates, the maximum drive time in the anomaly method becomes shorter, and the acceptable Δ_{cd} becomes larger; this means the constraints on interpolating become less strict.

We must also consider short-term magnetic field fluctuations on various timescales. Fluctuations on timescales much shorter than the drive time would average out, but could reduce the transition rate by increasing Δ_{cd} for part of the drive time. Fluctuations on timescales at or longer than the drive time would have a similar effect to the overall drift rate, limiting the maximum time on resonance.

We considered the effect of additional cyclotron-frequency measurement uncertainty on the anomaly-drive method. While the required sweep over the range of possible cyclotron frequencies increases the time required, the effect on the simultaneous SOF method is worse – uncertainty in ω_+ directly broadens the linewidth, by adding that uncertainty to the location of every spin-flip trial on the g-factor resonance. With the anomaly-drive method, both drives must be within the power-broadened linewidths regardless of the cyclotron frequency measurement.

Both methods propose making fast measurements of the axial frequency in the analysis trap, by transferring a particle to a trap with a known voltage and turning

on the SEO. In order to do this, we would have to be able to quickly settle the trapping potential in the analysis trap. We propose to do this with a "ringing in" procedure designed to get around the long electrode time constants, or alternatively by changing the relative potentials on other electrodes while keeping the analysis trap voltage constant. If neither of these works, it would be necessary to add settling time to each experiment cycle.

Measuring the cyclotron radius and the spin state after a trial each require extremely precise axial frequency measurements, so the time increase is roughly equal. However, the simultaneous SOF method requires an average of five additional analysis-trap frequency measurements per cycle, to select a sufficiently cold sub-thermal proton for spin-state readout. Overall, the anomaly-drive method would adapt better to this restriction.

We have also made some plausible but optimistic assumptions about how quickly we can determine the axial frequency in the analysis drive after an SOF measurement of the cycltoron frequency. If this takes longer than expected, it would affect both methods proportionally to the number of cyclotron-frequency measurements required. Thus, it would affect the anomaly-drive method in the low-drift case significantly less than other cases and methods; the higher rate of signal accumulation means fewer trials are required to attain the desired SNR. However, it would affect the simultaneous SOF method and the anomaly method in the high-drift case relatively equally, as both require relatively frequent measurements of the cyclotron frequency. (This is another case that improves if we can interpolate the magnetic field to make multiple anomaly-drive trials between cyclotron-frequency measurements.)

Other than analysis trap voltage settling and stability after transfers, several factors contribute to the resolution with which we can measure the cyclotron radius, including background cyclotron transitions, drift in the axial voltage source during the drive time, and the magnetron sideband cooling distribution width. None of these parameters affect the simultaneous SOF method, so if they are worse than expected that will favor SOF, and if they are better it will favor the anomaly drive.

Conversely, the simultaneous SOF method requires identifying single spin flips, while the anomaly drive method does not. Sub-unity efficiency or fidelity of spin-state identification thus make SOF a less attractive option. If adiabatic fast passage cannot be implemented, the time required for an SOF measurement would increase by a factor of 4 without significant improvements to axial stability.

6.8 Conclusion

We have proposed a new method for measurement of the proton and antiproton magnetic moments which relies on measuring the standard deviation of a quantum walk in the energy of a particle's cyclotron state. Detuning of either or both drives reduces the transition rate of this quantum walk and the spread of energies. Measuring its standard deviation at different drive frequency ratios gives a g-factor resonance. We have derived equations of motion for the spin and cyclotron wavefunctions acted on by simultaneous anomaly and spin-flip drives, which describe the quantum walk. We have solved these equations numerically and characterized the time required to make a g-factor measurement of a given target precision with different sets of experimentally achievable parameters. Finally, we have compared the proposed method to

the previously proposed method of simultaneous separated oscillatory fields, demonstrating the promise of significant improvement – an estimated factor between 3 and 7 – in the acquisition rate of statistics for a 0.2 ppb FWHM measurement. We also provided guidelines for evaluating which method is optimal under real experimental conditions when those are determined at Northwestern.

6.9 Appendix: rotating wave approximation for simultaneous drives

We start with the Hamiltonian for a proton addressed by drives near the spin and anomaly frequencies. Our goal is to use unitary transformations and the rotating wave approximation to remove the Hamiltonian's time dependence, so that we can obtain a simple set of differential equations governing the particle's quantum state.

$$\begin{aligned} H = \sum_n \Bigg[& -\frac{\hbar\Omega_A}{2}\left(e^{-i\omega_{ad}t} + e^{i\omega_{ad}t}\right)\left(|n,\uparrow\rangle\langle n+1,\downarrow| + |n,\downarrow\rangle\langle n-1,\uparrow|\right) \\ & - \frac{\hbar\Omega_S}{2}\left(e^{-i\omega_{sd}t} + e^{i\omega_{sd}t}\right)\left(|n,\uparrow\rangle\langle n,\downarrow| + |n,\downarrow\rangle\langle n,\uparrow|\right) \\ & + \hbar\left(-\frac{\omega_s}{2} + \left(n+\frac{1}{2}\right)\omega_c\right)|n,\downarrow\rangle\langle n,\downarrow| + \hbar\left(\frac{\omega_s}{2} + \left(n+\frac{1}{2}\right)\omega_c\right)|n,\uparrow\rangle\langle n,\uparrow| \Bigg] \end{aligned} \tag{6.34}$$

We re-define zero energy at the initial state $|n_0,\downarrow\rangle$ (and write the terms which only depend on the spin or cyclotron state in terms of $|n\rangle$ and $|\uparrow\rangle$):

$$\begin{aligned} \mathrm{H} = \hbar\sum_n & -\frac{\Omega_A}{2}\left(e^{-i\omega_{ad}t} + e^{i\omega_{ad}t}\right)\left(|n,\uparrow\rangle\langle n+1,\downarrow| + |n,\downarrow\rangle\langle n-1,\uparrow|\right) \\ & -\frac{\Omega_S}{2}\left(e^{-i\omega_{sd}t} + e^{i\omega_{sd}t}\right)\left(|\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow|\right) + (n-n_0)\omega_c|n\rangle\langle n| + \omega_s|\uparrow\rangle\langle\uparrow| \end{aligned} \tag{6.35}$$

We now apply a unitary transformation into a frame rotating at the spin-flip frequency. This should give us terms with no time dependence, as well as counter-rotating terms oscillating at $2\omega_{sd}$. The transformation matrix is

$$U_s = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\omega_{sd}t} \end{pmatrix} = |\downarrow\rangle\langle\downarrow| + e^{i\omega_{sd}t}|\uparrow\rangle\langle\uparrow| \tag{6.36}$$

The Hamiltonian under this transformation is given by

$$H_s = U_s H U_s^{-1} + i\left(\frac{\partial}{\partial t}U_s\right)U_s^{-1} \tag{6.37}$$

$$U_s^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & e^{-i\omega_{sd}t} \end{pmatrix} = |\downarrow\rangle\langle\downarrow| + e^{-i\omega_{sd}t}|\uparrow\rangle\langle\uparrow| \tag{6.38}$$

$$i\frac{\partial}{\partial t}U_s = i\begin{pmatrix} 0 & 0 \\ 0 & i\omega_{sd}e^{i\omega_{sd}t} \end{pmatrix} = -\omega_{sd}e^{i\omega_{sd}t}|\uparrow\rangle\langle\uparrow| \tag{6.39}$$

$$i\left(\frac{\partial}{\partial t}U_s\right)U_s^{-1} = \begin{pmatrix} 0 & 0 \\ 0 & -\omega_{sd} \end{pmatrix} = -\omega_{sd}|\uparrow\rangle\langle\uparrow| \tag{6.40}$$

Combining these gives the transformed Hamiltonian in the rotating frame, with spin terms at DC and at $2\omega_{sd}$ as expected:

$$\begin{aligned} &H_s = U_s H U_s^{-1} + i\left(\frac{\partial}{\partial t}U_s\right)U_s^{-1} \\ &= \hbar\sum_n -\frac{\Omega_A}{2}\left(e^{-i\omega_{ad}t} + e^{i\omega_{ad}t}\right)\left(e^{i\omega_{sd}t}|n,\uparrow\rangle\langle n+1,\downarrow| + |n,\downarrow\rangle\langle n-1,\uparrow|e^{-i\omega_{sd}t}\right) \\ &-\frac{\Omega_S}{2}\left((e^{2i\omega_{sd}t}+1)(|\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow|(1+e^{-2i\omega_{sd}t})\right) \\ &+(n-n_0)\omega_c|n\rangle\langle n| + (\omega_s - \omega_{sd})|\uparrow\rangle\langle\uparrow| \end{aligned} \tag{6.41}$$

If we only had time dependence at the spin-flip frequency, we could now apply the rotating wave approximation by dropping the $2\omega_{sf}$ terms. However, we still have

terms rotating at the anomaly frequency ω_a, keeping us without a clear separation into fast and slow components. We therefore perform another unitary transformation, to separate the anomaly-drive terms into constant and quickly rotating timescales in a new frame. Since our frame is are already rotating at the spin-flip frequency, to get into a frame rotating at the anomaly frequency we will apply a transformation rotating at the effective driven cyclotron frequency ω_{cd}, equal to the difference between the two drive frequencies $\omega_{sd} - \omega_{ad} = \omega_{cd}$:

$$\begin{aligned} U_c &= \sum_n e^{i(n-n_0)\omega_{cd}t} \left|n\right\rangle \left\langle n\right| && (6.42)\\ U_c^{-1} &= \sum_n e^{-i(n-n_0)\omega_{cd}t} \left|n\right\rangle \left\langle n\right| && (6.43)\\ i\frac{\partial}{\partial t}U_c &= -\omega_{cd}\sum_n (n-n_0) e^{i(n-n_0)\omega_{cd}t} \left|n\right\rangle \left\langle n\right| && (6.44)\\ i\hbar\frac{\partial U_c}{\partial t}U_c^{-1} &= -\hbar\omega_{cd}\sum_n (n-n_0) \left|n\right\rangle \left\langle n\right| && (6.45) \end{aligned}$$

The Hamiltonian under this transformation is given by

$$\begin{aligned} &H_a = U_c H_s U_c^{-1} + i\left(\frac{\partial}{\partial t}U_c\right)U_c^{-1} = \sum_n e^{i(n-n_0)\omega_{cd}t}\left|n\right\rangle\left\langle n\right| \times \\ &\Bigg(\hbar\sum_n -\frac{\Omega_A}{2}\left(e^{-i\omega_{ad}t} + e^{i\omega_{ad}t}\right)\left(e^{i\omega_{sd}t}\left|n,\uparrow\right\rangle\left\langle n+1,\downarrow\right| + \left|n,\downarrow\right\rangle\left\langle n-1,\uparrow\right| e^{-i\omega_{sd}t}\right) \\ &-\frac{\Omega_S}{2}\left((e^{2i\omega_{sd}t}+1)(\left|\uparrow\right\rangle\left\langle\downarrow\right| + \left|\downarrow\right\rangle\left\langle\uparrow\right|(1+e^{-2i\omega_{sd}t})\right) \qquad (6.46)\\ &+(n-n_0)\omega_c\left|n\right\rangle\left\langle n\right| + (\omega_s-\omega_{sd})\left|\uparrow\right\rangle\left\langle\uparrow\right|\Bigg)\times\sum_n e^{-i(n-n_0)\omega_{cd}t}\left|n\right\rangle\left\langle n\right| \\ &-\hbar\omega_{cd}\sum_n (n-n_0)\left|n\right\rangle\left\langle n\right| \end{aligned}$$

Distributing the sums gives

$$
\begin{aligned}
&H_a = \\
&\hbar\sum_n -\frac{\Omega_A}{2}\left(e^{-i\omega_{ad}t}+e^{i\omega_{ad}t}\right)\left(e^{i(\omega_{sd}+(n-n_0)\omega_{cd})t}\left|n,\uparrow\right\rangle\left\langle n+1,\downarrow\right| e^{-i((n-n_0+1)\omega_{cd})t}\right. \\
&\left.+\, e^{i((n-n_0)\omega_{cd})t}\left|n,\downarrow\right\rangle\left\langle n-1,\uparrow\right| e^{-i(\omega_{sd}+(n-n_0-1)\omega_{cd}t}\right) \\
&-\frac{\Omega_S}{2}\left((e^{2i\omega_{sd}t}+1)(\left|\uparrow\right\rangle\left\langle\downarrow\right|+\left|\downarrow\right\rangle\left\langle\uparrow\right|(1+e^{-2i\omega_{sd}t})\right) \\
&+(n-n_0)\omega_c\left|n\right\rangle\left\langle n\right|+(\omega_s-\omega_{sd})\left|\uparrow\right\rangle\left\langle\uparrow\right|-\hbar\omega_{cd}\sum_n(n-n_0)\left|n\right\rangle\left\langle n\right|
\end{aligned}
\tag{6.47}
$$

Distributing and combining terms using $\omega_{sd}-\omega_{ad}=\omega_{cd}$ gives

$$
\begin{aligned}
&\hbar\sum_n -\frac{\Omega_A}{2}\left(\left(1+e^{i2\omega_{ad}t}\right)\left|n,\uparrow\right\rangle\left\langle n+1,\downarrow\right|+\left|n,\downarrow\right\rangle\left\langle n-1,\uparrow\right|(e^{-i2\omega_{ad}t}+1)\right) \\
&-\frac{\Omega_S}{2}\left((e^{2i\omega_{sd}t}+1)(\left|\uparrow\right\rangle\left\langle\downarrow\right|+\left|\downarrow\right\rangle\left\langle\uparrow\right|(1+e^{-2i\omega_{sd}t})\right) \\
&+(n-n_0)(\omega_c-\omega_{cd})\left|n\right\rangle\left\langle n\right|+(\omega_s-\omega_{sd})\left|\uparrow\right\rangle\left\langle\uparrow\right|
\end{aligned}
\tag{6.48}
$$

Now we have clearly separated timescales, with slowly and quickly rotating terms. We can complete the RWA by neglecting the terms rotating at $2\omega_{sd}$ and $2\omega_{ad}$:

$$
\begin{aligned}
H_{RWA} &= \hbar\sum_n -\frac{\Omega_A}{2}\left(\left|n,\uparrow\right\rangle\left\langle n+1,\downarrow\right|+\left|n,\downarrow\right\rangle\left\langle n-1,\uparrow\right|\right) \\
&-\frac{\Omega_S}{2}\left(\left|\uparrow\right\rangle\left\langle\downarrow\right|+\left|\downarrow\right\rangle\left\langle\uparrow\right|+(n-n_0)(\omega_c-\omega_{cd})\left|n\right\rangle\left\langle n\right|+(\omega_s-\omega_{sd})\left|\uparrow\right\rangle\left\langle\uparrow\right|\right.
\end{aligned}
\tag{6.49}
$$

We can write this in terms of drive detunings, $\Delta_{sf}=\omega_s-\omega_{sd}$ and $\Delta_{cd}=\omega_c-\omega_{cd}$, and recombine the cyclotron and spin components, to give our final Hamiltonian:

$$
\begin{aligned}
H_{RWA} \approx &\sum_n -\frac{\hbar\Omega_A}{2}(\left|n,\uparrow\right\rangle\left\langle n-1,\downarrow\right|+\left|n,\downarrow\right\rangle\left\langle n+1,\uparrow\right|)-\frac{\hbar\Omega_S}{2}(\left|\uparrow\right\rangle\left\langle\downarrow\right|+\left|\downarrow\right\rangle\left\langle\uparrow\right|) \\
&+[\hbar(\Delta_{sf}+(n-n_0)\Delta_{cd})\left|n,\uparrow\right\rangle\left\langle n,\uparrow\right|+[\hbar(n-n_0)\Delta_{cd}\left|n,\downarrow\right\rangle\left\langle n,\downarrow\right|
\end{aligned}
\tag{6.50}
$$

Chapter 7

Cryogenic Single-Particle Detection

All single-particle detection in this experiment, whether through dips, driven signals, decays or self-excitation, is performed via measurements of the image currents induced in trap electrodes by the particle. These currents are small – a single proton, moving at 1 MHz with an axial energy of 4 K, induces a current of ∼2 femtoamperes. High-SNR current detection is thus essential. This chapter will focus on the RF electronics used to detect these currents: the low-loss tuned circuits which create a high resistance on resonance, and the cryogenic amplifier used to read out the signal. A significant focus of this work was improving the 1 MHz axial detection electronics, resulting in a factor of ~ 10 increase in effective resistance for the same coil size and frequency[1], as well as an improved understanding of losses and feedback.

[1]See table 7.3 for a breakdown of specific improvements. A factor of 9-10 improvement was measured at the operating parameters chosen for initial commissioning of the experiment, and limited by feedback through the FET amplifier. Without feedback, the improvement in R approaches a factor of 20. In section 7.2.8 we propose a method to achieve independent control over this feedback.

7.1 Detection Overview

The small induced current from a single particle must be transformed into a measurable voltage. We therefore send the current through a large resistance to induce a voltage. This voltage is applied to a cryogenic FET transimpedance amplifier. This increases the signal power, transmitting transforming it from a high to low impedance, so it can be sent through a transmission line cable to room temperature electronics. The induced voltage also acts back on the particle, providing the damping force described in section 2.3.

However, a large resistor would produce broadband Johnson noise, driving unwanted magnetron, cyclotron, and sideband transitions. We therefore maximize the impedance at the axial or cyclotron frequency using a resonant tuned circuit. Figure 7.1 shows how a large, low-loss inductor is connected in parallel to the capacitance between trap electrodes, creating an LC tank circuit with a large effective resistance at the selected frequency.

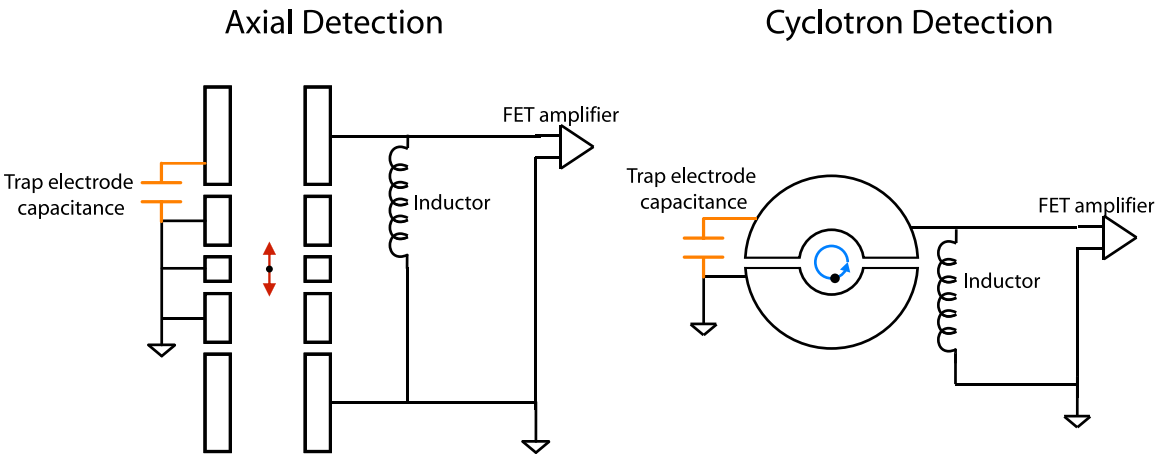

Figure 7.1: Schematic of the resonant detection system.

While the concept is similar for the cyclotron and axial motions, the frequencies differ by a factor of ∼100. The LC resonance condition $f = 1/(2\pi\sqrt{L * C})$ means a 10,000-fold smaller inductor is required to resonate with the trap capacitance at this higher frequency. The design principles, construction techniques, and limiting effects therefore differ significantly. Additionally, RF effects and magnetoresistances differ between 1 and 100 MHz. In the next section we will discuss the axial amplifier in detail, with a briefer discussion of the cyclotron amplifier following.

After the FET, the amplified signal is further impedance-matched to 50 Ω and sent either to a second cryogenic amplification stage or to the room-temperature detection chain. The room-temperature detection schematic remains unchanged from reference [35], although a new set of filters and amplifiers were constructed for the new apparatus.

7.2 The axial tuned circuit

Improving the construction and operation of the axial tuned circuit to increase its effective resistance was one focus of the work presented in this thesis. Higher effective resistance translates to improved signal-to-noise on measurements of the axial frequency, which improves the speed and accuracy of spin or cyclotron state measurements in the analysis trap. Up to a point[2], the increased effective resistance

[2]Higher effective resistance also increases the dip width. At a certain point, the wider dip introduces more uncertainty to the axial frequency measurement than the increased SNR resolves. Reference [100] found this to take place with Q values above 20,000. The relationship between dip width, frequency resolution, and axial stability has not yet been measured in our apparatus; the goal of the work in this thesis was to create resonators which could reach or exceed this limit. Later in this chapter we describe a method to controllably reduce the Q which would allow tuning of dip widths. The relationship between effective resistance and SEO stability and resolution also needs to be examined in the new apparatus.

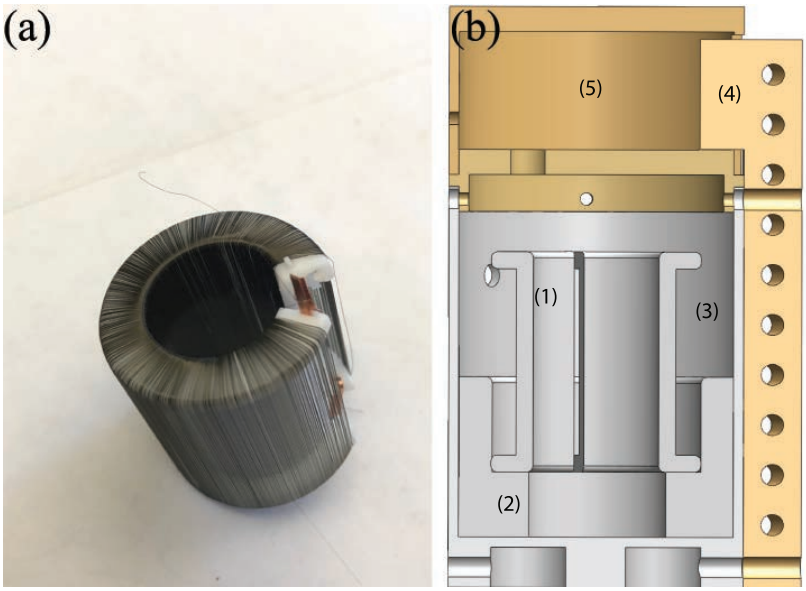

Figure 7.2: (a) Hand-wound NbTi axial coil, currently in use in the precision trap. (b) CAD model featuring (1) the PTFE coil form; (2) coil support, also PTFE; (3) NbTi superconducting shield can; (4) OFHC copper post for thermal and electrical anchoring; (5) copper shield for FET amplifier board.

also improves the axial frequency measurement precision and stability, as we can average down technical noise faster[3]. Additionally, improved axial signal-to-noise brings significant quality-of-life and experiment uptime improvements by making it easier to find and manipulate single-proton signals.

The work reported here can be conceptually separated into two parts. First, we systematically improved the design and construction of the tuned circuit to increase R. This included both improving the inductor coil and reducing loss in the trap capacitance due to structural materials. This work is reported in sections 7.2.1-7.2.3.

These improvements resulted in a resonator with low enough losses to reveal pre-

[3]Axial frequency stability in the analysis trap was found to depend on both short-timescale technical noise and long-timescale fluctuations due to cyclotron state changes [35], with an optimal averaging time that balanced these effects. Faster averaging of technical noise improves stability by allowing a shorter averaging time, meaning the long-timescale fluctuations have less effect.

viously unnoticed effects. Q and R were found to be limited by feedback through the FET amplifier, with a strong dependence on the FET operating parameters. In sections 7.2.4-7.2.6 we develop a full model of the resonator Q and R, starting with the resonator-FET coupling and then accounting for feedback through the FET.

7.2.1 Superconducting toroid design principles

The small size of a single-proton image current necessitates a large resistance R, given by

$$R = Q\omega L. \tag{7.1}$$

Here Q is the quality factor of the amplifier, determined by losses in the tuned circuit, ω is the axial frequency in radians, and L is the inductance of the superconducting coil. This resistance, together with the trap geometry, determines the damping width

$$\gamma = \left(\frac{e\kappa}{2z_0}\right)^2 \frac{R}{m} \tag{7.2}$$

This leads us to three main design principles for the axial circuit, which combine to maximize the effective resistance R. We want to maximize the inductance while minimizing both stray capacitance and RF loss. The inductance is determined by the number of turns, dimensions, and relative geometry of the coil winding, while loss and capacitance each depend both on the geometry and the construction materials for all relevant parts of the circuit.

Each axial coil is a hand-wound toroid, consisting of .0028" or .0030" diameter, PTFE- or formvar-coated NbTi superconducting wire wound over a PTFE form, inside a cylindrical NbTi can. Previous work [34] empirically found the toroidal

geometry superior to a solenoidal geometry. A toroid largely contains the RF fields inside the windings, allowing the can to act solely as a shield rather than part of the circuit, which reduces lossy couplings to other parts of the experiment. Additionally, in this chapter we show that magnetoresistance in the superconducting wire is a significant source of RF loss. The rectangular toroid, with its small aspect ratio (visible in figure 7.2), is less vulnerable to magnetoresistance. When aligned vertically, the majority of electron travel in the windings proceeds along the field axis, while in a solenoid all electrons travel perpendicular to the field.

The available space in the 4" magnet bore determines the size of the shield and coil. In the new apparatus, we need a total of three axial and two cyclotron coils, so the coil size was chosen such that two could sit adjacent to each other at the same height in the 2.6"-diameter tripod. This restricted the coil radius to about half that used in the previous generation proton experiment, but by improving construction techniques using low-loss materials we were able to get an equal or better effective resistance. Figure 7.3 shows the relevant dimensions and geometry for the axial amplifier.

Once the shield dimensions CanH and CanD have been determined, the optimal coil dimensions ID, OD, and CoilH can be calculated by maximizing the ratio L/c, where c is the total circuit capacitance to ground and L is the toroid inductance. L for a rectangular toroid is given by

$$L = \mu_0 \left(\frac{\pi * ID}{D_{wire}} \right)^2 \frac{CoilH}{2\pi} Log \left[\frac{OD}{ID} \right] \tag{7.3}$$

where D_{wire} is the wire diameter and other dimensions are defined in figure 7.3.

The capacitance to ground is the sum of the trap capacitance, parasitic capacitance on the signal path (largely in the vacuum feedthrough), the capacitance between the

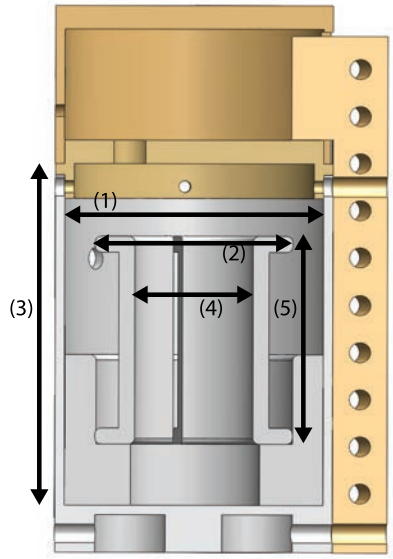

Figure 7.3: Cross-section of the axial amplifier with relevant dimensions for calculating inductance and capacitance. (1) Can diameter CanD. (2) Coil outer diameter OD. (3) Can height CanH. (4) Coil inner diameter ID. (5) Coil height CoilH.

coil and shield of the toroid, and the self-capacitance of the toroid. The shield-to-can capacitance can be approximated by treating the toroid-shield system as a set of three parallel-plate capacitors, with area equal to the surface area of the coil and coil-to-can separation as a parameter in the optimization.[4] This gives the equation

$$C_{Coil-Can} \approx \epsilon_0 \left(\frac{\pi * OD * CoilH}{\frac{1}{2}(CanD - OD)} + 2\frac{\frac{\pi}{4}(OD^2 - ID^2)}{\frac{1}{2}(CanH - CoilH)} \right) \tag{7.4}$$

For the self-capacitance, we approximate the toroid as though it were a solenoid with the same number of turns, length equal to the mean of the inner and outer toroid circumferences $L_{eq} = \frac{\pi}{2}(ID + OD)$, and diameter approximated by $D_{eq} = \frac{2}{\pi}(OD - ID + CoilH)$. We can then make use of an empirical formula for solenoid self-capacitance: [101]

$$C_{self} = 4\frac{\epsilon_0}{\pi} L_{eq} \frac{1 + .717\frac{D_{eq}}{L_{eq}} + .933\left(\frac{D_{eq}}{L_{eq}}\right)^{3/2} + .106\left(\frac{D_{eq}}{L_{eq}}\right)^2}{Cos\left[ArcTan\left[\frac{L_{eq}/N_{turns}}{\pi D_{eq}}\right]\right]^2} \tag{7.5}$$

With these parameters, as well as the predicted trap capacitance (calculated from electrode geometry) and the feedthrough capacitance (measured to be approximately 3 pF), we can numerically optimize the coil dimensions. Dimensions for the inductors used in this work are given in table 7.1, along with their electrical parameters.

7.2.2 Measuring and comparing amplifiers

The effective resistance R is the most important parameter for a resonator, because it sets the particle damping width (eqn. 7.2) as well as the signal at the input to the

[4] This underestimates the capacitance by ignoring the larger size of the can, but also overestimates the capacitance by ignoring the gaps between wires on the outside of the coil. In testing, the approximation proved close enough to allow final tuning with only minor adjustments.

Table 7.1: Dimensions and properties of axial amplifiers used in this work.

Amplifier	Precision Trap	Analysis Trap
ID	.62"	.62"
OD	1.05"	1.05"
CoilH	1.1"	1.1"
CanD	1.375"	1.375"
CanH	1.81"	1.81"
WireD	.0026"	.0030"
Nturns	760	660
L	1.55 mH	1.2 mH
Frequency	919 kHz	1075 kHz
Q (FET off)	24,000	16,000
Q (operating)[5]	12,000 or 16,000	7,000
R_{eff}	90 or 120 MΩ	56 MΩ

[2]See section 7.2.6 - 12,000 was used for initial searches, 16,000 for higher precision.

FET. However, for direct comparison of circuits and characterization of improvements, the quality factor Q is often used as a measure of loss in the system. Q is directly measurable, whereas R depends on an estimate of the total coil inductance and trap capacitance. Q can be measured by taking a Fourier transform of the Johnson noise in the resonator, amplified by the FET. The Johnson noise in a bandwidth Δ is

$$V = \sqrt{4k_b T R(\Delta)} \tag{7.6}$$

The Johnson noise in a resonator can then be fit to a Lorentzian, with the Q given by the width and center frequency[6]

$$V \propto \frac{\left(\frac{\gamma}{2}\right)^2}{(\nu - \nu_0)^2 + \left(\frac{\gamma}{2}\right)^2}, \qquad Q = \frac{\nu_0}{\gamma} \tag{7.7}$$

A noise resonance fit of our detectors is shown in figure 7.17 below. In addition to the Johnson noise, the amplifier Q can also be measured by applying an RF signal to a neighboring electrode and measuring the strength of pickup in the resonator as a function of drive frequency. The Johnson noise resonance is often preferable since it more closely mirrors the signal seen under experimental operating conditions; however, using a swept drive is faster and less sensitive to RF noise pickup, and so can be preferable for characterization. Driven resonances are shown in figure 7.11.

[6]Poor tuning of the impedance-matching circuit can introduce a reactance, which phase-shifts the signal coming out of the FET. This can result in a partially dispersive Lorentzian shape to the noise resonance. In this case, a more accurate Q can be extracted from a fit to the expression $V \propto \frac{\left(\frac{\gamma}{2}\right)^2 \cos(\phi) + \frac{\gamma}{2}(\nu - \nu_0)\sin(\phi)}{(\nu-\nu_0)^2 + \left(\frac{\gamma}{2}\right)^2}$. For operation with particles, we used component values and operating parameters that gave $\phi \approx 0$.

7.2.3 Improvements to axial coil construction and materials

In addition to systematizing coil design, several other improvements were made during this work which resulted in higher-quality coils. These included careful choice of low-loss structural materials in the trap and resonator; use of thinner wire and improved coil-winding techniques; and adjusting the coil's position in the magnetic field to reduce magnetoresistance.

The choice of trap and resonator construction materials has a significant effect on performance. The previous generation of the experiment used Macor spacers between trap electrodes (forming the trap capacitance dielectric), which we replaced with ground quartz. The motivations for this in terms of magnetic field homogeneity and trap design are discussed in section 8.1.1, but it also significantly reduced dielectric losses – on a test jig with actual trap electrodes in the magnetic field, replacing the Macor spacers with quartz was found to improve the quality factor by a factor of 2-3.

Structural materials inside the superconducting shield also have a large effect on Q. As noted above, the coil form is Teflon, chosen because of its low dielectric loss. The Teflon is machined so as to minimize the amount of material inside the toroid. After machining, the surface of the Teflon is sanded to remove tool inclusions, then scrubbed to remove sanding grit, before being cleaned in solvents. The previous apparatus used an alumina rod at the center of the toroid as heat sink. This was replaced with a Teflon coil support, noted on figure 7.3. While Teflon has lower thermal conductivity than alumina, it may also have lower dielectric loss.[7] Additionally, the Teflon support

[7]Measured dielectric loss tangent values for alumina are highly sample- and purity-dependant, ranging from half to quadruple the value for PTFE [102, 103]. The BASE collaboration found alumina-insulated feedthroughs to limit the Q of their superconducting resonator [104].

contacts the outside of the coil, reducing the effect of dielectric loss compared to a central rod. The increased contact surface area of the Teflon support compensated for the lower heat conduction, and over several months of operation we had no trouble keeping the wires thermally anchored.

The insulating materials for the wire and feedthrough were also replaced. The formvar insulation on the NbTi wire was replaced with Teflon, while the alumina in vacuum feedthroughs which form part of the RF circuits was replaced with sapphire. These materials were reported in [104] as a major limit to the quality of resonators on a similar Penning trap experiment. At this point we began to reach the limits on Q imposed by feedback through the FET, discussed later in this chapter. While an improvement of only 20-30% was measured when the formvar and alumina were replaced, the real improvement was likely greater.

Another improvement in the new apparatus was the use of thinner NbTi wire, either .0028" or .0030" instead of .0050" in diameter. This required improved coil-winding techniques[8], since working with this hair-thin wire is significantly more difficult – care must be taken to keep the wire taut enough to make dense, parallel windings without breaking it.[9] Comparative tests of coils with identical geometry and different wire thickness (figure 7.4) found no increase in loss in the narrower wire.

[8]In essence: the end of the NbTi wire was sanded bare and spot-welded to a copper strap, which was attached to the smooth, clean PTFE form. A length of wire was turned off its spool, which was then fixed. Each turn was wound through a slit in the form without twisting the wire. The smooth PTFE allowed the turns to be slid into place by hand or with tweezers. Once the available wire was wound onto the form, the spool was unfixed and more wire turned off of it. Throughout the process, the NbTi was kept under consistent tension to preserve winding geometry without breaking the wire. Once the coil was finished, the end of the wire was wrapped around a second copper strap to hold tension; it was then clipped, sanded, and spot-welded to the strap. The entire coil was sonicated, allowed to dry overnight, and inserted into the PTFE holder and superconducting shield.

[9]While this was not systematically studied, a coil with a patched broken wire was found to have a lower Q than similar coils, so we avoided breaking the superconducting wire during winding.

The larger number of turns thus increased the coil inductance and resonator R.

After the coil geometry and materials, there are two remaining contributions to losses in the resonator. Parasitic couplings may have an effect, even at the relatively low axial frequency of 1 MHz. A small amount of RF power is radiated from the copper straps connecting the trap to the coil. When other conductors are close enough to inductively or capacitively couple to the straps, they can provide a lossy parallel path to ground. To optimize the performance of the circuit, these straps must be separated from other conductors (as much as possible in the tightly packed tripod and trapcan).

The final major contribution to loss in the tuned circuit comes from magnetoresistance. In the absence of a magnetic field, the superconducting coil contributes nearly zero loss to the circuit. However, as the magnetic field increases, Q decreases. Measurements taken with a test jig in the bore of the solenoid (Fig. 7.4) demonstrate this – the resonator is lowered or raised in the bore, changing the magnetic field it experiences. These measurements confirm that the field-dependent loss is occurring in the coil, rather than the electrodes or signal lines – the effect is strongest when the coil is on field center and the test trap is past the high-field region.

As mentioned above, the toroidal geometry should mitigate the effect of magnetoresistance. Additionally, the coil is placed higher in the tripod. The tripod extends through the edge of the solenoid, where the field strength falls off quickly, so a shift of a few inches can mean a significant difference in magnetic field (see Fig. 7.5).

Measurements with a test circuit allow us to determine the losses in different parts of the system. A test coil was attached first to a lump-element capacitor; then a pair of feedthrough pins was added; then a trap and pinbase with Macor spacers

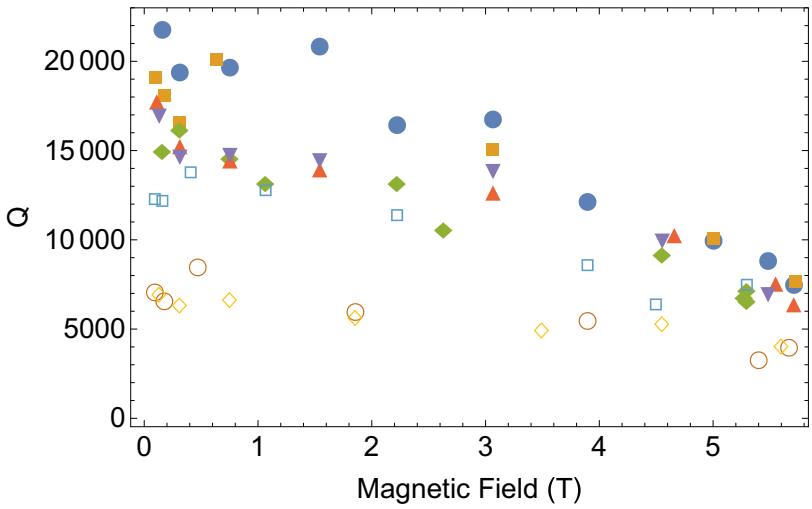

Figure 7.4: Measured Q values vs magnetic field strength at coil position, for a variety of different test resonators, demonstrating magnetoresistance.

was connected instead of the pins and capacitor; and finally the Macor spacers were replaced with ground quartz. Results for this test are presented in table 7.2.[10]

Table 7.2: Q values and calculated effective loss from different components of the test circuit. Effective loss due to feedback depends on FET parameters (see 7.2.6) but is held constant through this test.

Component	Q when attached	Effective loss
Coil with FET and test capacitor	20,500	0.35 Ω
Alumina-insulated feedthrough pins	19,000	0.07 Ω
Test trap stack with Macor spacers	7,000	0.74 Ω
Test trap stack with quartz spacers	14,000	0.19 Ω

From the measured quality factor in each test, we can calculate an overall loss and

[10]Note that in section 7.2.6 we calculate the series loss in the entire resonator on the experiment, at equal FET operating parameters, as ∼0.4 Ω, lower than the total of 0.61 Ω found for this test resonator; this demonstrates the importance of the individual coil on the resonator behavior.

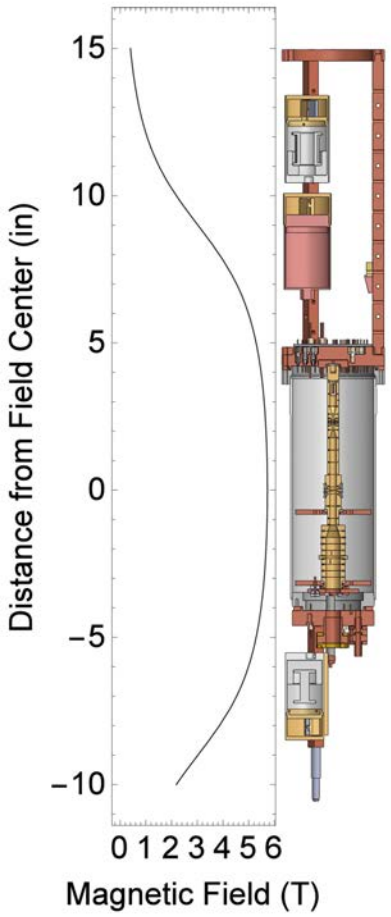

Figure 7.5: Tripod and trap can with amplifiers, compared with the approximate magnetic field from the superconducting solenoid vs distance from field center.

determine the loss in the changed component. This is calculated by treating the FET and circuitboard, at identical operating parameters, as a static loss R_b and converting the resulting circuit diagram (figure 7.6) to an effective parallel RLC circuit.

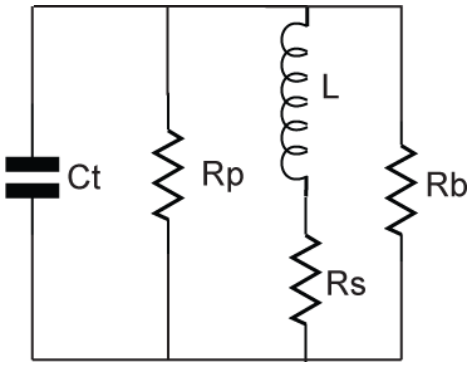

Figure 7.6: Resonator equivalent circuit, treating the FET and circuitboard as a constant parallel loss source R_b; L is the superconducting coil, C_t is the trap capacitance, R_s represents series losses in the coil and trap capacitance, and R_p represents generic parallel losses i.e. due to parasitic couplings.

The equivalent inductance in the effective parallel RLC circuit is modified by the series resistance to be $L * (1 + \frac{R_s^2}{\omega^2 L^2})$, but for reasonable component values we find $\frac{R_s^2}{\omega^2 L^2} \ll 1$. The series loss R_s converts to an equivalent parallel resistance $\frac{\omega^2 L^2}{R_s}$. The Q for a parallel RLC circuit is in general given by $Q = \frac{R}{\omega L}$, where ω is the resonant frequency. For the circuit in figure 7.6, we can write [11]

$$Q = \frac{1}{\omega L \left(\frac{1}{R_b} + \frac{1}{R_p} + \frac{R_s}{\omega^2 L^2} \right)} \tag{7.8}$$

In table 7.2 the effective parallel resistance R_p is treated as infinite; we include it to illustrate the potential effects of lossy contamination (ie, fingerprints) and radiative

[11] Note the third term, $\frac{R_s}{\omega^2 L^2}$, representing series loss in the resonator – as long as the losses in the circuit are dominated by this series term, the Q increases with higher inductance. However, the parallel loss R_p and the circuitboard and FET loss term R_b do not have the same scaling. Typically, our detector circuits are dominated by series loss, which is why Q can be approximated as $\frac{\omega L}{R}$ for our resonator (whereas typically Q for a parallel RLC circuit would be $\frac{R}{\omega L}$).

couplings, but with a carefully constructed circuit we find it to be negligible compared to R_b and R_s. We measure a DC resistance >10 GΩ and no radiative effect at 1 MHz.

Table 7.3 summarizes the design and construction changes made for the new amplifiers, and their approximate effect on resonator parameters compared to the previous amplifier with the same size and frequency based on tests described in this section. (Approximate R values are derived from proton dip widths and equation 2.21. Values are typical of the commissioning work reported in this thesis, where we chose FET operating parameters giving Q∼12,000. The actual Q increase due to specific upgrades could be larger, as our Q became limited by feedback.)

Table 7.3: Summary of axial resonator design and structural improvements.

Improvement	Effect	R (cumulative)
Increase coil windings	∼ 70% increase in L	factor of 1.7
Improve coil geometry	∼ 60% increase in Q	factor of 2.7
Replace macor spacers	∼ ×2 increase in Q	factor of 5.4
Relocate in magnetic field	∼ 30% increase in Q	factor of ∼7
Replace alumina & Formvar	∼ 30% increase in Q	factor of ∼9
Total	Q ∼ 2200 → 12000	R ∼10→90 MΩ

7.2.4 The cryogenic FET and circuitboard

At the high-impedance point between the inductor and trap capacitance, the induced current creates a voltage of order 100 nV. However, if we were to couple that

directly into a 50-ohm detector, it would load down the circuit to the point where the resonance would be non-existent. Therefore, we use a cryogenic field effect transistor as a transimpedance amplifier, to couple the signal from the high-impedance resonator out to room temperature. This FET is soldered to a heat-sink post and low-loss, homemade Teflon circuitboard, and capacitively coupled to the high-impedance point of the resonator. Figure 7.7 shows the full circuit for the precision axial amplifier; components to the right of the 0.5 pF capacitor are located on the circuitboard.

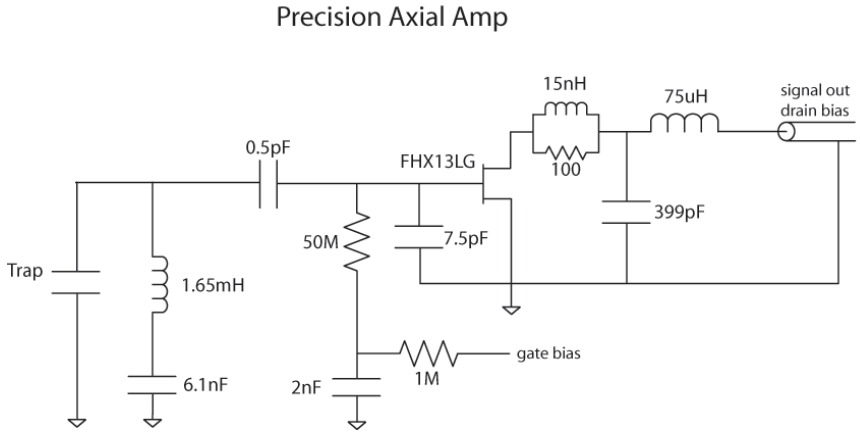

Figure 7.7: Circuit diagram for the precision axial amplifier, including trap and coil; tap ratio capacitors (0.5 and 7.5 pF); gate bias filter; high-frequency suppression circuit (100 Ω and 15 nH in parallel); and output matching network (399 pF and 75 uH).

The circuitboard has three essential functions: providing DC biases to the FET drain and gate; coupling the signal into the FET while minimally loading the resonator; and matching the FET output impedance to the 50-ohm cable impedance. Since the gate bias is applied to the same point as the signal, it is filtered to avoid noise, both with a standard filterboard (sec. 8.4) and additional filtering on the FET board. The drain bias is less sensitive, since noise would add to the amplified post-

FET signal, but it is still heavily filtered at room temperature.

The proton's signal is coupled to the FET input with a capacitive divider. The high-impedance arm of this divider is attached between the FET gate and the resonator, while the low-impedance arm goes from the gate to ground. This divider reduces the coupling between the FET and the resonator. This reduces the loading from the FET on the resonator, increasing Q and R, at the cost of reduced signal at the gate. We can optimize these competing effects by adjusting the capacitive divider or "tap ratio" (so called in analogy to higher-frequency resonators, where this effect is achieved by tapping partway along the coil).

In order to optimize the tap ratio, we need an estimate of losses in the circuitboard and FET. These were measured for a newly constructed coil by repeatedly cooling it to 4K with different tap ratios and identical FET parameters, as shown in figure 7.9. The curve of ratio vs Q was fit to obtain estimates for the loss before and after the tap capacitors. The approximate measured values at ∼1 Tesla were ∼1 MΩ parallel resistance to ground after the divider[12] and 0.35 Ω series loss in the inductor.

The effective circuit including tap ratio is shown in figure 7.8. The data in figure 7.9 is fit to the expression for total Q we get when modifying equation 7.8 to account for the effect of the tap ratio[13]:

$$Q = \frac{1/\omega L}{\left(\frac{1}{\left(\frac{C_1+C_2}{C_1}\right)^2 R_b + \frac{1}{\omega^2 C_1^2 R_b}} + \frac{1}{R_p} + \frac{R_s}{\omega^2 L^2}\right)} \approx \frac{1/\omega L}{\left(\left(\frac{C_1}{C_1+C_2}\right)^2 \frac{1}{R_b} + \frac{1}{R_p} + \frac{R_s}{\omega^2 L^2}\right)} \tag{7.9}$$

[12] In section 7.2.6 we show that this 1 MΩ loss comes from negative feedback, and depends on the FET gain and matching network tuning. Those parameters were held constant for these tests.

[13] The term $1/\omega^2 C_1^2 R_b^2$ arises from transforming the series C_1 and R_b into parallel impedances. With $R_b \geq 1$ MΩ and $C_1 \approx 1$ pF, this term is much smaller than the R_b term and can be neglected.

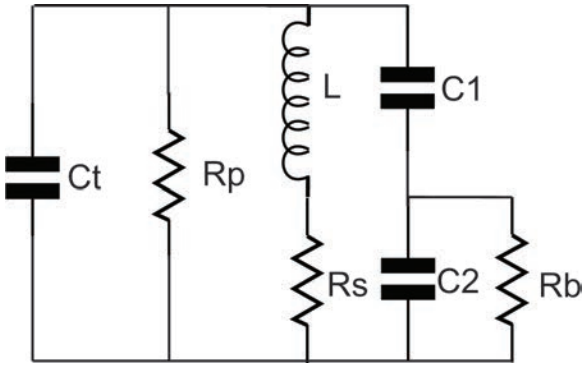

Figure 7.8: Resonator equivalent circuit, updated to include the tap ratio capacitors C_1 and C_2

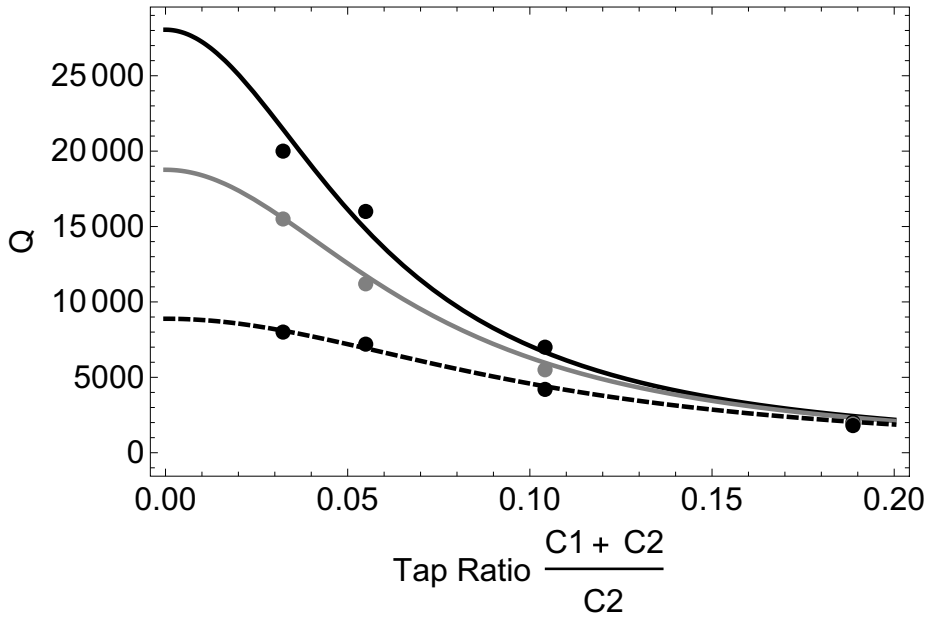

Figure 7.9: Measured Q values using a test resonator at different tap ratios, measured in low (∼1 T, dashed line), medium (∼2.5 T, gray line) and high (>5 T, black line) magnetic field, simultaneously fit to equation 7.9.

We optimized the tap ratio to maximize the signal-to-noise of the particle's axial motion. While this is the natural choice for detection (with benefits for spin and cyclotron state analysis), arguments can be made for other optimizations – for example, γ_z or the axial temperature could be optimized at a different ratio (see section 7.2.8).

The axial SNR is determined by the effective resistance of the resonator and the gain of the FET. The voltage across the resonator is given by equations 2.19 and 7.1. Incorporating equation 7.9, we obtain

$$V_{resonator} = I_{image} R_{effective} = \frac{q\kappa\dot{z}}{2z_0} \frac{1}{\left(\left(\frac{C_1}{C_1+C_2}\right)^2 \frac{1}{R_b} + \frac{1}{R_p} + \frac{R_s}{\omega^2 L^2}\right)} \tag{7.10}$$

The amplification from the FET is the product of g_{FET} and the tap ratio, which we label $T = \frac{C_1}{C_1+C_2}$. We can express the amplified signal from a thermal particle:

$$V_{signal} = g_{FET} * T * \frac{q\kappa\dot{z}}{2z_0} \frac{1}{\left(T^2 \frac{1}{R_b} + \frac{1}{R_p} + \frac{R_s}{\omega^2 L^2}\right)} \tag{7.11}$$

We maximize this by setting the derivative to zero and solving for T:

$$T_{optimal} = \sqrt{\frac{R_b R_p R_s + R_b L^2 \omega^2}{R_p L^2 \omega^2}} \approx \frac{\sqrt{R_b R_s}}{L\omega} \tag{7.12}$$

For the parameters in table 7.2, this gives $T_{optimal} \approx .075$. For the improved axial inductor used on the experiment, with estimated loss $R_s \approx 0.4\Omega$, this gives instead $T_{optimal} = .061$. We therefore used tap ratio capacitors with values $C_1 = 0.5$ pF and $C_2 = 7.5$ pF, for a tap ratio of $\frac{C_1}{C_1+C_2} = .063$.

Note that this result assumes a static R_b. We will show in the next two sections that R_b depends on the FET gain and impedance matching. Therefore, the implemented tap ratio could potentially be re-optimized using this procedure at different

FET operating parameters. To do this in situ, a varactor could be added in parallel with C_2 (as in [35]). This would require careful evaluation of loss in the varactor.

7.2.5 Feedback and stability

Feedback through the FET implies an effect on the input from the output. The effects of feedback on FET stability are described in detail in [44]. We note the result that the FET input impedance at the gate, Z_{in}, depends on the impedance seen at the drain and the FET transconductance as follows:

$$Z_{in} = \frac{Z_D + Z_m}{1 + g_m Z_D} \tag{7.13}$$

where Z_D is the impedance seen by the drain of the FET (determined by the matching network and the 50 Ω load), Z_m is the internal drain-to-source impedance (often referred to as the Miller capacitance), and g_m is the FET transconductance. Note that with g_m=0 (the FET "off"), Z_{in} is simply the series sum of Z_m and Z_D.

When the FET drain sees an inductive load and $g_m > 0$, the real part of Z_{in} is negative. This corresponds to positive feedback, which artificially increases the resonator Q factor, leading to unpredictable behavior and increasing the effective particle temperature [35]; or it can cause saturated oscillation at a single frequency, rendering the amplifier useless. When the drain load is capacitive, the real part of Z_{in} is positive, corresponding to negative feedback and reducing the Q of the resonator.

The drain is loaded by the impedance-matching network which couples the FET output to a 50-ohm output coax[14]. This is achieved using either an L- or pi-network

[14]The suppression circuit of figure 7.7 also contributes to the drain load, but its impedance at 1MHz is equivalent to a shift of ∼0.5% of the matching network inductor, less than variation in component values and cooldown effects, so we neglect it in this analysis for simplicity.

of matching components, as indicated in the circuit diagram of figure 7.7. The first capacitor and the inductor determine both the frequency at which the impedances are matched, and the value of the drain impedance best matched to 50 ohms. (Since output impedance depends on gain, this determines an optimally matched gain value.) The "second leg" capacitor of the pi-network is often omitted, leaving an L network. When precise tuning is desired, this capacitor gives control over the optimally matched gain value with less effect on the frequency; but similar results can be achieved by changing the values of the first two components while keeping their product constant. In figure 7.7, this capacitor would connect after the 75 μH inductor and go to ground.

To avoid feedback, we would in theory keep the drain impedance purely resistive at the amplifier frequency. However, on cooling to 4K, shifting component values can add a reactive component to a resistive load. To avoid an inductive load causing regeneration or oscillation, we intentionally make the drain load slightly capacitive, by mistuning the matching network slightly lower in frequency than the resonator.

While it prevents positive feedback, this mistuning does introduce some negative feedback. This is preferable to positive feedback, because it encourages stability, but it does affect the resonator – especially for high-Q amplifiers.

7.2.6 Feedback and transconductance-dependent amplifier Q

We can take equations 7.13 and 7.9 together for a full calculation of the Q factor. This is useful for understanding limitations on detection and identifying avenues for improvement. Figure 7.10 shows the complete circuit diagram, incorporating a re-defined R_b representing the 50 MΩ gate bias resistance (and parasitic losses through

the circuitboard material), as well as effective parallel FET input capacitance and resistance $C_{in,p}$ and $R_{in,p}$ (which sum to Z_{in} as given in equation 7.13).

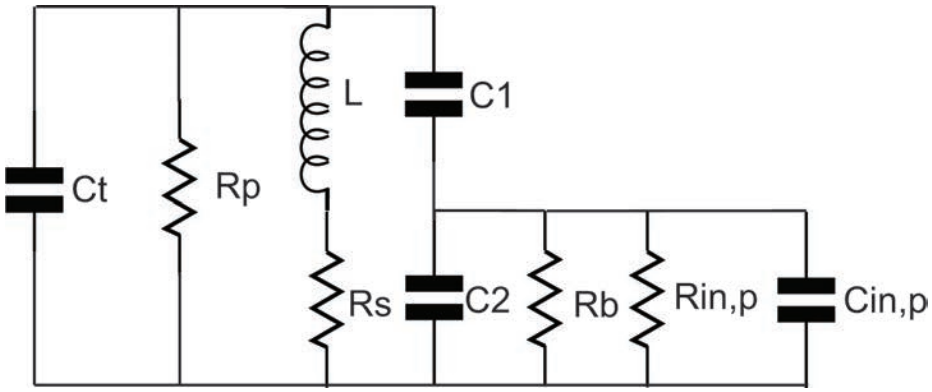

Figure 7.10: Resonator equivalent circuit, updated to include the transconductance-dependent FET input impedance.

Using series and parallel circuit conversion rules, we transform the resonator circuit to get an equivalent parallel RLC circuit. Modifying equation 7.8 we obtain

$$Q \approx \frac{1}{\omega L} \frac{1}{\frac{1}{R_{board+FET}} + \frac{1}{R_p} + \frac{R_s}{4\pi^2 f^2 L^2}} \tag{7.14}$$

$$R_{board+FET} = \left(\frac{1/R_{in,p} + 1/R_b}{\omega^2 C_1^2} + \frac{(C_1 + C_2 + C_{in,p})^2}{C_1^2} \frac{1}{1/R_{in,p} + 1/R_b} \right) \tag{7.15}$$

As in equation 7.9, the first term arises from converting the series C_1 to a parallel impedance. However, we cannot necessarily neglect it, because we want to allow for arbitrary $R_{in,p}$. To investigate of the effect of feedback on the amplifier front-end resonance, we can insert realistic parameters and see how the system changes when we vary them. We will use the component values in table 7.4.

The parasitic series and parallel losses R_s and R_p are necessary to complete this calculation. Because we can measure a DC leakage resistance of several hundred GΩ, we approximate $R_p >> 1/R_s$. R_s itself depends on losses in the trap and resonator

Table 7.4: Component values used to calculate the effects of transconductance on the front-end resonance.

	L	C_t	C_1	C_2	f	R_b	C_{pi}	L_{pi}
Value	1.65 mH	17.5 pF	0.5 pF	7.5 pF	.924 MHz	20 MΩ	470 pF	62.3 uH

materials, and can only be measured on the experiment while cold. We can do this by measuring the Q of the resonator with the transconductance set to zero, eliminating the effect of negative feedback from the resonator. When we apply a strong drive to the trap and sweep its frequency, it resonates between the inductor coil and the trap electrode capacitance, traveling through the copper straps connecting them and emitting radiation. This radiation induces currents in the parallel straps connecting the analysis trap inductor and electrode capacitance. By measuring at the drive frequency on the analysis trap FET, we can thus measure the current in the precision trap resonator with its FET turned off.[15] This measurement, shown in figure 7.11, yields a "FET-off" or "feedback-free" Q $\approx$ 24,000 for the precision axial amp, and $\approx$ 16000 for the analysis axial amp. We can then calculate the series loss in the resonator; for the precision trap, we derive $R_s \approx 0.4$ Ohms.

We can now model the entire detection circuit to see how feedback affects the front end resonator. We observed the effect of feedback on the new, low-loss resonators when we measured a dependence of Q on the FET operating current. In figure 7.12 we show the measured and calculated dependence of amplifier Q on transconductance during one cooldown. (The small-signal transconductance is measured at each point

[15]The analysis trap resonator impedance could distort this measurement, but the resonances are separated in frequency by nearly a hundred times their bandwidth, so this effect is negligible.

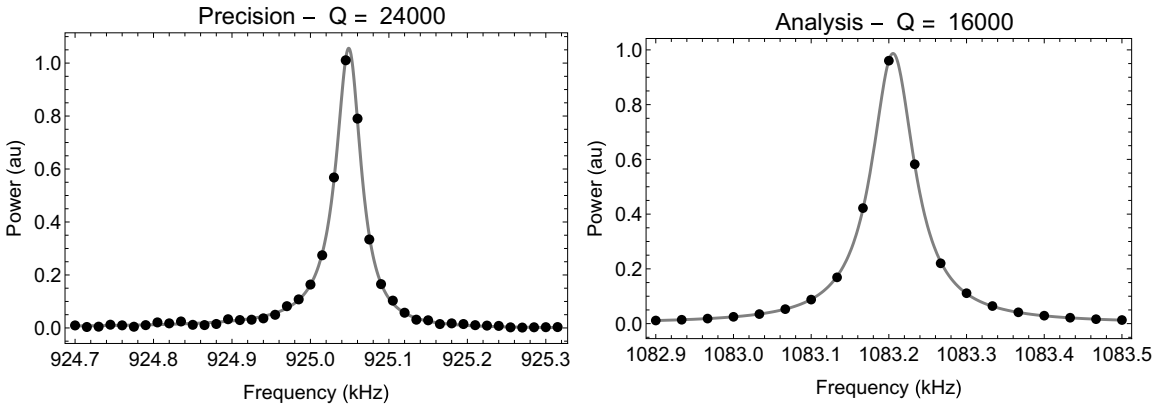

Figure 7.11: Driven resonances of the precision and analysis axial amplifiers, with FETs off (current and transconductance set to zero).

by changing the bias voltage by a small amount δV and measuring the resultant change in current, yielding $g_m = \frac{\delta I}{\delta V}$.) Figure 7.13 calculates this behavior for systems with larger losses, in which the smaller effect could have passed unnoticed.

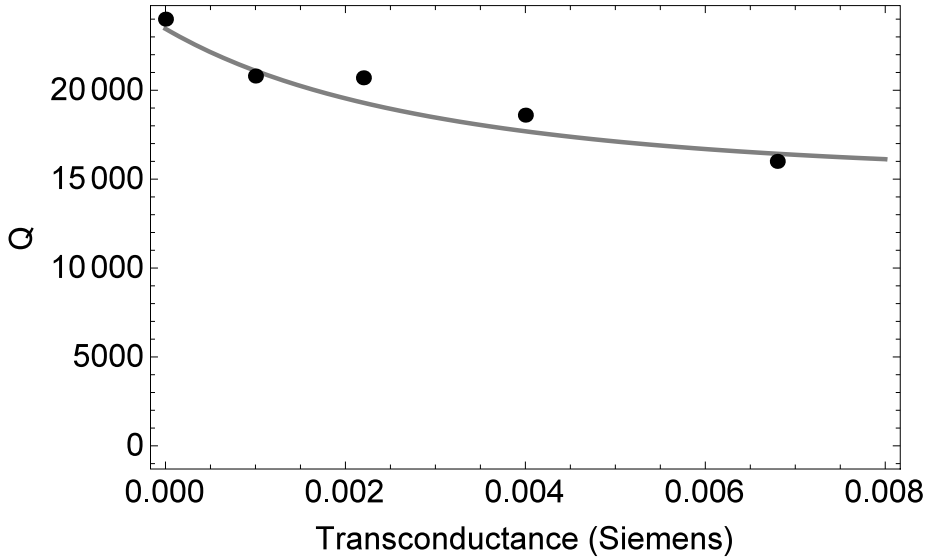

Figure 7.12: Points: Measured Q values at different FET operating points. Line: Calculated Q from equations 7.14 and 7.15, using parameters in table 7.4

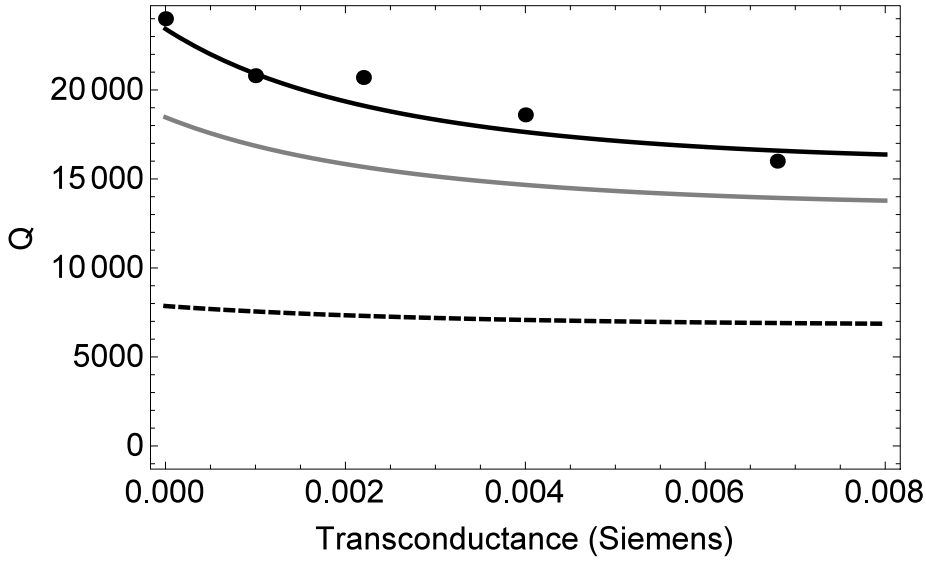

Figure 7.13: Points: Measured Q values at different FET operating points.
Black Line: Calculated Q for parameters in table 7.4 and 0.4Ω series loss
Gray Line: Calculated Q value with 0.5 Ω series loss
Dashed line: Calculated Q value with 1.2 Ω series loss

The strength of this transconductance dependence also depends on the tuning of the matching network. The further the matching network frequency is detuned from the front-end resonance, the more capacitance is seen by the FET output, and thus the more resistive loss is present at the FET input. Figure 7.14 shows the change in Q versus transconductance for different matching network capacitances, calculated with a 62.3 uH inductance and for an amp front-end at .924 megahertz. This implies a method to control this effect – adding a varactor in parallel with the matching network capacitor would allow us to tune the strength of feedback, independent of transconductance. This is discussed in section 7.2.8.

We have thus far assumed that the value of the Miller capacitance is relatively constant with respect to transconductance. This assumption is motivated by the

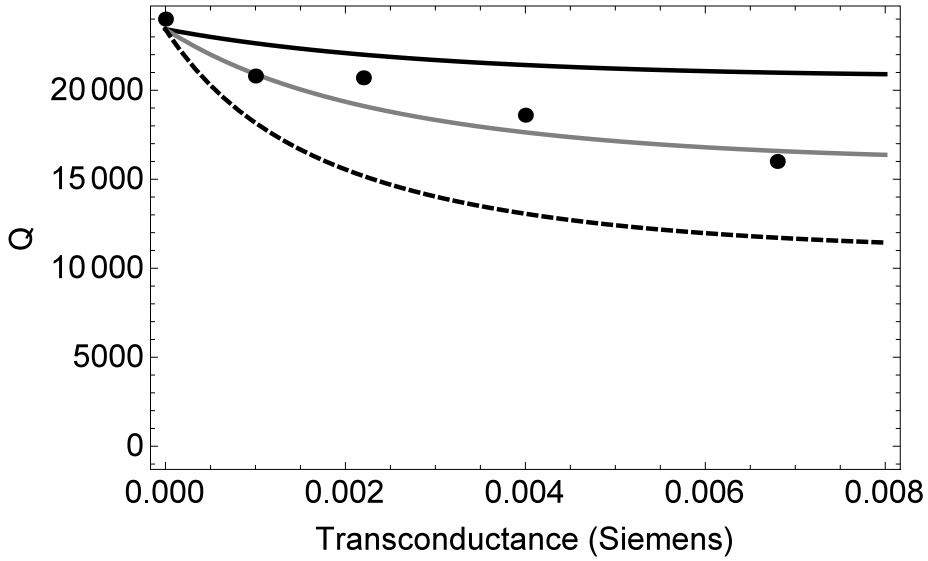

Figure 7.14: Points: Measured Q values at different FET operating points. Black Line: Calculated Q with a 468 pF pi-net capacitor, detuning 0.7 kHz Gray Line: Calculated Q with parameters in table 7.4, detuning 2.7 kHz Dashed line: Calculated Q with a 474 pF pi-net capacitor, detuning 6.7 kHz

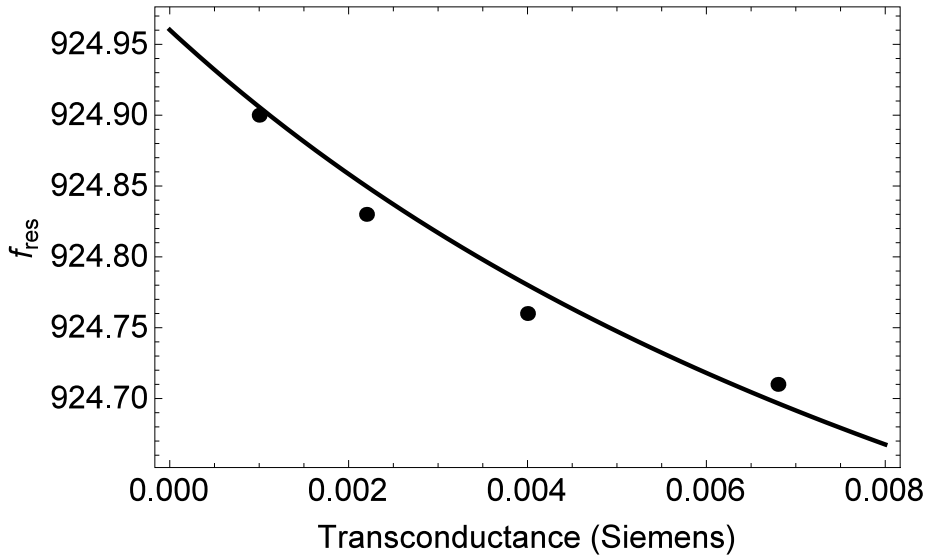

Figure 7.15: Resonator frequency as a function of transconductance. Calculation is for parameters in table 7.4 using 0.25 pF Miller capacitance.

regime in which we are operating the FET – it is designed to operate in the tens of gigahertz with 10 mA of current, so at 1 MHz and 45 uA we are starving it and operating it near DC. We therefore expect changes in its performance to be small with respect to small changes in operating conditions. However, we can check this assumption by looking at the effect of feedback on the resonator's frequency as well as its Q. Using equation 7.13 for the input impedance of the FET, we can calculate a transconductance-dependent capacitance which adds in parallel with the 7.5 pF capacitor in the capacitive divider. Figure 7.15 shows that a Miller capacitance of 0.25 pF gives a good simultaneous match to the measured transconductance-dependent Q and frequency data. Additionally, figure 7.16 shows that the Q calcualtion is relatively insensitive to the exact value of the Miller capacitance.

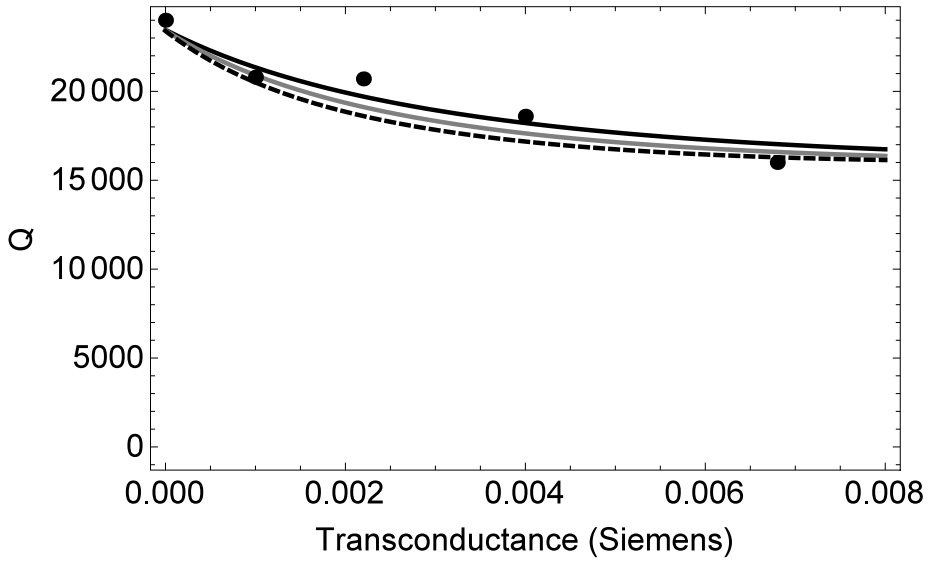

Figure 7.16: Q factor vs transconductance, data and calculations for different values of Miller capacitance:
Black line: c_m=0.2 pF; Gray line: c_m=0.25 pF; Dashed line: c_m=0.3 pF.

7.2.7 Experimentally Realized Axial Amplifiers

Three superconducting axial amplifiers were constructed and tested for use on the apparatus. One is attached to the precision trap and one to the analysis trap; the third will be attached to the cooling and/or loading traps when those traps are in use. The precision axial amp has the highest quality factor and thus the highest signal-to-noise ratio of the three. The additional sources of loss which reduce the analysis amplifier Q factor are as yet undetermined, but are likely related to coil geometry and/or trap materials. The circuit diagram for the precision-trap axial amplifier is above in figure 7.7. Figure 7.18 shows the analysis-trap axial amplifier circuit diagram.

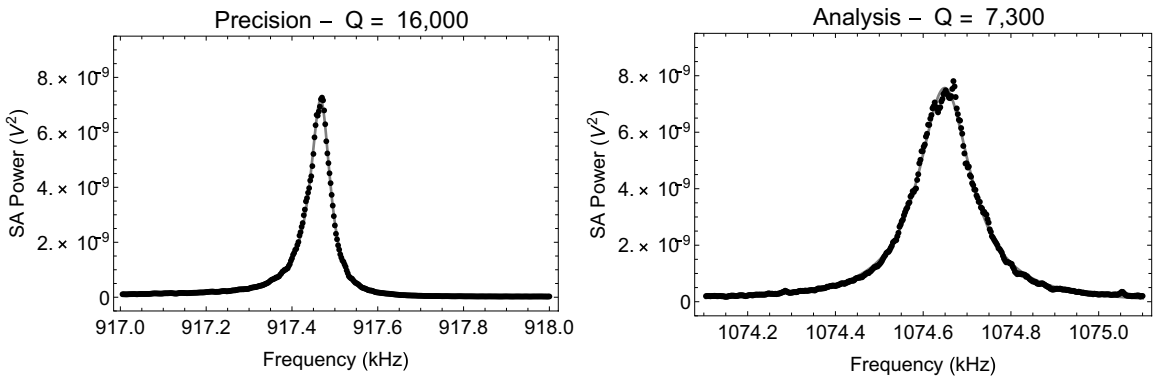

Figure 7.17: Noise resonances for the precision and analysis trap axial amplifiers, using typical FET operating parameters.

For high-precision measurements of the magnetic moment, the most important trap for axial frequency measurements will be the analysis trap. For a 0.1 ppb measurement of the cyclotron frequency, the axial frequency in the precision trap only needs to be measured to around 0.5 Hz. However, to accurately measure the spin state or cyclotron quantum number, we would like to quickly determine the axial frequency in the analysis trap to 0.1 Hz or better. Therefore, it will be beneficial to improve

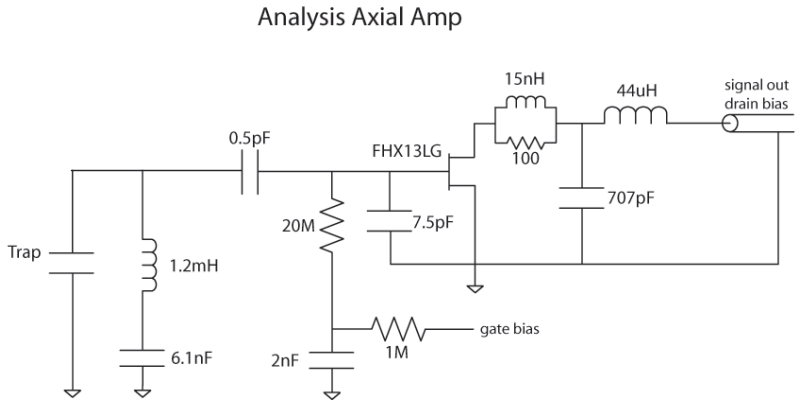

Figure 7.18: Circuit diagram for the analysis axial amplifier

analysis trap axial detection by switching the two amplifiers or by reducing loss in the analysis trap circuit. The same materials and techniques were used in both traps, so the lower Q in the analysis trap is likely due to imperfect coil-winding; however, there may also unexpected additional RF loss, potentially in the iron or from coupling to the spin-flip drive line (sec. 8.5). In that case, we could move the amplifier from a compensation electrode to an endcap; the reduced geometric coupling should be compared to the increased R to determine if this would be beneficial.

7.2.8 Prospects for further improvement of axial detection

Further optimization of the axial detection system is still possible. As we noted above, beyond a certain point simply increasing the Q and R stops being helpful – the bandwidth of the amplifier becomes too narrow and the dip width γ_z too broad. The amplifiers achieved in this work may have high enough R to reach this limit, especially the precision trap amplifier with a Q of up to 24,000. However, currently a choice of R

comes at the cost of reduced signal output due to low transconductance. Independent adjustment of R and gain should be a primary goal of amplifier upgrades.

We note two other effects which could motivate future amplifier work. In the first generation of this experiment, negative feedback applied through the axial drive line was shown to reduce the axial temperature, and positive feedback applied through the FET was shown to increase the axial temperature [43, 35]. Negative feedback applied through the FET should therefore decrease the temperature. If this can be demonstrated using particles thermalized with a transconductance-dependent amplifier, it would motivate further work to reduce loss in the resonator, in order to apply more negative feedback without losing signal. Second, improvements to the FET or its circuitboard may improve the SNR without affecting the Q or R of the resonator.

Active control of feedback and optimizing FET operating parameters

The first step towards improving the detection using the principles outlined in this chapter would be to develop in situ control of feedback. Figure 7.14 shows that the tradeoff between Q and transconductance depends on the matching network detuning. Adding a varactor in parallel with the matching network tuning capacitor would allow us to control the slope of this effect, giving increased signal without sacrificing Q.

A second varactor, added in parallel with the pi-network "second leg" capacitor, would allow us to match the signal after choosing both a transconductance and matching-network detuning, to simultaneously optimize damping and signal-to-noise. Without a second varactor, SNR would be lost when tuning the transconductance, because the pi-net would only optimize matching at one value of FET output impedance.

By controlling these two varactors, we could independently control the quality factor, transconductance, and signal matching efficiency. Without these varactors, we either must balance between FET gain and front-end quality factor by choosing a transconductance, or iteratively change capacitor values over multiple thermal cycles.

A simple step after that would be to combine the understanding and control of feedback with the optimization of the tap ratio. The procedure described above for optimizing the tap ratio relied on an estimate of loss in the FET at given operating parameters; we now understand that loss to be in large part due to negative feedback.

Additionally, it should be possible to use feedback to improve the detection SNR without changing the Q or bandwidth of the amplifier. Figure 7.19 shows that the effective input resistance at the FET gate is nonlinear with respect to transconductance. Therefore, by simultaneously changing transconductance and the tap ratio (and thus the FET-resonator coupling) while holding the resonator loading (equation 7.15) constant, it may be possible to increase the signal amplitude, up to a limit determined by the signal chain noise floor and non-feedback losses in the resonator.

Further materials and wiring improvements

Assuming we determine that the axial temperature is reduced by negative feedback through the FET, or some other way to benefit from reduced resonator losses while keeping a large enough bandwidth, another avenue for improvement in axial detection would be to continue increasing the Q of the resonator itself.

Areas which have not yet been fully analyzed include the effects of dielectric loss in circuitboard materials; resistive loss in spot welds and solder joints; losses in the

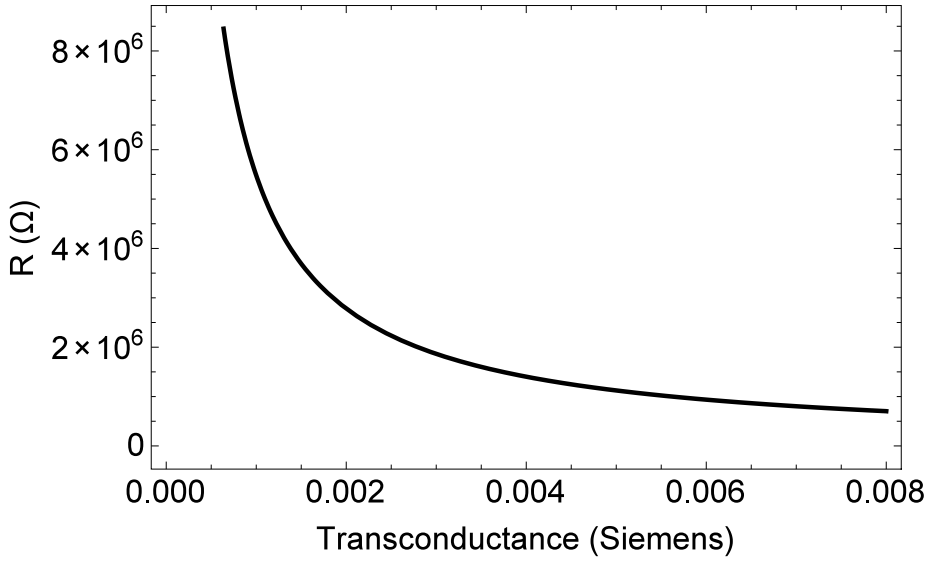

Figure 7.19: Resistive component of FET input impedance vs transconductance, using parameters in table 7.4.

grounding capacitors on current return paths; and dielectric and resistive losses in the trap spacers and electrodes. Parasitic couplings to other elements in the tripod represent another potential loss source; however, similar Q values measured on a test trap and on the experiment indicate this should be a small effect.

FET noise and temperature

The active portion of the detection system can also be improved. The electronic temperature of the FET is known to be higher than its physical temperature [105]. This electronic temperature affects the noise bath in the resonator, which sets the Boltzmann distribution of axial states. The Johnson noise from the FET input impedance, modified by the tap ratio, adds in parallel with the 4 K Johnson noise in

the tuned circuit. The effective temperature for parallel resistors adds as

$$T_{total} = \left(\frac{T_1}{R_1} + \frac{T_2}{R_2} \right) R_{total} \tag{7.16}$$

The increased electronic temperature due to the FET was measured for the 2013 antiproton result [35]. It was improved by the addition of a second cryogenic amplification stage between the axial FET and room temperature, which acted as a buffer against noise being communicated to the particle through feedback. A similar second stage has been built for the new apparatus, but not yet attached. Additionally, because the FET's input impedance (R_2 in equation 7.16) is modified by the tap ratio, the axial temperature could be improved with a different tap-ratio optimization from that of section 7.2.4.

Another parameter that could be improved is the FET noise factor, which helps determine the SNR and has not yet been systematically studied in this experiment. The noise floor and electronic temperature could potentially be improved either by using a different FET, or using two FETs in a cascode configuration. For the moment, satisfactory Q and gain are simultaneously achievable with the new low-loss resonator; but in the longer term the experiment may benefit from improvements to the FET.

SQUID amplifier

A more radical upgrade would involve replacing the FET entirely with a SQUID. While this would add complexity to the experiment, it would have the benefit that a SQUID's electrical temperature is equal to its physical temperature. This could narrow the axial Boltzmann distribution.

Besides the complexity of using DC or RF phase locking electronics to control

the SQUID rather than simply applying bias voltages, a SQUID can only operate in near-zero magnetic field. Achieving this without altering the solenoid is possible, but only if we were to locate the SQUID at the top of the helium dewar. Even there, at 200-500 G the field is too high for the SQUID, so we would have to null the residual field. We could do this during cooldown with a room-temperature bucking coil located outside the neck of the magnet. An NbTi shield around the SQUID would lock in the zero field via flux trapping, at which point the bucking coil could be turned off.

The persistent currents in the NbTi shield would have a much smaller effect on the field at the particle's location than the bucking coil. Nulling the field at the SQUID location with a bucking coil would shift the field by between 3 and 5 PPM, while doing so with persistent currents in the shield would only be a shift of .06 to .15 PPM. Additionally, disengaging the room temperature power supply for the bucking coil would improve stability. The contribution to axial field inhomogeneity would be negligible at this distance, and the shield could be centered to avoid radial gradients. Figure 7.20 shows a potential placement of the SQUID shield and bucking coil.

Locating the SQUID at the top of the dewar would mean a large separation from the trap. This would either require moving the coil up to the same position as the SQUID, or separating the coil and the SQUID. If they were both moved together, the long wire between trap and coil could pick up noise and carry it to the proton. However, if the coil was kept at the tripod, then noise picked up along the long wire would have a larger magnitude relative to the particle signal by a factor of the tap ratio. A Penning trap experiment at MIT [106] implemented their coil near the SQUID, far from the trap; this has the additional advantage of eliminating loss due to

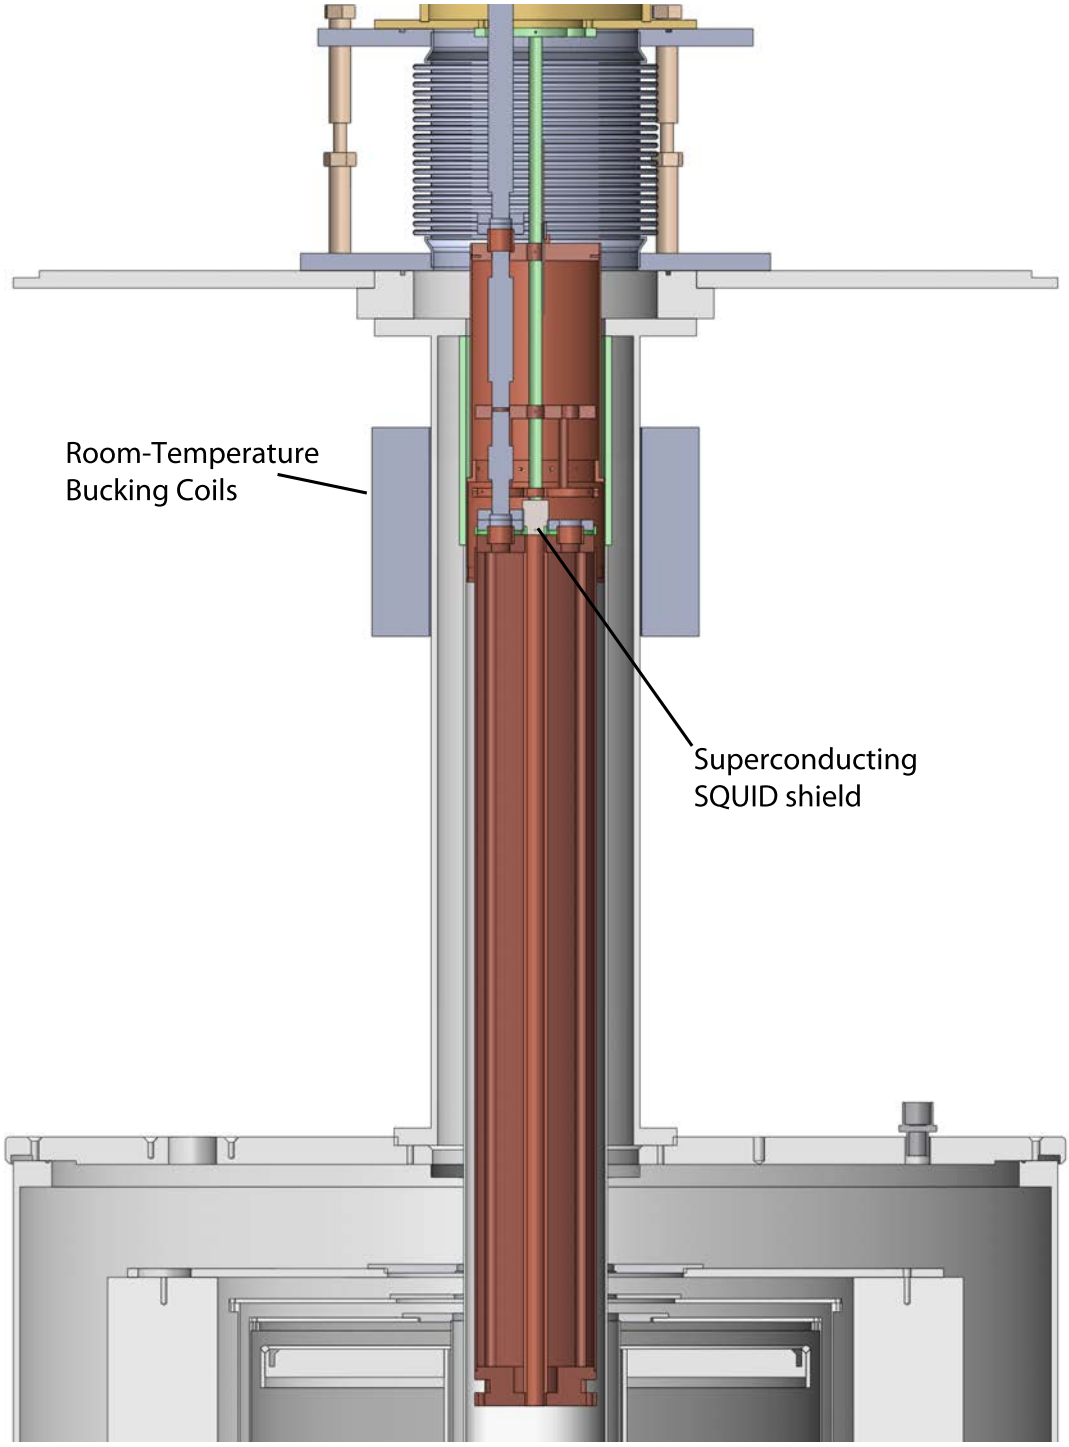

Figure 7.20: Potential placement for bucking coils and superconducting SQUID shield. Note that if the resonator were moved to the top of the dewar, it would require more space than the SQUID shield pictured here.

magnetoresistance in the superconductor. However, over a year of work was required to reduce the noise pickup enough to operate the experiment. In our experiment, this would also require redesigning the thermal isolation baffles to fit the amplifier coil.

7.3 Second stage amplifiers

A second stage of amplification can be added between the output of the FET and room temperature. This adds an additional stage of cryogenic amplification, improving SNR. More importantly, it blocks room-temperature noise from being communicated to the particle through the FET via feedback. In the previous apparatus, adding a second stage amplifier reduced the axial temperature by 2-3 K, down to 8 K. The effects of second stage amplifiers are discussed further in [35] and [44].

The second stage amplifier has not yet been implemented on the new apparatus, but a copper support and two amplifier circuits have been constructed to attach at the top of the dewar. A small OFHC shield will protect the second stage amplifiers from RF noise pickup. Figure 7.21 shows the second stage circuit diagram, while figure 7.22 shows the CAD model of the support post and shield.

The addition of a second stage amplifier affects the behavior of the first stage amplifier. As discussed in sections 7.2.4-7.2.6, the impedance seen by the FET drain can cause FET instability (through positive feedback) or reduce the Q and effective resistance of the front end (through negative feedback). Additionally, poor impedance matching through the pi-network can cause a reduction in signal amplitude.

In the absence of a second stage amplifier, the role of the FET input pi-network is to impedance match between the FET output impedance and a 50 Ohm microcoaxial

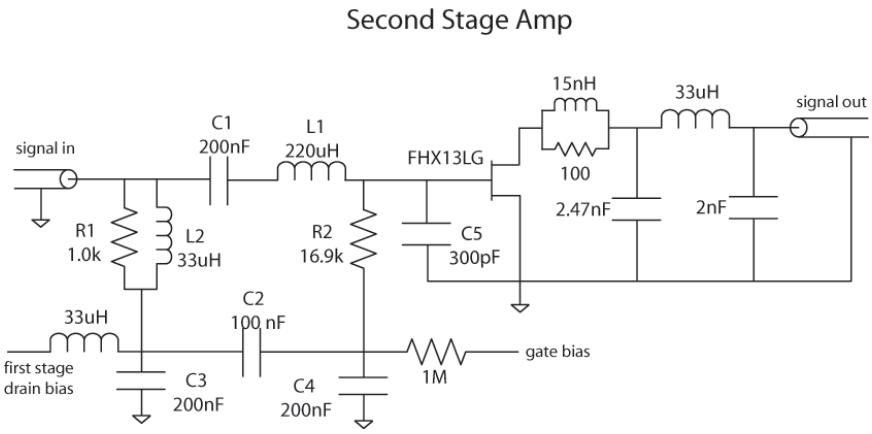

Figure 7.21: Circuit diagram for the second stage axial amplifier

cable to room temperature. The pi-network is then intentionally detuned by a small amount, to ensure stability. The second stage amplifier has an input matching network with a similar role, consisting of $C_1,\ L_1,\ R_1, C_5$ and L_2 in figure 7.21. Together, these impedance-match between the 50 Ohm microcoax and the second stage FET gate.

If the microcoax stage were many times longer than the signal's wavelength, then we could treat the two matching networks separately. However, that is not the case, and each matching network influences the other. We are primarily concerned about the effect of the second-stage input matching network on the first stage output pi-net, because of the risk of FET instability/oscillation and the potentially significant effects of matching network detuning on the amplifier's Q and R.

The second stage input matching network presents a partially reactive impedance to the first stage output network. This shifts the reactance of the first stage output matching network away from its design value. A negative reactance to the second stage input matching network will shift the first stage output matching network higher

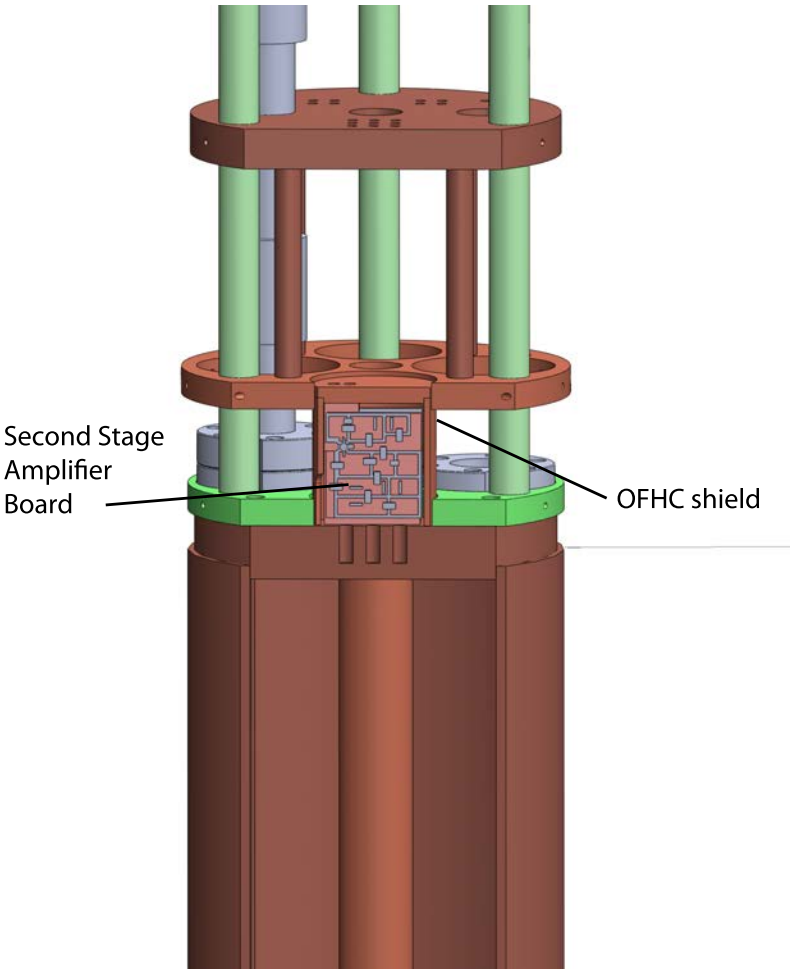

Figure 7.22: Second stage amplifier board CAD model, mounted on helium dewar in its copper shield.

in frequency, risking regeneration or oscillation. This happened in 2014 - a layer of Q-dope fell off of L_1, shifting the center frequency of the input matching network. This shifted the first stage output matching network and caused regeneration.

Additionally, if the resistive part of the second stage matching network input impedance is far from 50 Ohms, it can also shift the first stage pi-network. Impedance $> 50\ \Omega$ will result in an effective shift upwards of the pi-network center frequency. Poor second-stage input matching can thus risk causing regeneration or oscillation.

The second stage input matching network should therefore be tuned using a network analyzer before it is attached to the experiment. The input reactance at the amplifier front-end frequency should be made positive to avoid regeneration, and should be as close to zero as possible to optimize signal matching. The input resistance should be at or below 50 Ω to avoid regeneration, and as close to 50 Ω as possible to maximize signal. With an approximately correct set of component values, fine-tuning can be done using only two components. Using a network analyzer in resistance-reactance mode, the value of C_5 should be adjusted to change the reactive component of the input impedance. (The reactance also depends on C_1 and L_1 but adjusting C_5 is the most straightforward way to make small changes in reactance without affecting other circuit parameters). Near the resonant frequency of the matching network, increasing C_5 will increase the input reactance at a given frequency.

Once the reactance is small and positive, the resistance can be adjusted by changing the value of R_2. Using the network analyzer in gain-phase mode, the resistor can be changed to minimize reflection and maximize the signal coupled into the FET.

Discussion of second stage amplifiers in section 4.3 of [107] includes an improved

tuning algorithm for the second stage amplifier output matching network to avoid positive feedback while improving signal matching, by offsetting the output center frequency above the input center frequency. After the second stage amplifier's input matching network has been tuned, its output matching network should be tuned following that procedure before it is implemented on the experiment.

7.4 Cyclotron amplifiers

The cyclotron amplifier operates similarly to the axial amplifier – the particle's motion induces a current in the electrodes, which is forced through a large resonant impedance to provide damping and pick up the signal. However, the higher cyclotron frequency implies significant technical differences. For example, producing a resonant frequency 100X higher with a similar capacitance requires a 10,000X smaller inductor.

This inductance is provided by a helical coil of thick, silver-clad copper wire inside a copper shield, together comprising a coaxial resonator with a helical center conductor (as described in [108]). The circuit board for the cyclotron amplifier is identical to that for the axial amplifier, with component values adapted to the higher frequency.

While the axial frequency is set by the electrode voltages, and can thus be adjusted to match the detector, the cyclotron frequency is set by the magnetic field. The field is chosen to separate the proton's frequency from nearby radio or television stations, and to give high relative precision. Once the field value is chosen, it is difficult to adjust with the experiment in place. We therefore adjust the cyclotron amplifier frequency to match the field by changing the inductance or adding parallel capacitance.

The trap and coil contract on cooldown, moving the resonator frequency by around

1.5 MHz – several times its linewidth. This cooldown shift is mostly predictable, but can vary by a significant fraction of the linewidth, so even with properly chosen components we still need in situ tuning. A varactor in parallel with the trap capacitance allows us to tune the front end over a range of ∼100 kHz.

The generation 2 proton experiment has two cyclotron amplifiers. One is connected to the precision trap, while the other is connected to the cooling trap.

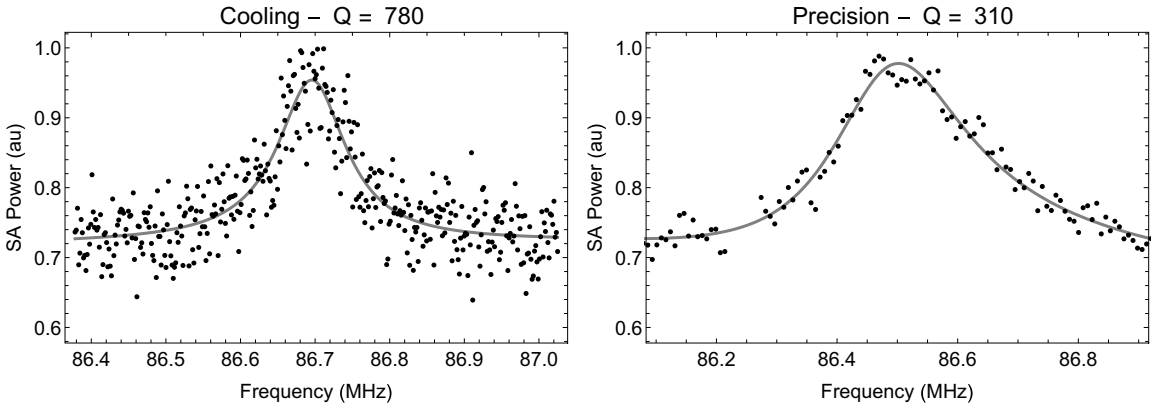

Figure 7.23: Noise resonances for both cyclotron amplifiers.

7.4.1 Precision trap cyclotron amplifier and FET switch

The precision trap cyclotron amplifier is used for damping the cyclotron energy of protons during loading; for reducing to a single proton using relativistic decays; and for diagnostic tests which make use of those decays. However, for a precision measurement we intend to use one of the techniques described in chapters 5 and 6, which do not require a cyclotron amplifier while taking data. The precision trap has been optimized for these methods in ways which add loss to the cyclotron amplifier. Part of the current return path follows the shield of the anomaly-drive coax (sec. 8.5);

this has a higher inductance than a strap, and contributes some loss at the cyclotron frequency. Additionally, for the simultaneous spin-anomaly method detailed in chapter 6, we need to measure relatively small changes in the cyclotron quantum number. The cyclotron motion must evolve undisturbed except when subjected to resonant drives. We therefore added a second FET in parallel with the resonator, a method developed in [109]. When biased to allow current, this FET shorts the resonator and eliminates the interaction between the particle and the thermal bath. When the FET is biased to block current, the resonator operates normally; however, this also contributes loss to the circuit. Figure 7.25 demonstrates the FET switch shorting the resonator.

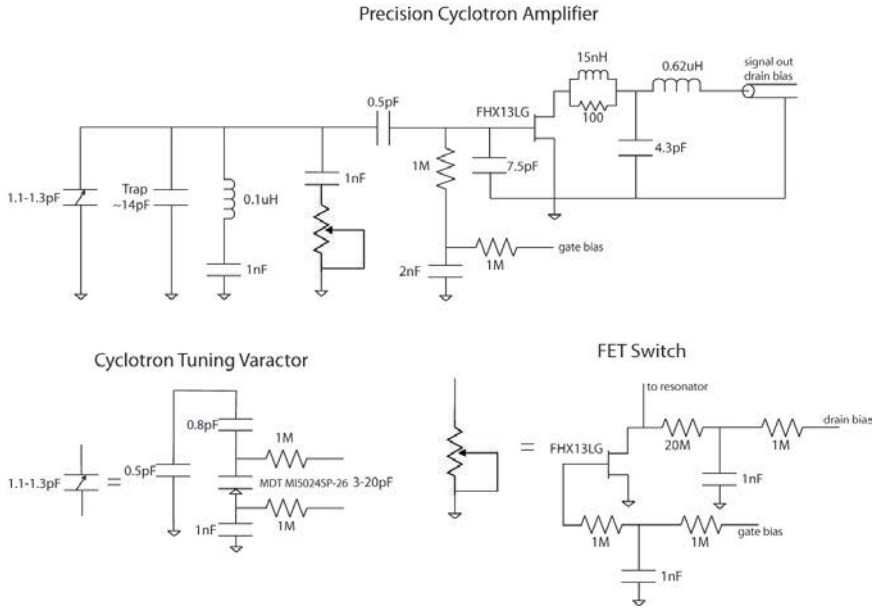

Figure 7.24: Circuit diagram for the precision cyclotron amplifier, with tuning varactor and FET decoupling switch.

The precision trap cyclotron amp currently has a Q factor of ∼320 and a measured damping time of 220 seconds. This is sufficient to measure ν_c via decays and reduce

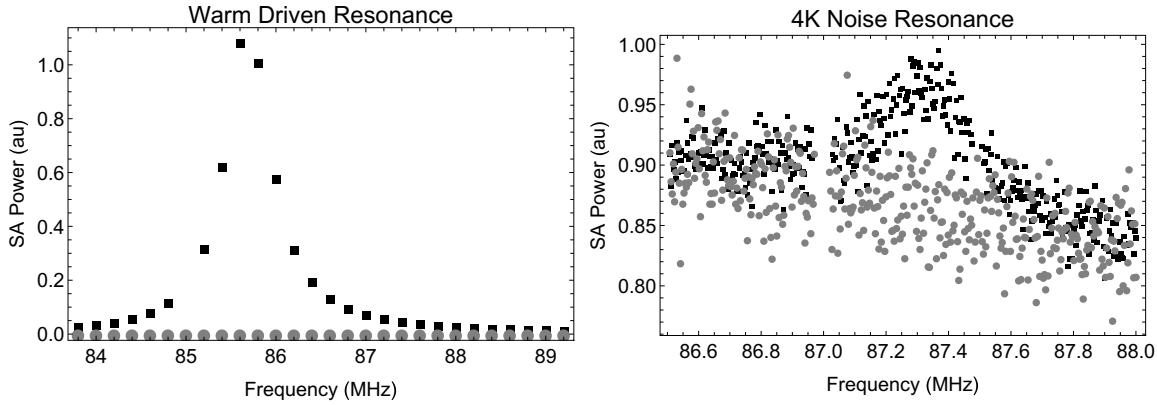

Figure 7.25: Precision cyclotron amplifier resonance profiles with the FET switch open (black squares) and closed (gray circles), for room-temperature driven and 4 K noise resonance scans.

to a single particle – see chapter 9 for results with the new apparatus.

7.4.2 The cooling trap cyclotron amp

The cooling trap will be used to initialize protons by selecting new cyclotron states from the thermal distribution. This is essential for spin-flip methods, where a sub-thermal state is needed to resolve spin flips in the analysis trap. The limit on acceptable cyclotron states for the anomaly-drive method depends on experimental parameters, but it will also need to be re-selected if it grows too large. The effective resistance of the cyclotron amplifier sets the damping time [1]

$$\tau_c = \frac{1}{\gamma_c} = \frac{1}{2}\left(\frac{e\kappa_1}{2\rho_0}\right)^{-2}\frac{m}{R} \tag{7.17}$$

where κ_1 is a constant depending on trap dimensions, calculated in [45].

To improve the Q factor of this amplifier, we adopted some of the same changes made to the axial amplifiers. Macor spacers and alumina feedthroughs were replaced with quartz and sapphire. Additionally, we studied the effect of magnetic field on the

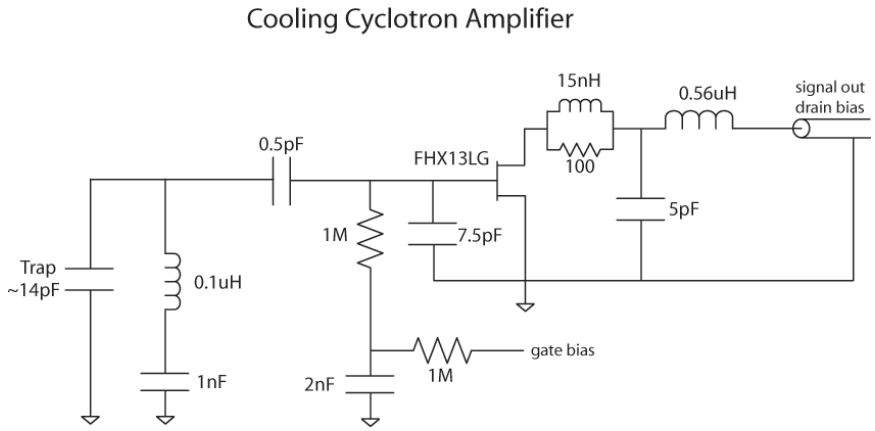

Figure 7.26: Circuit diagram for the cooling cyclotron amplifier

resonator with a test capacitor and with a test trap (figure 7.27). While there is some effect at high fields, the magnetoresistive loss at 3-4T does not limit resonators with Q values of over 1000, so we conclude it is not significant for our amplifiers.

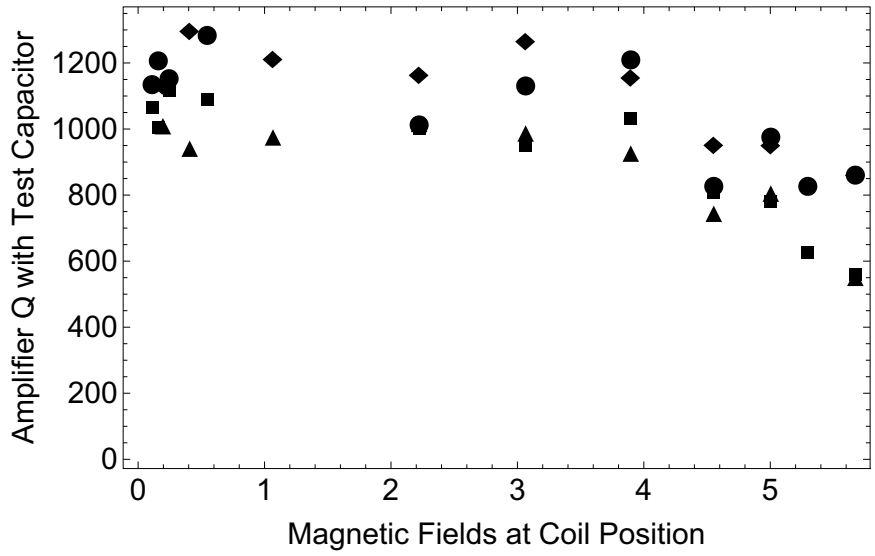

Figure 7.27: Q factor vs magnetic field for multiple 85-90 MHz cyclotron amplifiers, on test setup.

For the cooling cyclotron amplifier, we have not included either a varactor or a

FET switch. The FET switch is unnecessary, because we will only transfer protons into the cooling trap when we want them to interact with the amplifier. The varactor is also thought to be unneeded in this trap, although more testing is required. The cooling trap is located in the magnetic field gradient from the iron ring in the analysis trap. Therefore, at different axial positions in the cooling trap, the particle will have a different cyclotron frequency. The cyclotron damping rate is much slower than the axial frequency, so the average field over the course of an axial oscillation should be the relevant quantity for interaction with the amplifier. Therefore, using asymmetric voltages to shift the particle's center position in the gradient should adjust the cyclotron frequency. This technique has not yet been demonstrated, so a varactor may prove necessary, but its effect on the amplifier Q and resistance should be considered.

The current cooling cyclotron amplifier has a Q of 780 and effective resistance ∼60-90 kΩ. While this is good (a Q of 340 was used in [34] on the previous experiment, for example), it can be improved – a more carefully optimized cyclotron amplifier recently constructed in our group [62] has a Q ∼1200. Techniques used to construct that coil can be adopted for this apparatus (or the coil itself could be relocated).

7.4.3 Superconducting cyclotron amplifiers and magnetic field

We studied construction of a superconducting cyclotron amplifier for the cooling trap. The initial results were promising, with extremely high Q measured in a test setup. However, when the resonator was inserted into the magnet, its Q dropped quickly and dramatically. Our results indicate that the strength of magnetoresistance in NbTi is frequency-dependent; at a field with almost no magnetoresistance in our

1 MHz coil, the loss in the 86 MHz coil was already increasing dramatically. Other experiments have used superconducting cyclotron coils in Penning traps [19] at lower magnetic fields and frequencies; we believe that their results and ours are consistent.

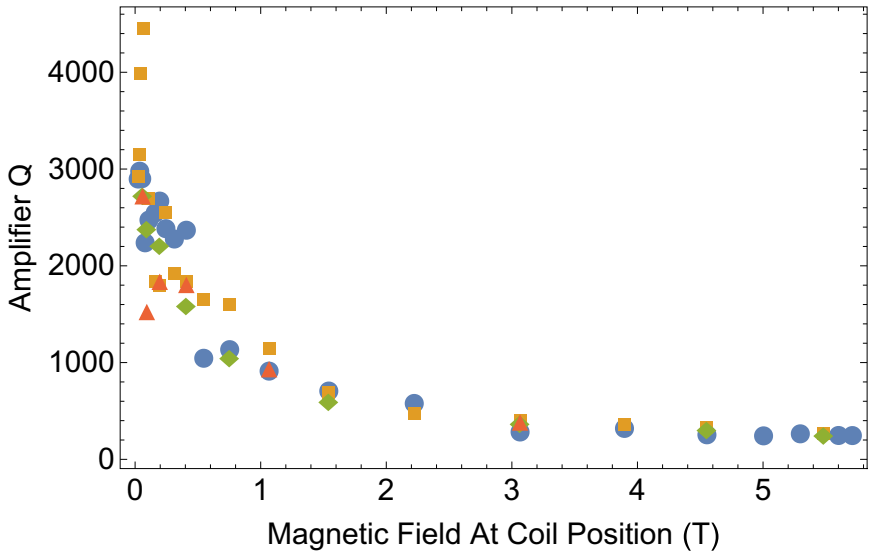

Figure 7.28: Q factor vs magnetic field for two superconducting NbTi cyclotron amplifiers – one toroid and one solenoid – on a test setup.

The magnetoresistive limitations of the amplifiers beg the question of whether the magnetic field strength should be reduced. While this would allow us to use a superconducting cyclotron amplifier, my conclusion is that we should keep our field at its current level. Reducing the field would hinder the quantum walk measurement, by reducing the acceptable Rabi frequency for a given level of power broadening. Similarly, for the single-spin-flip measurement a higher field gives improved relative precision for the same absolute frequency uncertainty. Essentially, we have shown that we can produce amplifiers with high enough effective resistance to make a precision measurement at our current field. Additionally, the increased damping rate in the new loading trap (section 8.1.2) should make the cyclotron Q less critical. We therefore

would lose more in relative precision than we would gain with a lower field value.

Chapter 8

A New Apparatus for Sub-ppb Measurements

A large part of this thesis work was devoted to constructing an apparatus for a sub-ppb measurement using the new methods proposed in chapters 5 and 6. The new apparatus constructed included Penning trap electrodes, cryostat, electronics and wiring, and vacuum system. In this chapter, we provide details on the new system, focusing in particular on elements relevant to enabling a high-precision measurement.

- A precision Penning trap electrode geometry was designed and implemented to minimize magnetic field inhomogeneity.
- Three additional collinear Penning traps were constructed, to measure spin states, allow fast cyclotron damping, and create an antiproton reservoir.
- Finishing methods were devised to produce electrodes with high dimensional precision and low surface roughness (which could heat the cyclotron motion).

- An electron-beam welded titanium cryostat was constructed for the apparatus, as well as a cryogenic support structure and thermal insulation.
- To accommodate the four traps and new measurement methods, the experiment wiring was freshly designed for both DC trapping and RF control. (Improvements to RF detection are discussed in more detail in chapter 7.)
- A cryogenic positioning system was designed and built to enable in situ alignment between Penning trap electric and magnetic fields, and to stabilize the position of the particle relative to the magnet.

8.1 Penning trap design and construction

The new apparatus includes four Penning traps in the electrode stack, each of which is designed to enable high-precision measurements. The precision trap was designed and constructed with a new geometry, which will improve magnetic field homogeneity. The analysis trap is designed to improve axial stability by inhibiting RF noise couplings to the cyclotron state. The new cooling trap will provide fast cyclotron damping, allowing quick selection of particles with small cyclotron radii and high axial stability (see sec. 3.2.3), and the loading trap has a large radius for efficient antiproton trapping (and will act as an antiproton storage reservoir). All of these were finished with a new polishing procedure.

In this section, we will discuss in detail the design of the precision trap. We will also summarize the goals and design of the loading and cooling traps, which will be discussed in more detail in [62]. All of these traps were constructed and finished with

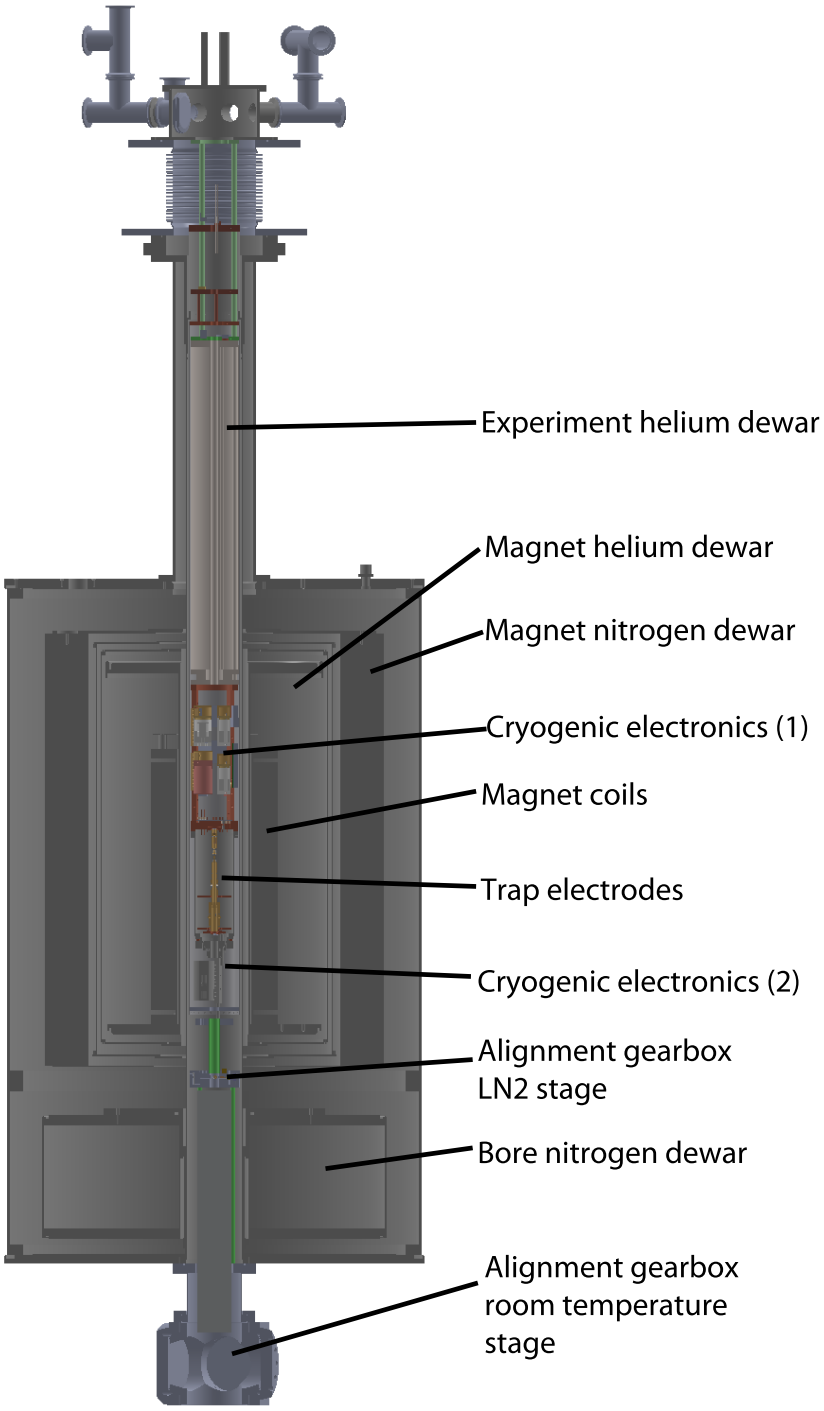

Figure 8.1: CAD model of the experiment with the magnet, electronics, and cryogen dewars.

procedures developed during this thesis, which will also be described here.

In Section 2.1, we discussed the electrostatics of the Penning trap. Orthogonality and compensation determine the height-to-diameter ratios of the particle-facing surfaces of the Penning trap (figure 2.2). Balancing coupling to the amplifier with sensitivity to noise and voltage instability, we choose a trap radius $\rho = 3$mm for the precision trap, completing the definition of the particle-facing dimensions.

8.1.1 Magnetic design in the precision trap

The remaining electrode dimensions do not affect the electrostatic trapping potential seen by the particle. Rather than choose them for ease of precision machining, we considered the effects of trap materials on the magnetic field. A careful choice of electrode dimensions allows us to minimize the residual magnetic field inhomogeneity in the precision trap, which will yield a narrow linewidth and ultimately a more precise measurement.

When considering the magnetic field distortions due to trap materials, we follow section VI.B of [1], integrating the effect of small rings of magnetic dipoles. This yields an expansion in terms of Legendre polynomials for the field perturbation at $\mathbf{r}$

$$\Delta\mathbf{B}(\mathbf{r}) = \sum_{l=0}^{\infty} B_l r^l \left[P_l(cos(\theta))\hat{z} - (l+1)^{-1} P_l^1(cos(\theta))\hat{\rho}\right] \tag{8.1}$$

where P_l and P_l^1 are a Legendre polynomial and an associated Legendre polynomial, and B_l is the l-order magnetic field perturbation strength, given by

$$B_l = (l+1)(l+2)2\pi \int \rho' d\rho' dz' M\left(\rho', z'\right) \left(r'\right)^{-l-3} P_{l+2}(cos(\theta')) \tag{8.2}$$

where $M(\rho', z')$ is the magnetization of the material located at (ρ', z').

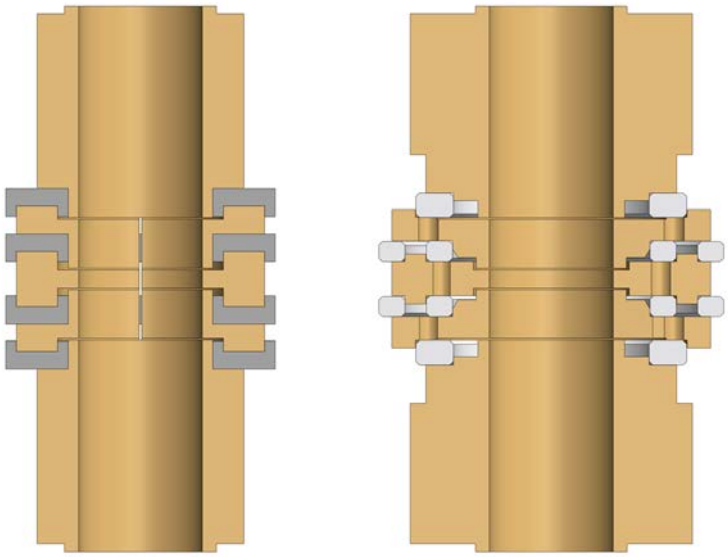

Figure 8.2: 2011 (left) and 2018 (right) versions of the precision trap – particle-facing dimensions are identical but electrode and spacer shapes have been altered to optimize magnetic field homogeneity and mechanical stability.

In general, when considering magnetism in the trap, we are mostly concerned with the B_0 and B_2 terms. As long as the Rabi frequency of a given drive is small compared to the axial frequency, the transition frequency will only depend on the average field seen during the axial oscillation. We can therefore neglect the effect of terms which are odd in z, since they average to zero. Meanwhile, the effects of terms with $l \geq 4$ fall off quickly with distance from trap center, and are small compared to B_2 [1].

B_0 is a shift in the homogeneous field from the solenoid, so we aren't particularly concerned with the value (except insofar as it affects the tuning of the cyclotron amplifier). However, we are concerned with stability in B0. The effect of magnetic field drifts on the future measurement are described in section 6.4.3 for the quantum walk method, and section 5.3 for the simultaneous SOF method. We expect there to be slow, predictable drifts due to thermal expansion and contraction of structural

materials and pressure changes in the various cryogen spaces. However, if a piece of material with nonzero magnetization changes position – for example, due to vibrations in the tripod or experiment dewar – then the B_0 seen by the particle due to that material would also change. Therefore, materials with nonzero magnetism must be carefully anchored to structural support so vibration amplitudes are kept low.

That being noted, B_0 is not a significant concern when designing the trap itself. The electrodes cannot shift position relative to trap center because they set the position of trap center. In trap design, we are therefore concerned with B_2 – the magnetic field gradient needs to be very large in the analysis trap to enable spin and cyclotron state analysis, and very small in the precision trap to reduce the transition linewidths.

For our analysis of magnetic field curvature in the precision trap, we look to the second-order term in equation 8.2:

$$B_2 = 24 \int \rho' d\rho' dz' M(\rho', z')(r')^{-5} P_4\left(Cos(\theta')\right) \tag{8.3}$$

Plotting this as a function of ρ and z in figure 8.3, we note regions of positive and negative sign, with multiple zeroes. By careful placement of trap materials, we can use their magnetism to "shape" the magnetic field at the center of the trap. The dimensions of our spacers are chosen via numerical calculations of B_2 such that their contributions cancel. This both reduces B_2 overall and allows us to change spacer materials without affecting the linewidth[1].

The shapes of the copper electrodes are chosen in a similar way. However, rather than canceling to zero, we use the copper's magnetization in our favor. We can calculate the residual B_2 in the precision trap due to the presence of the analysis

[1] In chapter 7 we discuss significant improvement to detection from replacing our spacer material, a modification enabled by this design.

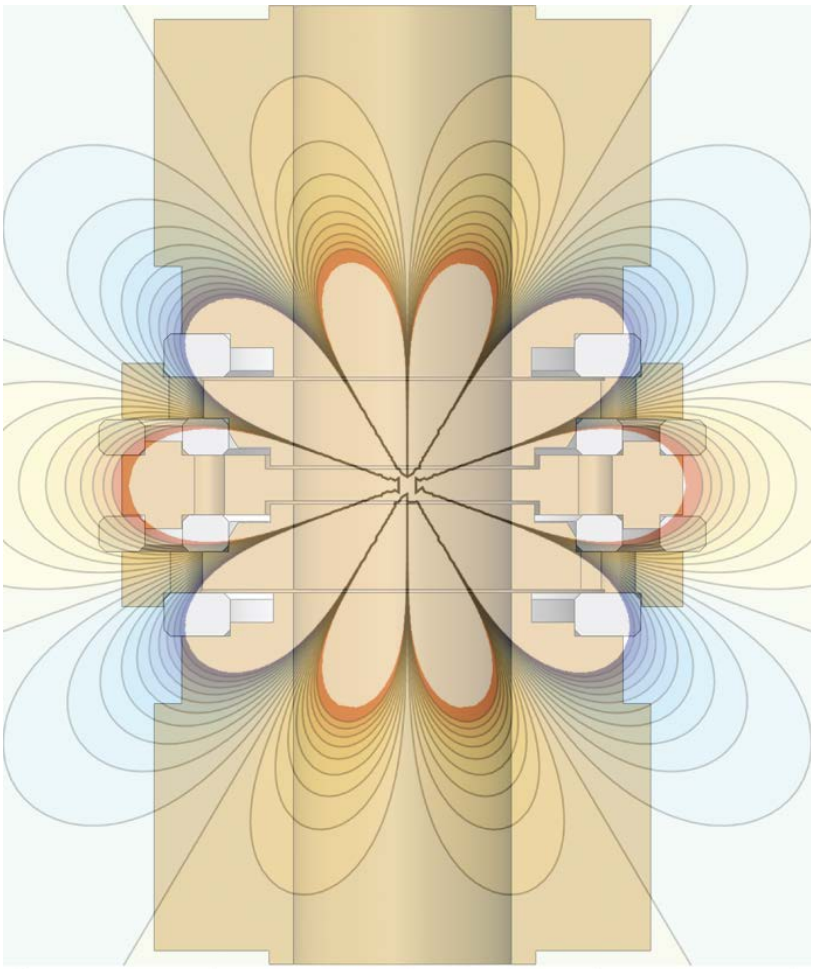

Figure 8.3: New precision trap superimposed with contours of the B_2 contribution, equation 8.3. Spacer geometry is chosen such that B_2 contributions sum to zero, while electrode geometry is chosen to cancel the residual effect from the iron ring in the analysis trap.

trap iron ring a few inches away, and design the copper electrodes to cancel out this residual bottle. The predicted B_2 field in the precision trap due to the iron ring is 0.28 T/m^2; the copper electrodes have been designed to give a contribution with approximately equal magnitude and opposite sign.

Different values are available in the literature for the magnetization of copper at 4K and 5.6 T. In designing the trap to cancel the residual bottle from the iron ring, we used the value presented in [1]. However, the uncertainty on this value implies that the bottle-field cancellation may not be as exact as desired. The sign and approximate value are agreed on, so we expect the net B_2 to be less than the 0.28 T/m^2 due to the iron ring. If B_2 is found to be larger than desired when precisely measured at Northwestern (see section 10.1.7), a new set of electrodes could be designed relying on a different value of M for copper in equation 8.3. However, a better solution would be to cancel the bottle field in the existing trap as proposed in sec. 10.2.4, either by adding material with known magnetization or using current coils around the endcaps.

8.1.2 Analysis, cooling, and loading traps

The new apparatus also includes a high-gradient analysis trap, as well as two new traps. The design of the analysis trap is largely similar to that used in the 2013 apparatus, and design of the new traps is discussed in detail in [62]. We will provide a brief discussion of each here.

Analysis trap

In the analysis trap, we maximize the magnetic field gradient with a high-purity iron ring electrode and a 1.5mm trap radius. With a 3mm trap radius in 2009, the field inhomogeneity was too small to resolve the axial difference due to a spin flip given experimentally achieved axial frequency stability [34].

This choice of radius may be worth re-evaluating going forward. If spin-state identification is limited by the FET noise floor or 1/f noise in the detection chain, then the factor-of-4 increase in B_2 from the smaller trap is a significant advantage. However, the smaller electrodes increase the coupling to magnetron and cyclotron noise. As discussed in sec. 3.2.3 and [35], the n_c dependence of the axial frequency stability seems to indicate it is limited by cyclotron state transitions in the bottle. In this case, increasing B_2 does not help spin-state identification – the instability per cyclotron transition scales at the same rate as the frequency shift per spin flip. Additionally, the rate of cyclotron state change is larger in a smaller trap – the coupling to RN noise scales with $1/\rho^2$, while the anomalous heating by surface effects increases with $1/\rho^4$ (see sec. 5.2.1) – if n_c instability is mostly responsible for problems with spin-state identification, then we may benefit from returning to a 3mm radius.

This could be demonstrated by a careful comparison of axial stability in the precision and analysis traps. if we find that the frequency stability in the precision trap is sufficient to see the smaller spin-flip frequency shift that would come from a larger analysis trap, it could be worthwhile to experiment with one.

The iron ring and copper compensation electrodes feature inner dimensions scaled by a factor of 1/2 from those of the precision trap, to preserve compensation and

orthogonality. However, each endcap is divided into two sequential electrodes, which combined are 1/2 as long as the precision trap endcaps. This effectively creates a seven-electrode trap, rather than a five-electrode trap like the precision trap. If the endcap segments are held at the same DC potential, or both grounded through cryogenic resistors, the analysis trap will still behave like an orthogonal, compensated five-electrode Penning trap. Dividing the endcaps in this fashion allows us to split only the outer segment, to provide the radial asymmetry required for magnetron sideband cooling. This increased distance to the particle significantly reduces cyclotron heating rates due to Johnson noise or RF leakage, while requiring strong but realistic magnetron cooling drives – see [62] for a detailed calculation.

Cooling trap

A cyclotron cooling trap was also constructed for this apparatus. As explained in sections 3.2.3 and 4.3, selecting a sub-thermal cyclotron radius is essential for single spin flip detection, but requires multiple cycles of thermalization and radius measurement. While the quantum walk measurement may allow us to re-thermalize less frequently, the new apparatus is designed to enable either a single-spin-flip or a quantum walk measurement. The cooling trap would significantly speed up a single-spin-flip measurement by reducing the time required to thermalize the cyclotron state.

The cooling trap has the same radius as the analysis trap, 1.5 mm. Additionally, the cooling trap uses a single split electrode, rather than a ring and two compensation electrodes. This makes it uncompensated, and thus axially anharmonic, but increases the geometric coupling between the amplifier and particle by an additional factor of

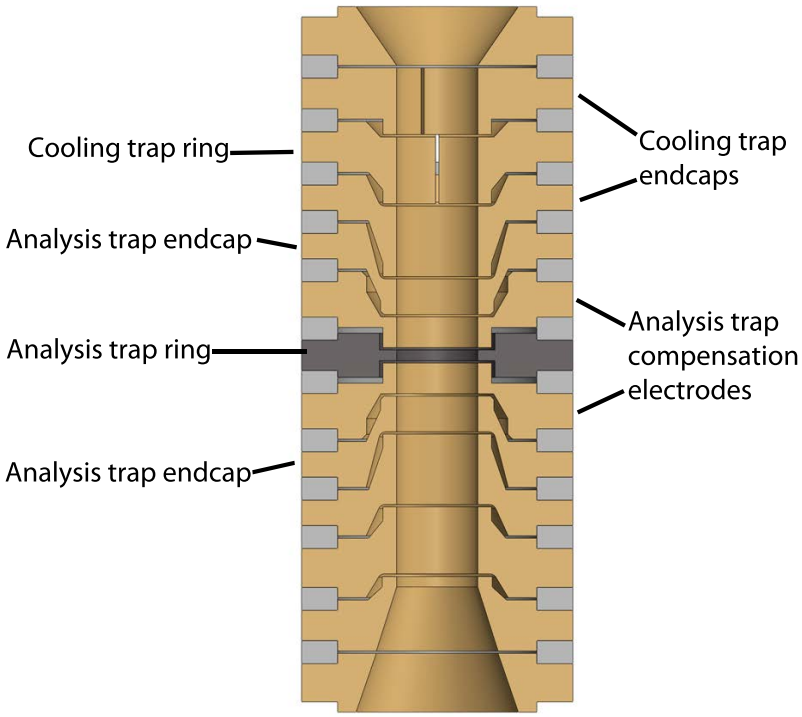

Figure 8.4: The 1.5 mm radius electrodes comprising the analysis and cooling traps.

2.5 [62] compared to a five-electrode trap. The damping rate of energy into the cyclotron amplifier is

$$\gamma_c = \left(\frac{q\kappa_c}{2\rho_0}\right)^2 \frac{R}{m} \tag{8.4}$$

so the reduced ρ and increased κ should reduce the damping time by a factor of $\sim$27.

Loading trap

The loading trap is located below the precision trap, and is intended to simplify antiproton loading and serve as a reservoir for antiprotons. We attribute some of the difficulty we experienced in 2012 loading antiprotons (section 3.1) to the small trap radius. We only were able to load antiprotons by trapping and electron-cooling them in a short potential well on the electrode adjacent to the degrader. However, other experiments in the ATRAP collaboration have been able to catch many antiprotons in wells with the same or larger length:diameter ratios than our initial attempts. We concluded that whatever effect caused antiproton loss was likely nonlinear in length:diameter, and that a larger-diameter trap is likely to catch and cool more antiprotons. Therefore, the loading trap was constructed with twice the radius (maintaining the orthogonalized length:diameter ratios of the trap electrodes.)

In addition to simplifying antiproton catching, this trap is intended to serve as an antiproton reservoir. Demonstrated in [110] by the BASE collaboration, the reservoir would store a cloud of antiprotons for a long period of time, and deterministically split off low numbers of antiprotons from the cloud for transfer up to the precision and analysis traps. This protects against disruptions in the antiproton beam during runtime, and allows measurements to continue while the beam is shut off. The loading

trap endcaps are segmented to allow splitting of stored antiproton clouds.

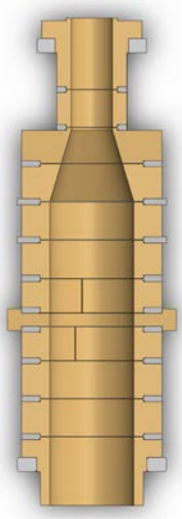

Figure 8.5: The 12 mm diameter electrodes comprising the loading trap, together with the conical transfer electrodes connecting to the precision trap.

8.1.3 Trap electrode construction

A new procedure for finishing and plating the electrodes was designed to reduce the potential effect of surface inhomogeneity on the background cyclotron transition rate (section 5.2.1), as well as to achieve high dimensional precision.

The Penning trap electrodes for this experiment were fabricated in Harvard University's SEAS instrumentation shop. Several copies of each piece were fabricated and measured in the lab; we could thus select a set of electrodes with particle-facing dimensions within .0002" of the ideal. The effect of these remaining imperfections on the measurement will be small because of the invariance theorem in equation 2.14.

However, the electrodes initially have a rough surface finish from machine tooling. In order to maintain the precise dimensions while also creating an extremely smooth and precise bore, we finish the inner surfaces of the electrode ourselves in the Harvard student machine shop. This allows us to apply an abrasive while holding the electrode in a soft teflon lathe collet, ensuring that the soft high-purity copper of the electrodes doesn't get deformed during polishing.

While constructing this apparatus, the quality of the electrode surface finish received particular attention. As discussed in section 6.4.2, heating of particle motions has been observed in ion traps due to interactions with the trap surface. We hypothesize that surface effects could be responsible for some of our observed heating. We can then conjecture that improving the surface finish would reduce the density or intensity of surface noise sources, and therefore reduce the rate of cyclotron transitions.

We decided to seek a polishing compounds which held grit in solution with lubricant, and apply pressure using a soft cotton swab. Well-known methods for Penning trap construction rely on applying abrasive sheets to the inner bore, and applying pressure with a wooden cuticle stick or dowel rod. We found, however, that removed material ("swarf") would occasionally embed or clump up in the sheet, leading to small scratches on top of the mirror finish, and that uneven pressure, inevitable when applied by hand using a flexible wooden rod, would give a measurably and visibly convex shape to the bore.

The particular requirements of a Penning trap experiment made finding the appropriate combination of grit, lubricant, speed, and pressure a nontrivial task. We rejected nearly 20 compounds before settling on a satisfactory procedure. Jeweler's

rouge is usually used for polishing soft metals like copper, but it is mostly ferric oxide; using it would risk impregnating the electrodes with magnetic impurities. Commonly used compounds for polishing harder metals include garnet solutions, which left scratches in the copper, and alumina, which polished extremely slowly[2]. Heavier oil-based lubricants caused the cotton fibers in the swab to clot together and build up swarf swarf buildup, scratching the surface. Some alumina and silicon carbide solutions were too acidic or basic and etched the electrode surfaces. Other polishing compounds left a residue which was either difficult to remove (other oil-based lubricants) or difficult to identify (Simichrome paste, previously used as a final step, advertises that it leaves an "invisible protective coating" on metal).

We finally concluded the best combination of smooth finish, low residue and flat surface was obtained using water-soluble diamond lapping paste[3]. The paste is spread onto a cotton swab; a film of lubricant[4] is then coated over the paste and swab using a wash bottle. The swab is then inserted into the bore of the electrode, which is mounted in a teflon collet on a rotating lathe. Soft pressure is applied as the swab is run back and forth evenly over the surface of the bore. Whenever the lubricant begins to dry up or run out of the bore, additional lubricant is sprayed from the wash bottle. (Several swabs are used at each grit size, as the grit also slowly runs out of the electrode bore.) This creates a slurry composed of lubricant, diamond particles, and the binding agent from the paste. When applied to an electrode, the diamond particles both elastically deform and cut into the copper surface. A paste with a particular diamond particle size should leave a somewhat finer surface finish than its

[2]Likely due to elastically deforming rather than abrading the copper
[3]Sandvik Hyperion, particle sizes 80-1 μm, available from McMaster-Carr
[4]Hyprez W water-based lubricant for diamond lapping paste, also available from McMaster-Carr

particle size, because the particles are suspended in the slurry and only the edges will cut into the copper.

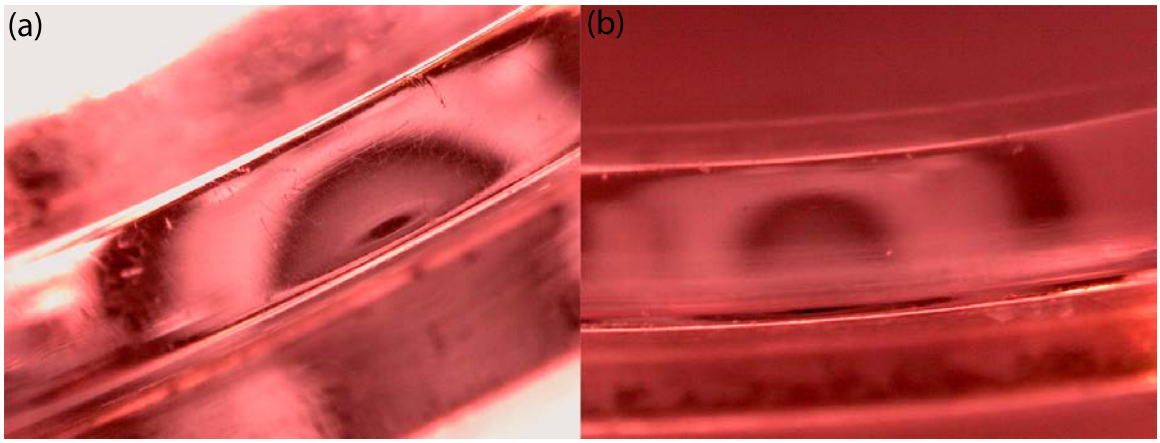

Figure 8.6: Ring electrodes polished using (a) polishing paper, and (b) diamond paste slurry. While both methods achieve a smooth mirror finish to the eye, the diamond paste is more dimensionally precise and smoother under a microscope.

Because the particles only cut with their edges and corners, the paste removes much less material than the fixed-grit film or paper. If the machining marks on the bore before polishing are too large, their edges will be smoothed but the actual surface will not be flat. Therefore, the electrode bore must be reamed or honed to a smooth initial condition and an appropriate pre-polish diameter[5].

Since dimensional precision is important as well as surface finish, we carefully chose an initial diameter and calibrated the amount of material removed at each polishing step. To measure the removed material we sequentially inserted precise gauge rods to find the radius where they contact the bore – measurements using an optical microscope or collimation of a laser failed due to diffraction and the nonzero

[5]We had success making a small cut with a reamer or a fine-grit FlexHone with thick honing oil.

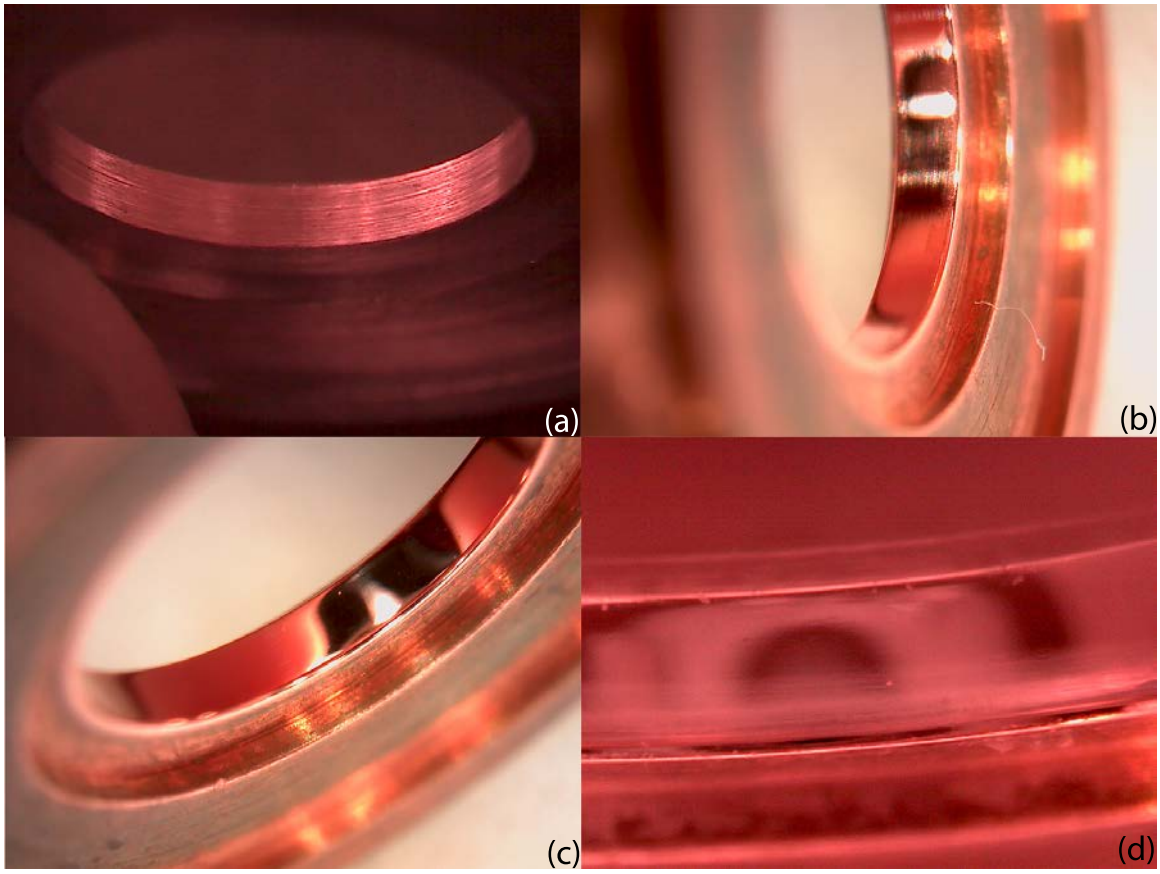

Figure 8.7: Precision ring electrode at different stages of polishing with diamond paste slurry. (a): as returned from reaming. (b): After polishing with diamond grit down to <4 μm. (c): After polishing with diamond grit down to <2 μm. (d): Final polish.

length of the bore, and other methods like ball gauges and calipers were inconsistent. While the gauge rods scratch the surface, they can be used as long as enough polishing remains to remove the scratches. Up to that point, the remaining material should be measured after each step to achieve the correct final diameter.

To calibrate the final steps, we measured the final diameter of a test piece for each electrode type and recorded the amount of material removed after each grit size. We then followed an identical procedure for identical electrodes, assuming the material removed would be similar. Little material is removed during the final steps, so this adds minimally to machining tolerances on the bore diameter.

The main dimensional flaw of electrodes polished with diamond paste is that the corners are somewhat rounded. The soft cotton swab spreads the grit around the corner of the electrode rather than keeping it on a plane. However, rounded corners give us an electrostatic surface closer to an ideal cylindrical Penning trap than when the entire surface is concave.

After the surfaces are polished to a mirror finish, leads were brazed into the electrodes in our hydrogen oven, using a silver-copper eutectic alloy. In addition to attaching the leads, this process also cleans the electrodes of surface impurities (since the hydrogen reduces the surface of the copper), and anneals residual tension in the copper from machining. However, the brazing process also causes the grains in the copper to re-align, which causes grain boundary growth in the copper surface [111], as shown in figure 8.8. After brazing, these grain boundaries can be removed by re-polishing using the smallest grit size. However, because the annealing process softens the copper, this re-polishing can leave microscopic pits from clumped swarf.

The same issue is repeated with gold plating – the gold also forms grain boundaries, which can be polished out; however, the gold is even softer than the annealed copper, so the pitting issue can be worse. Because the material is so soft, essentially no material is removed during this step; most of the polishing is accomplished via elastic deformation of the surface rather than abrasion. The pits are much smaller than the grain regions or boundaries, so we decided to perform this post-polishing rather than leave the grain boundaries.

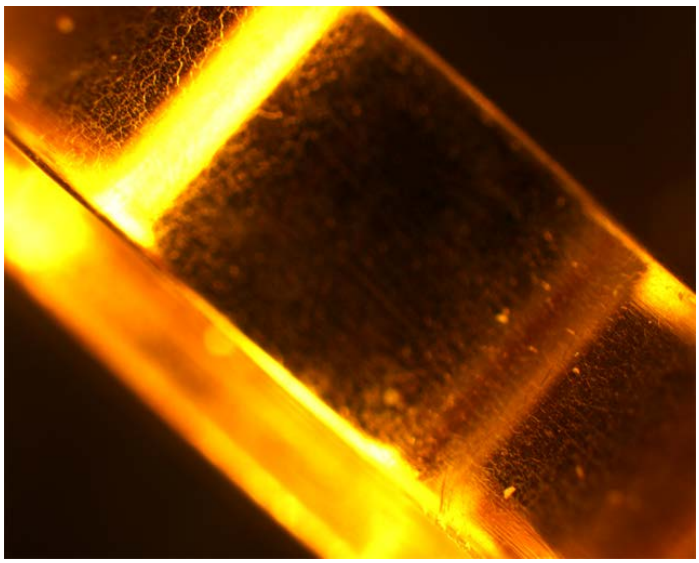

Figure 8.8: Compensation electrode test piece from the generation 1 trap, evaporatively plated, showing the "scale" pattern due to grain formation while brazing; to remove these features, the smallest grit paste was applied by hand to the electrodes after brazing and again after gold plating.

The electrodes are gold plated to improve conductivity and to prevent corrosion or oxidation. Standard electroplating procedures were rejected because of the use of a magnetic nickel barrier layer. Thermal evaporation [34], used in the first-generation proton/antiproton apparatus, gives a thinner, smoother gold layer, and is thought

to reduce the formation of grains in the gold. However, over time the thin gold layer loses adhesion or defuses into the copper, and we observed bare spots forming on evaporatively plated electrodes over the course of several years. We therefore opted to use nickelless electroplating for the new trap. Compared to evaporation, electroplating forms a stronger chemical bond between the gold and the copper, and grain formation can be counteracted by post-polishing. We gold plated the analysis and precision trap ring and compensation electrodes in our lab, while other electrodes were sent out to a professional plating company[6]. We plated the most significant ones ourselves because the grain boundaries formed with our (slower) procedure were smaller under the microscope than those plated at the company.

8.2 Trapcan, pinbase, and tripod

A precision measurement requires a high-quality Penning trap vacuum. Incidentally trapped ions disturb the trapping potentials; and collisions with background gas can cause particle loss from the trap, either kinematically or (if the trapped particle is an antiproton) through annihilation. The Penning trap is therefore kept in a separate vacuum container or "trapcan", inside the lower-quality bore vacuum. This trapcan is sealed with compressed indium wire and pumped out through an annealed copper tube, which is then pinched closed to form a cold-welded seal. At room temperature the can is pumped to 10^{-7} Torr or below. When cooled to 4.2K, residual gas cryo-pumps onto the cold surfaces, improving the vacuum to the XHV regime; a previous experiment measured $\sim 10^{-17}$ Torr [112]. Because helium is the only gas

[6]Hi-Tech Plating in Everett, MA

with a significant vapor pressure at 4.2K, and the lab atmosphere is relatively rich in helium (due to cryostat exhaust and occasional helium leak-checking), we flush the trapcan – pumping down to 10^{-2} Torr, then filling the volume with dry nitrogen – between 7 and 10 times, until on average <1 molecule of helium remains[7].

The trapcan consists of a CP3 titanium tube and two OFE copper flanges, bolted togehter and sealed with compressed indium wire. The copper flanges contain oven-brazed ceramic-to-metal feedthroughs, to couple electrical signals into the vacuum. Above the trap can is the "pinbase", a flange with 52 electrical feedthroughs which is taken as the experiment's common electrical ground. The pinbase is sandwiched between the trapcan and the "tripod" region, where most of the experiment's cryogenic electronics are located, including the amplifiers discussed in chapter 7 and the filters and drive lines discussed in section 8.4. The tripod consists of three OFHC copper posts which provide structural support for the trapcan, conduct heat out of the trap and electronics to the helium dewar, and ground the electronics to the pinbase.

Below the trapcan, the bottom flange includes more electrical feedthroughs, as well as a port for vacuum pumping, a port for a thin titanium window for antiproton access, and copper posts to support additional electronics and interface with the alignment system. Electronics space below the trapcan allows us to fit five resonator coils, to accommodate the cooling and loading traps. Additionally, the simplified trapcan vacuum includes the minimum possible four indium seals (which must be tightened each cooldown) – top, bottom, vacuum port, and antiproton window.

In the future, the titanium trapcan tube could be replaced with thick copper.

[7]Ignoring diffusion through rubber seals and backflow through the scroll pump, both of which will contribute a nonzero but effectively negligible level of helium contamination.

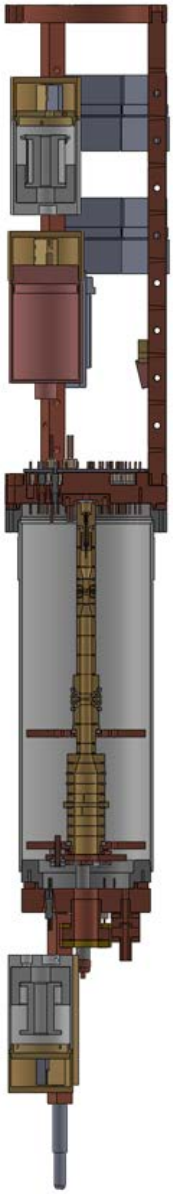

Figure 8.9: CAD model of tripod and trapcan with electronics and trap electrodes.

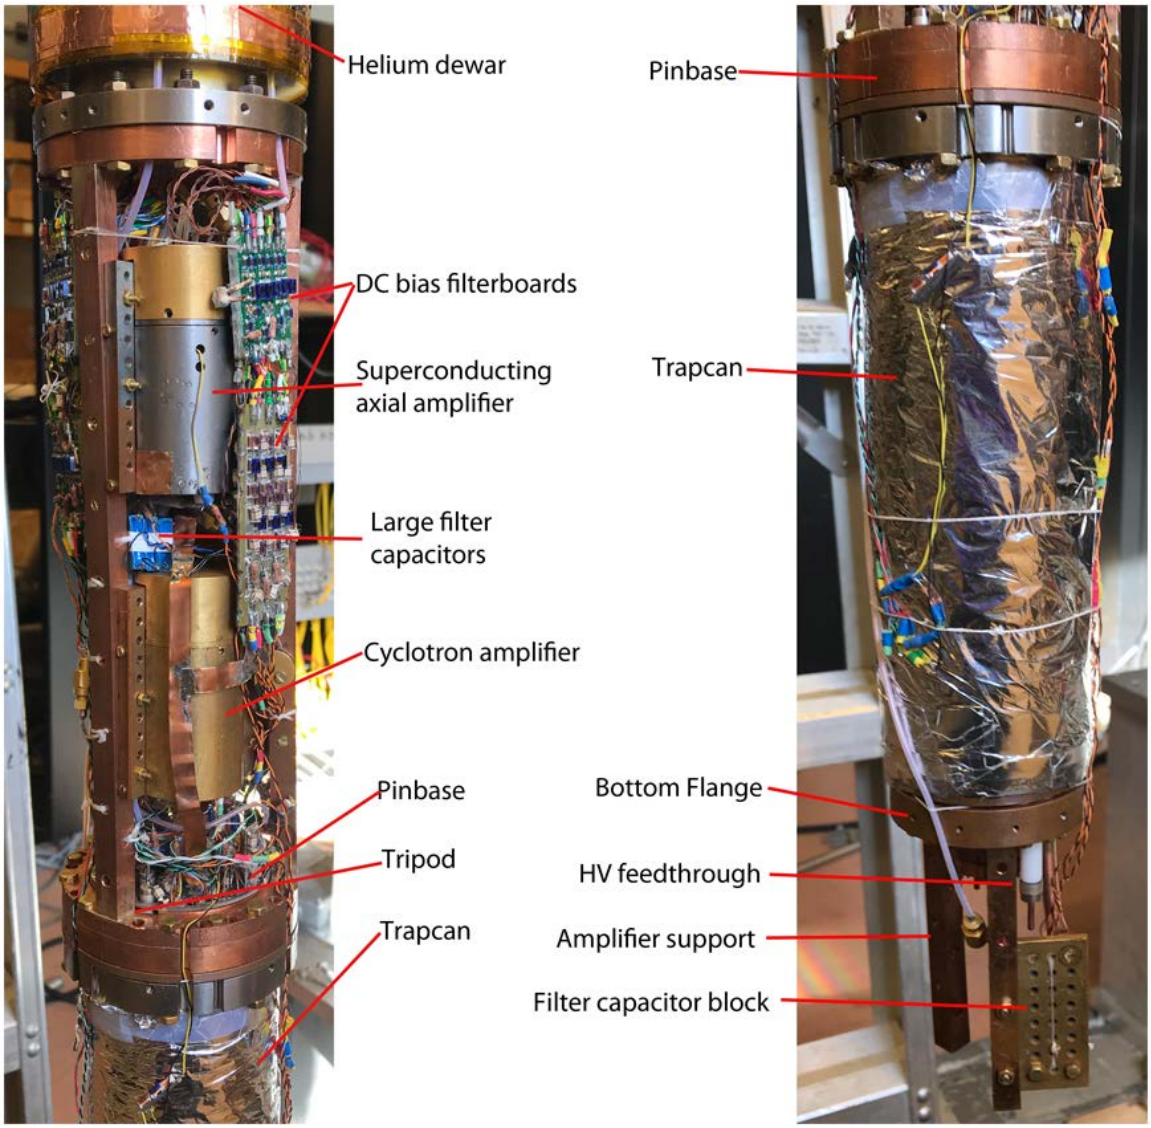

Figure 8.10: Left: Tripod and cryogenic electronics. Right: Trap can and secondary cryogenic electronics region.

Eddy currents in this highly conductive copper would help shield magnetic field noise, potentially increasing short-term field stability. The copper tube could be electron-beam welded to titanium indium-seal flanges, to avoid having soft copper on both sides of a large, tight seal (risking cold-welding or deforming the sealing surfaces).

8.3 Experiment cryostat and mechanical structure

8.3.1 Titanium cryostat

A new cryostat was constructed from titanium for the second-generation apparatus. A copper dewar was initially constructed, but the technique previously used at the machine shop to make simultaneous leak-tight seals was not replicated. Attempts to fix leaks in our hydrogen oven also failed. Switching to titanium allowed the joints to be electron-beam welded together rather than soldered, localizing heating to the beam spot and reducing joint-damaging thermal gradients.

We chose titanium because of experience electron-beam welding the material[8] for cryogenic vacuum joints, and succeeded in creating a vacuum-tight helium dewar. It should also be possible to electron-beam weld copper, despite the greater heat conduction. While we should continue to use electron-beam welding for new cryostats, if a new cryostat is required copper may be preferable. We summarize two concerns with the use of titanium, both of which are resolved in this apparatus but could be avoided in the future – thermal gradients and thermoacoustic oscillations[9].

[8] Welding performed by Joining Technologies Inc, East Granby, CT

[9] Note that over many years of use, titanium will deform less than copper; for new cryostat design, that should be balanced against the thermal considerations described in this section

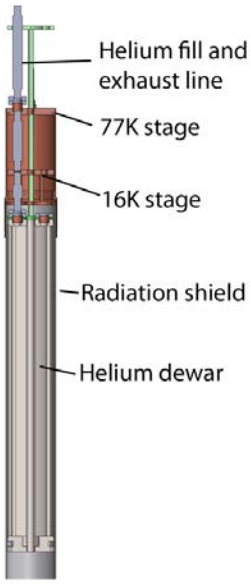

Figure 8.11: Titanium cryostat and thermal isolation stages

8.3.2 Thermal gradients in titanium

The thermal conductivity of titanium is significantly less than that of copper. This raises a possibility of temperature gradients in the experimental structure which change as the helium in the dewar boils. To establish whether this is a significant concern, we will calculate an upper bound on the size of this possible effect. We start with some conservative approximations – we will ignore the cooling power of the cold helium gas as it travels through the dewar and exhaust system, treating only conductive cooling through the titanium; and we will assume the entire heat load on the experiment is incident at the top of the dewar. We want to find the temperature gradient between the top and bottom of the helium dewar at the end of the hold time, when there is a small amount of liquid helium left in the tank. We therefore make the

further assumption that only the bottom of the dewar is held at 4.2K, and attempt to find the temperature at its top.

We can estimate the heat load on the experiment from the helium hold time. The 3.2 liter dewar takes over 27 hours to boil off, or .12 liters per hour. The cooling power of boiling helium is 1.38 liters per hour per Watt [113], giving a heat load on the helium dewar of 86 mW. The temperature gradient can then be found using

$$\dot{q} = \frac{A}{L}\int_{T_1}^{T_2} \Lambda(T)dT \tag{8.5}$$

where A and L are the cross-sectional area and length of the conducting material, $\Lambda(T)$ is the temperature-dependent thermal conductivity of the material, and T_1 and T_2 are the temperatures at either end of the thermal gradient.

The cryogenic thermal conductivity of titanium is approximately linear from 4.2K through 20K, with a value of $\sim .014 * T$ W/(cm K) [114, 115]. The total cross-sectional area of the titanium dewar tubes is 0.96 square inches, and the distance between flanges is 23.4 inches. Integrating equation 8.5 and setting the result equal to 86 mW gives a temperature estimate of 11K at the top of the dewar.

The temperature gradient across the dewar is only problematic insofar as it affects the particle or some aspect of the measurement. The trap and cryogenic electronics will be held at 4.2K by thermal conduction to the bottom flange of the dewar, but the elevated dewar-top temperature could affect the electronic temperature of a second-stage amplifier (see sec. 7.3). However, a second-stage amplifier at 11K should reduce room-temperature thermal noise nearly as much as one at 4K, especially since the physical temperature is only part of what determines the electronic temperature. If the difference between 4K and 11K is found to be significant, a new second stage

could be constructed and located in the tripod region.

A more substantial concern regarding temperature gradients is the possibility of thermal expansion. If the physical dimensions of the dewar change as the helium boils off, the particle's position would shift in the magnetic field. Fortunately, the thermal expansion coefficient of titanium at cryogenic temperatures is low. Specific data in the range 4-11K is hard to find, but [113] notes that titanium expands by a total of 0.001% between 4.2K and 40K. This would correspond to a shift of 0.00023 inches in particle position. A shift in position of 0.00023 inches would approximately[10] correspond to a fractional change in field strength of $1.1 * 10^{-10}$ over 27 hours, or 0.004 parts per billion per hour. We conclude that even this upper bound on thermal expansion gives negligible field drift, compared to the expected 0.1 to 0.5 PPB per hour due to pressure and other temperature changes.

8.3.3 Thermoacoustic oscillations

When we first cooled down the experiment with the titanium dewar, we observed periodic temperature spikes (shown in figure 8.12), together with degradation of the vacuum and sudden shifts in the Q and frequencies of the superconducting resonators.

This behavior can be attributed to thermoacoustic oscillations in the exhaust gas. Ordinarily, the boiling helium in the cryostat exhausts into the recovery system via a column of gas rising through the fill line (see figure 8.11). The temperature of the gas increases as it rises, cooling the bellows walls and counteracting the heat load from room temperature. Under certain conditions, as the gas warms and expands the

[10] Calculated by treating the magnet as a simple solenoid, neglecting the flattening effect of the shim coils.

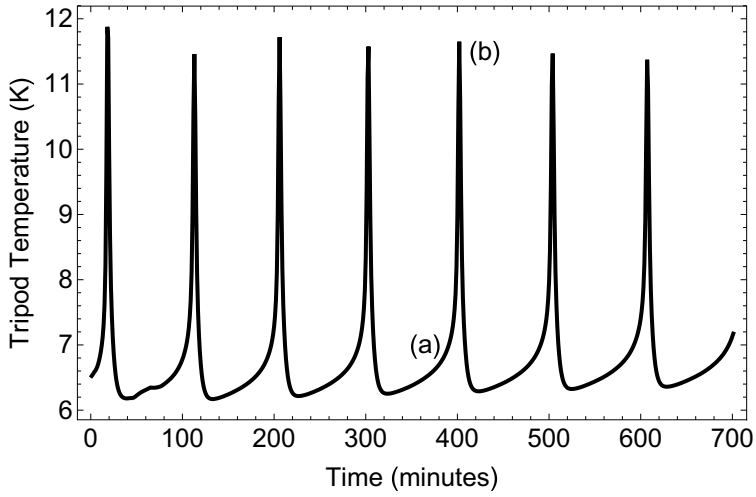

Figure 8.12: Periodic temperature spikes due to thermoacoustic oscillations in the helium dewar. During period (a), the temperature rises slowly as the fill line returns to thermal equilibrium after being cooled during a helium fill (or by a previous oscillation). During the spike, period (b), a positive feedback loop develops, creating thermoacoustic oscillations. The high boiloff rate from the oscillation breaks the feedback loop when it re-cools the fill line.

pressure at the top of the column can grow higher than the pressure of the boiling gas in the cryostat. When this happens, hot gas is pushed back down the fill line. When it reaches the reservoir, it boils off more helium, further increasing the pressure at the top, and so on. The warm and cold gas travel as sound waves in the fill tube – hence, thermoacoustic oscillations. This is the same principle underlying the "thumping" behavior in manual helium level sensors. These oscillations can occur whenever a relatively large cold volume exhausts through a narrow tube ending in a flow restriction. In our case, the cold volume is the helium tank, and the flow restriction comes from the re-entrant layered tubing and the connection to the recovery line.

The conditions for these oscillations depend on the length-to-diameter ratio of the tube and the temperature gradients of the tube walls [116, 117]. The temperature and

geometry of our titanium cryostat is very close to the boundary between stability and oscillation – see for example figure 4 of [116]. If the system passes between stability and instability depending on temperatures in the fill line and gas column, it would explain the periodic nature of the temperature spikes. The bellows walls are cooled by the increased exhaust during a helium fill, then warm up as the tube returns to thermal equilibrium. At some critical wall-temperature ratio, the gas column would become unstable, inducing oscillations; these raise the temperature of the helium dewar and increase the boiloff rate further, until the cold boiled-off gas alters the fill line temperatures back into the stable regime.

Across various experiments our lab has used at least six cryostats with similar design parameters, and until this cryostat all previous ones were stable. One potential cause would be the change in dewar materials – the lower thermal conductivity of titanium may have affected the thermal anchoring of the fill line, which could have pushed it into the unstable regime.

Once we understood the principle behind the oscillations, we were able to prevent them by changing the temperature ratios along the tube. To do this we added copper braid to the side of the tube and anchored it at the liquid nitrogen baffle, as shown in figure 8.13. This extended the length of tube anchored to 77K, and reduced the length of tube between 77K and room temperature . While this slightly increases the heat load of the experiment, by adjusting the height of the copper braid on the fill line we eliminated the oscillations while maintaining a 27 hour hold time. Even though this problem is solvable, it may be preferable to avoid it altogether by constructing the next cryostat out of electron-beam welded copper.

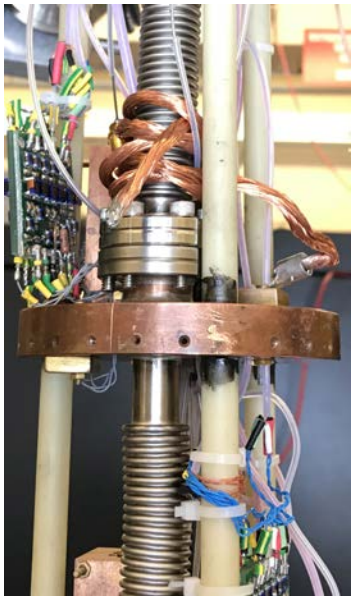

Figure 8.13: Copper braid partially anchoring the exhaust tube to the 77K stage.

8.3.4 Thermal isolation stage

Between the cryostat and the room-temperature "hat" of the experiment are several stages of thermal isolation, shown in schematic in figure 8.11 and pictured in figure 8.14. Structural support is provided by G10 tubes, which have high tensile strength and low thermal conductivity. The 77K stage is anchored to the liquid nitrogen-temperature magnet bore through a copper tube and compression fingers, and hosts filters for the temperature sensors. The 16K stage consists of two flanges which float at $\sim$ 16K and anchor the radiation shield, an aluminum tube wrapped with several layers of mylar superinsulation, which reflects thermal radiation from the magnet bore. This stage stage hosts cryogenic 50Ω attenuators, intended to reduce the strength of RF pickup and replace room temperature with 16K Johnson noise[11].

[11]Note that as of this thesis, the cryogenic attenuators are not yet in use on the drive lines.

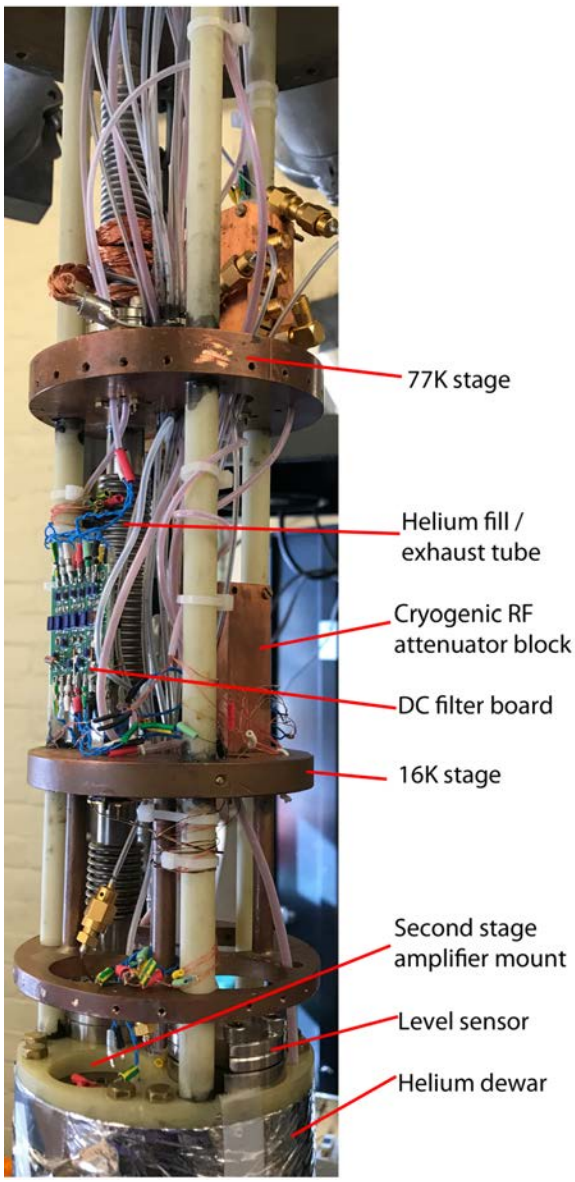

Figure 8.14: Thermal isolation stages of the experiment (77K cold fingers and 16K radiation shield removed).

At each end, the G10 tubes are glued into a G10 flange, bolted to either the hat or the helium dewar. The relative straightness of the hat, thermal insulation stage, and helium dewar are adjusted to $< .04$ degrees[12].

8.4 DC Wiring

The addition of the cooling and loading traps, new amplifiers, FET switch, and alignment sensor more than doubled the number of DC connections from the previous generation. DC voltages are carried from room temperature to the experiment on twisted pairs of 0.003" thick Constantan alloy wires. Figures 8.15 and 8.16 show the wiring diagrams for the trap electrodes. (Wiring diagrams for the resonators and FET circuits are shown in chapter 7).

The high and low sides of each DC signal pass through a low-pass filter. DC signals entering the trapcan also traverse an RC low-pass filter directly attached to each vacuum feedthrough pin.

8.5 Tuned circuit drive lines

Measuring the magnetic moment requires driving spin transitions, with RF fields at hundreds of megahertz. The AC behavior of structural and electrical components in this frequency range can differ significantly from their DC properties. The Penning trap is optimized for the DC properties of the electrostatic fields used to trap particles. This does not preclude the use of Penning trap electrodes as circuit components at

[12]To accomplish this, the assembly is rotated and measured with a micrometer, and thin brass shim-stock is added between stages as they are bolted together.

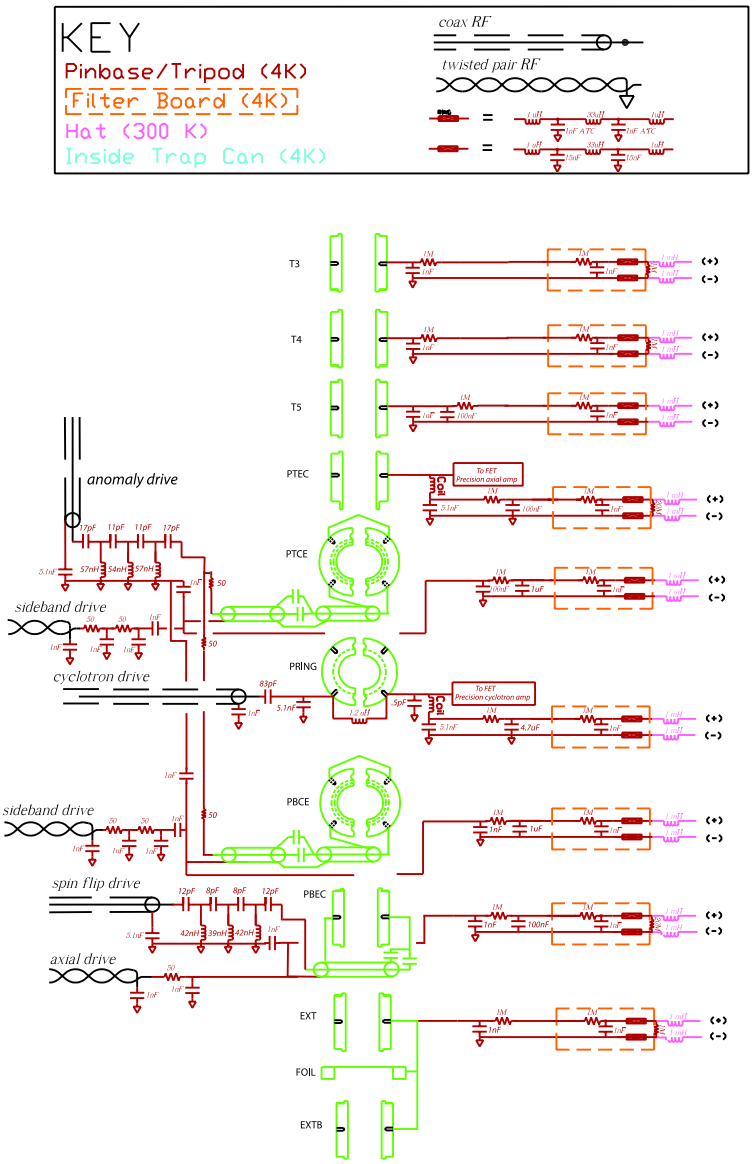

Figure 8.15: Wiring diagram for the precision trap.

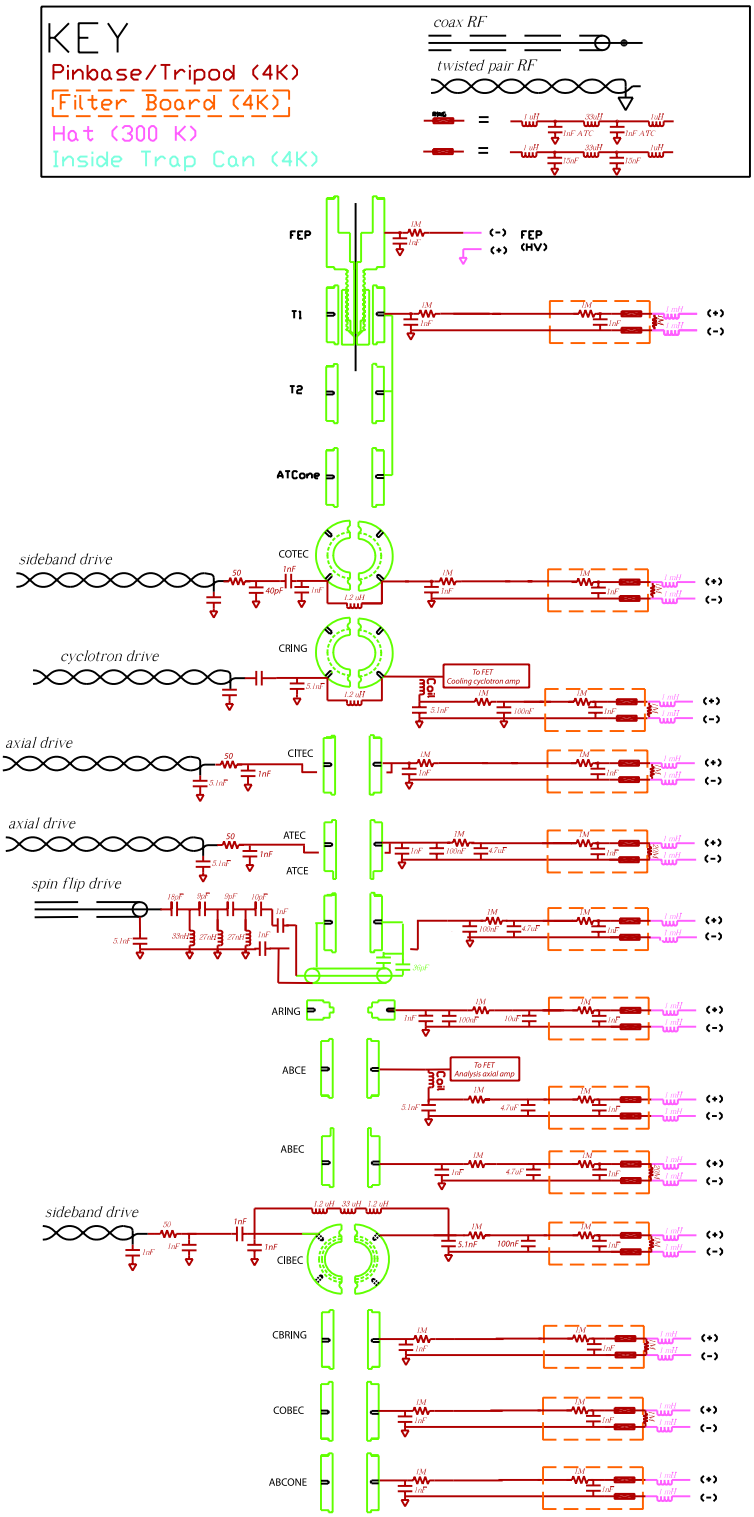

Figure 8.16: Wiring diagram for the analysis trap.

some frequencies – for example, as capacitances for detection resonators [46], or as the walls of a microwave cavity in the electron g-2 experiment [118, 119]. At the intermediate range of hundreds of megahertz for a proton or antiproton spin-flip drive, the electrodes' RF impedance makes coupling power to the particle challenging.

Different methods have been used in ion Penning trap experiments to drive RF transitions. In earlier Penning trap experiments with ions [47, 17], currents were applied through DC leads, driving transitions but reflecting or absorbing most RF power in the load and electrode impedances (and inductive couplings to other trap elements). Other Penning trap experiments [8, 120] use a disc antenna located near the trap. In this geometry, the RF field is coupled efficiently into the trap vacuum and penetrates into the trap through the slits between electrodes. However, eddy currents in the electrodes screen the field and cause heating in the electrodes [121].

For the antiproton magnetic moment measurement described in chapter 3, we used the electrode itself as our RF antenna. Making its inductance part of a low-Q resonator allowed us to efficiently couple current into the trap vacuum, and directly into the electrode, where its magnetic field drives transitions of the particle. Initially, we thought to use a half- or quarter-wavelength coaxial cable to impedance-match the drive line and minimize reflected power. However, impedance-matching to 50 Ω drive would maximize the total power input to the trapcan, rather than the current in the electrode (which sets the driving magnetic field seen by the particle).

Rather than 50 Ω matching, a controlled impedance mismatch can increase this drive current. Coaxial cable with a 10 Ω characteristic impedance[13] inside the trapcan

[13] 10Ω is the nominal characteristic impedance; measurements on several samples consistently yielded results closer to 9 Ω.

forms reflected waves at the interface between different impedances. This allows a standing wave to build up inside the cable, effectively creating a low-Q transmission-line resonator. Designing this so a current maximum of the standing wave travels through the electrode at the drive frequency gives a greater Rabi frequency than 50 Ω impedance matching.

This system was implemented in 2012-13 for the antiproton magnetic moment measurement [6], and improved in 2017 for the new apparatus. In this section, we first present the 2012 version of the drive line, where the load electrode was at the end of a transmission-line standing wave resonator. We then discuss the results from this drive line in terms of the improved drive current and reduced systematic shifts. Finally, we present the improved drive line resonator implemented in the new apparatus, where the electrode and a nonmagnetic lumped-element capacitor form a moderate-Q resonator, coupled through a capacitive divider to the 10Ω transmission line.

8.5.1 Transmission line resonator

We use transmission-line equations to calculate the steady-state current and voltage in the load electrode [122]

$$V_L = V_i^+ \frac{1+\Gamma_L}{1-\Gamma_L\Gamma_S e^{-2i\beta L}}, \qquad I_L = \frac{V_i^+}{z_0}\frac{1-\Gamma_L}{1-\Gamma_L\Gamma_S e^{-2i\beta L}} \tag{8.6}$$

where V_L and I_L are the voltage and current, respectively, through the load once the standing wave has built up to a steady state; Γ_L and Γ_S are the reflection coefficients, looking out from the transmission line at the load and source respectively; V_i^+ is the voltage entering the load, before building up to a steady state, as the initial voltage wave travels down the cable; β is the wavenumber of the RF signal; and z_0 and L are

the characteristic impedance and length of the transmission line. The transmission line parameters Γ and β are given by

$$\Gamma = \frac{z - z_0}{z + z_0}, \qquad \beta = \frac{2\pi}{\lambda} = \frac{2\pi f}{c_{cable}} \tag{8.7}$$

The initial voltage across the load, V_i^+, is given by

$$V_i^+ = V^+ e^{-i\beta L}, \quad V^+ = \frac{V}{1 + \Gamma_{in}}, \quad V = \frac{z_{in} V_s}{z_{in} + z_s} \tag{8.8}$$

where V^+ is the amplitude of the initial voltage wave entering the transmission line; V is the amplitude of the initial voltage wave heading towards the transmission line from the source; Γ_{in} and z_{in} are the input reflection coefficient and impedance, respectively, from the source into the line; and V_S is the open-circuit voltage from the source, disconnected from the load or transmission line. We can then rewrite equation 8.6:

$$V_L = \frac{\frac{z_{in} V_s}{z_{in} + z_s}}{1 + \Gamma_{in}} \frac{e^{-i\beta L}(1 + \Gamma_L)}{1 - \Gamma_L \Gamma_S e^{-2i\beta L}}, \quad I_L = \frac{\frac{z_{in} V_s}{z_{in} + z_s}}{z_0(1 + \Gamma_{in})} \frac{e^{-i\beta L}(1 - \Gamma_L)}{1 - \Gamma_L \Gamma_S e^{-2i\beta L}} \tag{8.9}$$

To use this equation, we need to calculate the reflection coefficients based on the source and load impedances. The circuit diagram for the transmission line resonator is shown in figure 8.17. The source impedance includes the 50 Ω coaxial drive line from room temperature; we also treat the impedance of the vacuum feedthroughs as part of the source, for the purpose of evaluating the standing wave in the 10 Ω cable. Measured impedances of the feedthroughs are approximately 12 nH on either side of a 3 pF capacitance to ground. We can then add the source components of figure 8.17 using series and parallel rules to get z_{in}.

On the load side, each electrode shape gives a slightly different inductance. For an analysis trap compensation electrode (the load for the spin-flip drive line implemented

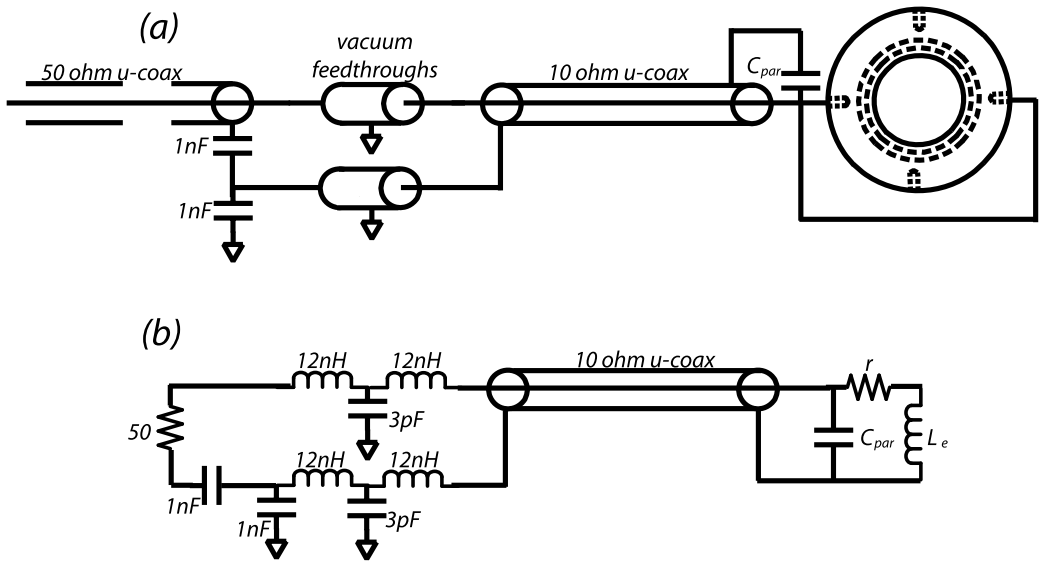

Figure 8.17: Circuit diagram for the tuned circuit drive line. (a): drive line circuit as implemented. (b): equivalent circuit elements used to calculate drive currents.

in 2012), we measure an inductance of around 18 nH. There is also capacitive coupling on the order of a few pF to neighboring electrodes; however, those electrodes are only connected to ground through long copper straps, with inductances of several hundred nanohenries. At the spin-flip and anomaly frequencies those straps have a high impedance and very little current flows through them. We add two more elements to our model of the load inductance – we allow for a nonzero RF loss at the electrode, r, due to magnetoresistance, dielectric losses, and other miscellaneous effects, estimated at between 0.03 and 0.3 Ω; and we allow the addition of a small parallel capacitance C_{par} to tune transmission-line resonator frequency without physically modifying the 10 Ω cable length. The expressions for load voltage and current in equation 8.9 are for the total signal across the load; to account for the flow of current between the

parallel capacitor and the electrode, we take the current only through the electrode[14],

$$I_{electrode} = \frac{V_L}{Z_{electrode} + r} \tag{8.10}$$

With these impedances, setting $V_S = 1V$ and $c_{cable} = 0.7c$[15], we can use equation 8.7 to calculate the reflection coefficients, and equations 8.9 and 8.10 to calculate the drive current through the electrode as a function of signal frequency, tuning capacitance, electrode inductance, cable length, and resonator loss. Results of these calculations are shown in figure 8.18.

8.5.2 2012-13 results with transmission line resonator

In the 2012 proton and 2013 antiproton measurements, we used a spin-flip drive to induce transitions in the analysis trap. The probability of flipping the spin, with a drive Rabi frequency Ω_r and linewidth $\Delta\omega_z$ is [1]

$$P_{sf} = \frac{1}{2}\left(1 - e^{-\frac{\pi\Omega_R^2 T}{\Delta\omega_s}}\right) \tag{8.11}$$

The magnetic bottle induces a large linewidth $\Delta\omega_z \sim 2\pi * 25$ kHz in the analysis trap, so inducing transitions requires a large drive power. In the 2012 proton measurement [17], we observed a significant drive-strength-dependent shift in the axial frequency while applying the spin-flip drive. RF power was carried to the trap on inductive copper straps, leading to high power dissipation in the electrode and straps carrying the drive. The unshielded straps also coupled inductively to other trapcan elements, dissipating more power. Because the circuit was not optimized for current coupling,

[14]This expression treats the electrode as a lumped-element inductor, since its length is much less than the wavelength of the signal.

[15]Manufacturer's specification for propagation speed of EM waves in the cable

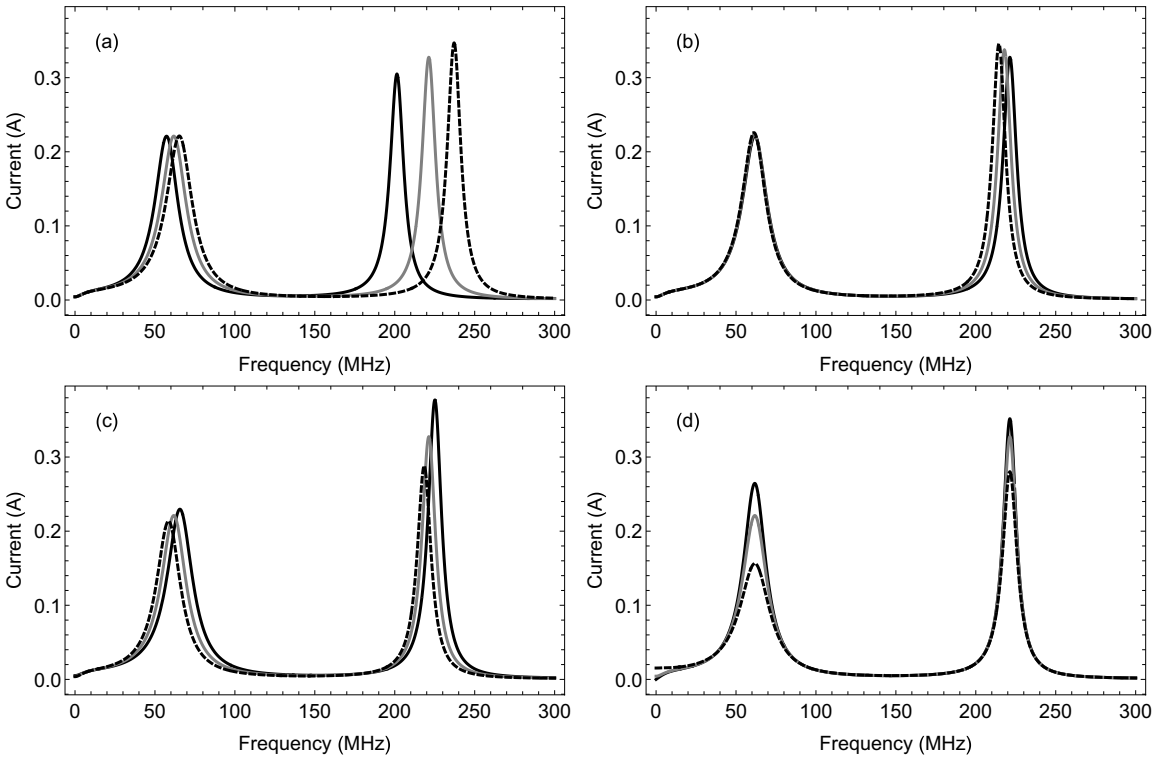

Figure 8.18: Current profiles for different transmission line parameters; open-circuti input voltage assumed to be 1V peak-to-peak.
(a): different transmission line lengths – black .6 m, gray .54 m, dashed .5 m
(b): different tuning capacitances - black 0 pF, gray 5 pF, dashed 10 pF
(c): different electrode inductances - black 15 nH, gray 18 nH, dashed 21 nH
(d): different resonator loss values – black 0 Ω, gray 0.3 Ω, dashed 1 Ω

RF voltages on various electrodes could also shift the effective trapping potential. Additionally, we were unable to fully saturate the spin-flip transition.

For these reasons, we implemented a tuned-circuit coaxial drive line before beginning the 2012-2013 antiproton magnetic moment measurement. This significantly reduced the total power dissipated in the trap by reducing net inductance, controlling induced currents, and improving coupling efficiency. Figure 8.19 shows the axial frequency shift before and after implementation of the transmission line resonator.

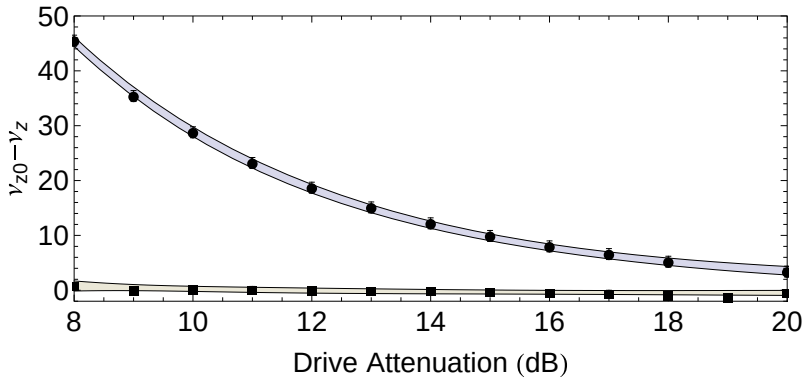

Figure 8.19: Axial frequency shift versus spin flip drive attenuation in the 2012 proton (upper line) and 2013 antiproton (lower line) measurements. Confidence intervals in exponential fits are shown as shaded bands. Data for the antiproton measurement was taken with 12dB drive attenuation, consistent with zero axial frequency shift.

In addition to reducing or eliminating this power shift, the coaxial resonator also significantly increased the driving field at the particle per power input at the hat. We can calculate the current from measured transition rates with and without the new drive line, using equation 8.11. In the 2011 proton measurement, a 4 second drive time with drive strength 24 dBm (2 dB attenuation in figure 8.19) yielded a 20% spin-flip probability on resonance; equation 8.11 gives a Rabi frequency of 80 rad/s.

In the 2013 antiproton measurement, the transition was saturated with a 2 second drive time and input drive strength of +14 dBm (12 dB attenuation in figure 8.19). As the transition rate approaches 50%, Ω_r of equation 8.11 approaches infinity, meaning data points near 50% will yield an unreliable Rabi frequency estimate. Instead, we calculate Ω_r at each point on the lineshape, using equations 5.27 and 6.24 of [1]. Averaging over all nonzero measured transition rates in the antiproton lineshape (figure 3.9) yields a Rabi frequency of approximately 410 rad/s. The ratio of Rabi frequencies of $\frac{410}{80} \approx 5$ gives the increase in electrode current between the two measurements.

To compare the two drive systems, we account for both the 10 dB reduction in input power and this increased drive current (corresponding to a factor of 25 in power, or an increase of 14 dB). In total, the spin-flip transmission-line resonator thus increased the power delivered to the electrode by 24 dB for equal input power, while reducing the axial frequency shift to a level consistent with zero.

8.5.3 2017 improvements – LC drive resonator

For the new apparatus reported in this thesis, we further improved the drive circuit. The driving circuit is modified as shown in figure 8.20. The electrode forms as the inductor in a medium-Q parallel LC resonator with the lumped-element capacitor C_{par}. A second capacitor C_{ser} forms a capacitive divider, partially decoupling the LC resonator from the source impedance (otherwise, Q would be limited by the 50Ω source). Current is then coupled into this LC resonator through a transmission line resonator. For the precision trap, a short coax separates the electrode physically

from the two capacitors, to reduce the effect of capacitor materials on magnetic field homogeneity. The circuit diagram for the new drive line is shown in figure 8.20.

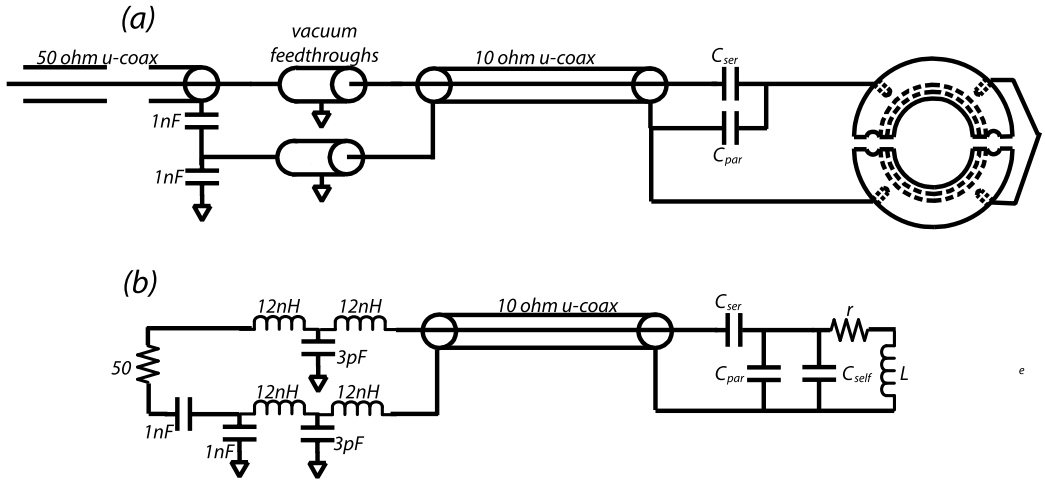

Figure 8.20: Circuit diagram for the LC resonator drive line. (a): drive line circuit as implemented. (b): equivalent circuit elements used to calculate drive currents.

The calculation for this circuit largely follows section 8.5.1. Equation 8.9 is still valid for current and voltage across the entire load; however, the load impedance Z_L is modified by the addition of the series capacitor and the change in the (previously small) parallel tuning capacitor. Additionally, the geometry required to drive anomaly transitions (section 2.7) connects two halves of a split electrode at one end, leaving a capacitance between the two halves at the other end, which we will label C_{self}. Finally, we modify equation 8.10 to isolate current through the electrode, giving

$$I_{electrode} = \frac{V_L - Z_{c_{series}} I_L}{Z_{electrode} + r} \tag{8.12}$$

Figure 8.21 shows the result of this calculation, giving current through the electrode for a 1 V open-circuit amplitude at the source. We see two peaks in the LC

resonator profile – a relatively broad transmission-line resonance and a narrower LC resonance. Current in the electrode is not maximized when these resonances overlap; impedance-matching into the coaxial resonator would load down the LC resonator.

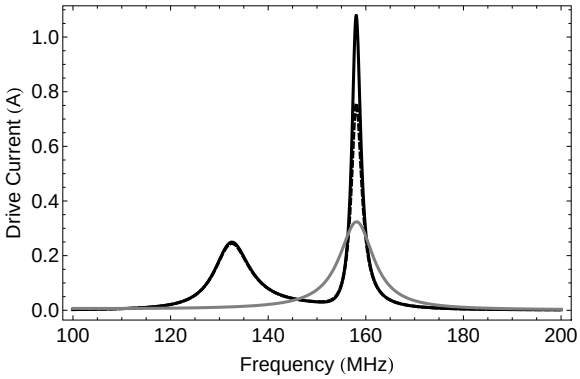

Figure 8.21: Current profile for the parallel LC drive resonator at $\nu_s - \nu_c \approx$ 158 MHz. Grey line is the transmission line resonator of section 8.5.1. The improvement from the LC circuit depends on the estimated loss in the resonator; 0.03 Ω gives the solid black line, while 0.1 Ω gives the dashed line.

Drive strength improvement from the LC resonator

Figure 8.21 demonstrates that the improvement over the transmission line resonator depends on the loss in the LC resonator. The best way to evaluate this would be to drive spin-flips and measure the Rabi frequency, but that would also be very time-consuming. The drive current can be measured with a network analyzer; current is supplied to the drive line and signal is received with a pickup coil held near the electrode. This gives the frequency of the resonance, but the measured amplitude depends strongly on the coupling to the pickup coil, so without extensive normalization we cannot compare these measurements directly to calculated amplitudes.

That said, we can estimate the loss in the resonator by observing the change in

resonance amplitude as we shift one of the circuit parameters. Figure 8.22 shows measured and calculated current on resonance, with a set electrode inductance, cable length, and capacitance. Calculated values are normalized to the measured value. The curvature and peak location of the measured data indicate the loss in the resonator to be 0.015-0.05 Ω. This implies the LC resonator drive line should deliver a factor of 3-4 more current than the transmission line resonator alone (9-12 dB more power).

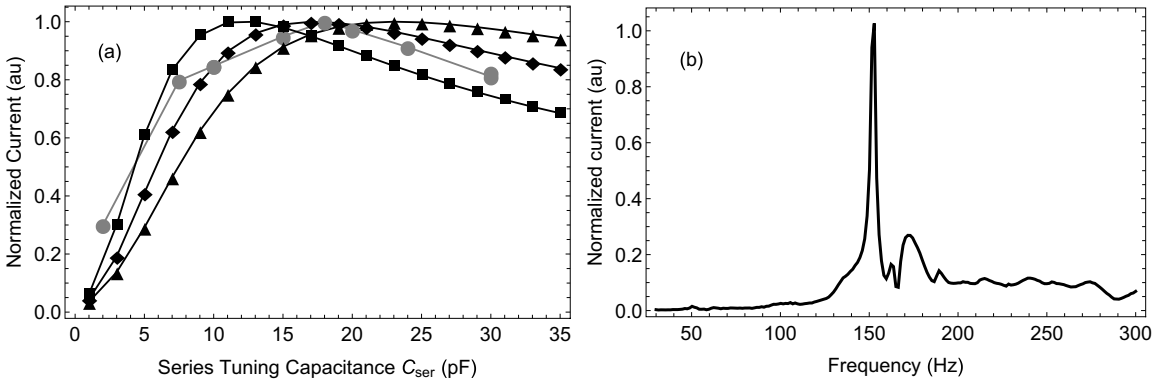

Figure 8.22: Empirical tests of the LC resonator drive line.
(a): Resonant current amplitudes versus series capacitances, with other parameters held constant. Gray circles are measured values on a test setup; black points are calculated values with different estimated resonator losses. Square: 0.015 Ω; diamond: 0.03 Ω; triangle: 0.05 Ω.
(b): Network analyzer scan for one set of parameters out of the data in (a).

Tuning the LC drive resonator

The resonance profile for this drive coupling is narrower than for the simple transmission line, and there are more component values to tune. Figure 8.23 shows the effect on the resonator's current profile of changing the different components.

This circuit has four tuning parameters. The values of the series and parallel capacitances and the coaxial cable length can be directly adjusted. Additionally,

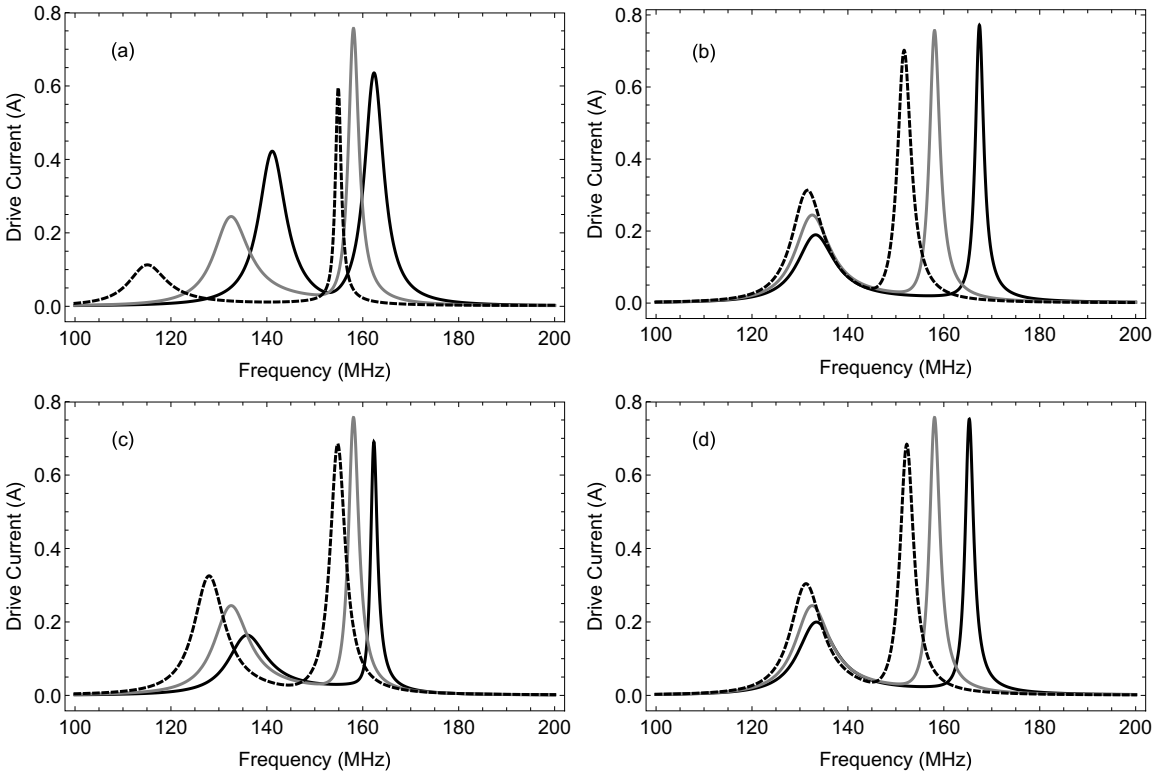

Figure 8.23: Current profiles for different LC drive resonator parameters.
(a): different transmission line lengths – black .7 m, gray .77 m, dashed .9 m
(b): different parallel capacitances - black 30 pF, gray 36.5 pF, dashed 42 pF
(c): different series capacitances - black 10 pF, gray 14.5 pF, dashed 20 pF
(d): different electrode inductances – black 18 nH, gray 20 nH, dashed 22 nH

the effective electrode inductance can be slightly adjusted by changing where on the copper leads we solder the coax shield and center conductor. With this large of a parameter space, we would like an algorithm for practical optimization of the circuit.

The effective electrode inductance can be changed only by a small amount, and for initial tuning we will treat it as set. At a given electrode inductance, a pair of capacitances can always be found for any cable length to give the same, optimum current value at the drive frequency, as shown in figure 8.24.

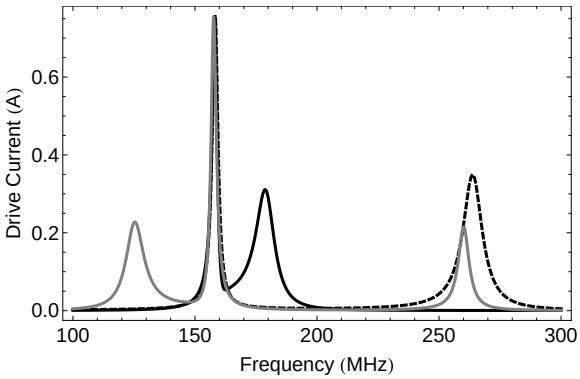

Figure 8.24: Current profiles with different cable lengths. Capacitor values are numerically optimized for a set cable length and electrode inductance.
Black line: 60 cm coax, 37 pF series, 8.7 pF parallel.
Gray line: 40 cm coax, 20 pF series, 26 pF parallel.
Dashed line: 80 cm coax, 32.6 pF series, 19.3 pF parallel.

We now propose a procedure for tuning the drive resonator. First, measure the electrode inductance (or estimate it based on geometry). Then, choose a length of 10Ω coax[16] based on the geometry of the electrode stack – preferably, a couple of inches longer than required to reach the electrode from the pinbase. With the chosen cable length, optimize equations 8.9 and 8.12 to select initial capacitor values. Using a network analyzer, adjust the series capacitance C_{ser} to maximize the current on

[16] again, nominal - the actual characteristic impedance should be measured before use.

resonance – as figure 8.23 shows, the series capacitance has a larger effect on the amplitude than the parallel capacitance. Once the series capacitance is chosen to maximize the amplitude, the frequency can be coarsely tuned by adjusting the parallel capacitance C_{par}. Finally, the frequency can be fine-tuned by changing the position where the coax is soldered to the electrode, which adjusts its effective inductance.

8.6 Cryogenic alignment system

For the new experiment, we have constructed a cryogenic alignment and support system. A gearbox, operating at liquid nitrogen temperature, shifts the bottom of the apparatus in the bore, thus tilting the experiment and changing the alignment of the trap electrodes. An alignment sensor, which detects the angle of an electron beam using a resistive position sensor, has been designed and is currently being tested for use together with this gearbox. This system serves two purposes in the experiment. By anchoring the bottom of the experiment with respect to the bore, we reduce the amplitude of vibrations which could shift the trap position relative to the magnetic field. Additionally, the gearbox and sensor give us in situ measurement and control of the alignment between the trap electric and magnetic fields, with implications for experiment precision, antiproton trapping efficiency, and particle transfer performance.

8.6.1 Cryogenic gearbox

Inspired by work performed at the University of Washington and reported in [123] we have implemented a cryogenic gearbox for experiment alignment. Our gearbox differs significantly in design and operation from that reported in [123], but shares the

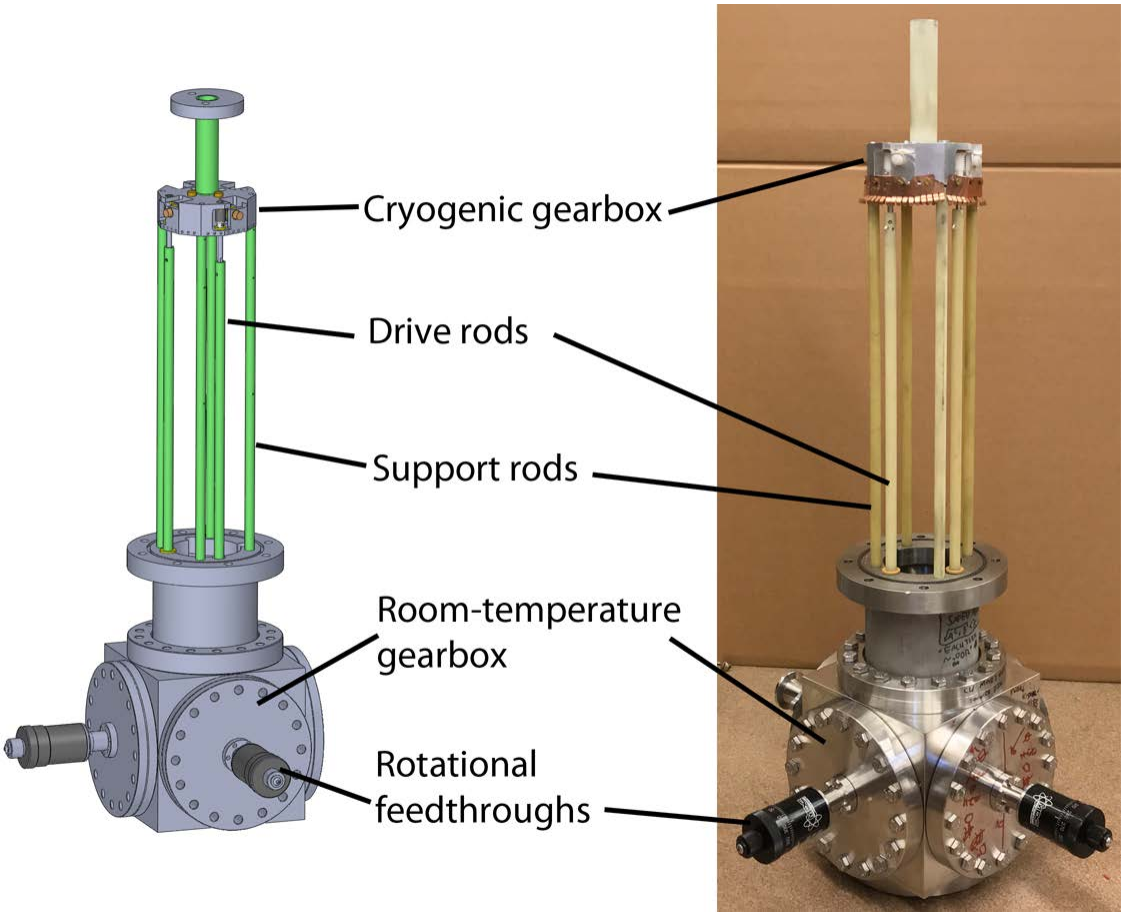

Figure 8.25: Cryogenic gearbox assembly.

underlying principle of using cryogenic worm gears to adjust the alignment between Penning trap electrodes and the homogeneous magnetic field.

Our gearbox is supported from underneath by four long, semi-flexible G10 support rods. Room-temperature rotational feedthroughs control angle gears, which turn a pair of drive rods. These rods in turn control a pair of nylon worms housed inside the gearbox. Each worm turns an acetyl worm gear around an axis perpendicular to the magnet bore. Worms and worm gears are held in place by thrust bearings made of Rulon-J, a glass-filled Teflon material which maintains a low coefficient of friction at cryogenic temperatures.

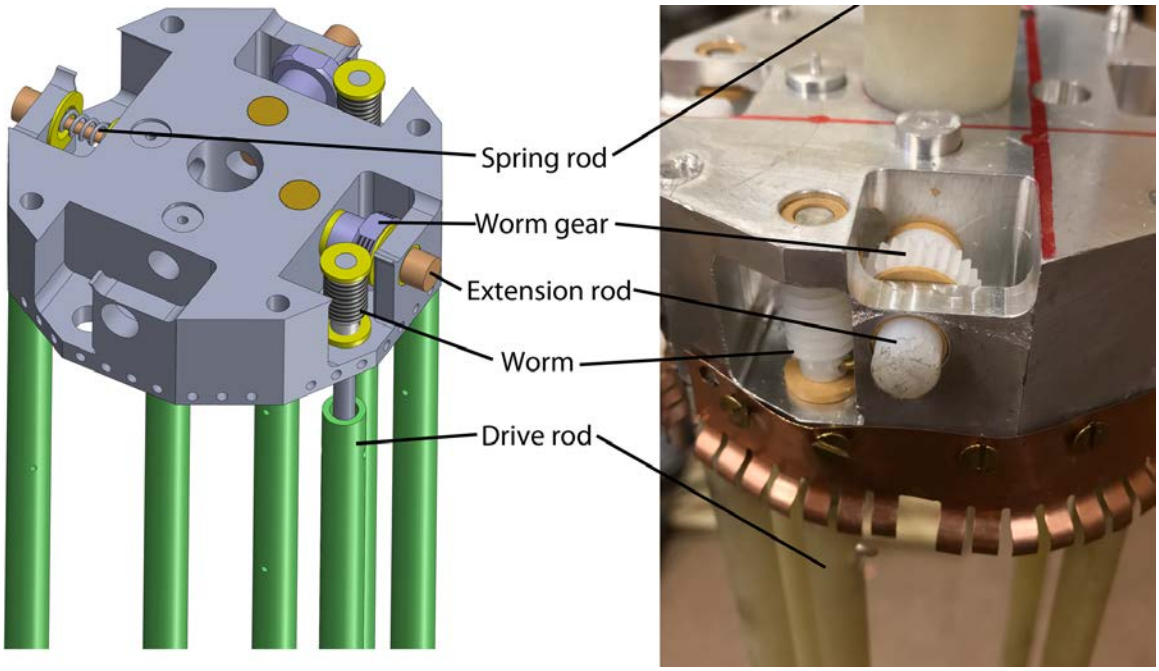

Figure 8.26: Diagram and picture of the cryogenic gearbox featuring worm, worm gear and thrust bearings.

These worm gears are threaded along the central axis, and mate with a threaded nylon rod, which is held in a fixed orientation by a pin. When the drive rod turns,

the worm rotates the worm gear, which extends or retracts the threaded rod. An additional nylon rod is extended by a beryllium copper spring opposite each threaded rod. The threaded and spring-loaded nylon rods press against the magnet bore. When a threaded rod is extended, the opposing spring contracts and the entire assembly moves in the bore (flexing the G10 support rods). The worms have two leads each to the 30 teeth of the worm gear, meaning a ratio of 15:1. The worm gear axis is tapped with 28 threads per inch, giving a total extension or retraction of .0024 inches per turn of the drive rod.

A rigid G10 tube is held in place by thermal contraction of the surrounding aluminum, extending upwards towards the experiment. PEEK rods hold an aluminum plate at the bottom of the experiment, which slides onto this G10 tube when the experiment is inserted into the bore. When the gearbox is translated by the threaded rods, the bottom flange of the experiment is also translated by the G10 tube and aluminum plate. Figure 8.27 shows this assembly.

The G10 rods supporting the experiment through the thermal isolation stages (figure 8.14) allow sufficient flexion to make these small adjustments. Because the hat is bolted in place, a horizontal translation of the bottom of the experiment yields an angular adjustment of the trap electrode alignment relative to the magnetic field.

This method allows only small angular adjustments, at most 0.5-1 degree. However, the experiment dewar shimming discussed in section 8.3.4 gives an initial alignment of the experiment structure to better than that level. Additionally, larger misalignments can be corrected while warm by addition of shim stock at the tripod-dewar interface. The in situ cryogenic alignment is required to consistently correct for small,

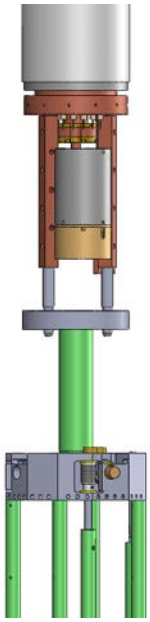

Figure 8.27: Assembly for adjusting the position of the experiment bottom flange with the gearbox.

uncontrolled factors like compression of cold fingers or contraction of support materials, which otherwise give inconsistent angular offsets that change with each cooldown.

The cryogenic gearbox operation at liquid nitrogen temperatures has been demonstrated, first in a nitrogen-filled test dewar and later in conjunction with the full experimental apparatus. Protons are loaded to the trap by excitation of adsorbed gas using a field emission point electron beam. With enough time and current, the field emission point will entirely clean off the target spot, resulting in reduced loading rates. When this was observed during the initial run of the experiment at Harvard, adjustment of the alignment using the cryogenic gearbox successfully shifted the target point and refreshed proton loading.

8.6.2 FEP-resistive anode alignment sensor

The cryogenic gearbox allows in situ adjustment of the alignment with small, controlled translations of the trap. However, a full alignment system should not only adjust the alignment, but also be able to repeatably return to a known position. Thermal contraction and inelastic compression of structural elements like G10 rods, beryllium copper compression fingers, and gear teeth mean that no purely mechanical alignment system can avoid some amount of hysteresis. Before the gearbox was implemented, the hat angle would be adjusted using differentially threaded screws; hysteresis came from the compression of rods and cold fingers, which did not return to the same position when an alignment adjustment was reversed. With the cryogenic gearbox, there is hysteresis because the gear teeth and axial threads contract on cooldown, meaning rotation in different directions cannot be perfectly transmitted.

Therefore, a full alignment system requires not only a mechanical adjustment, but also some way to measure the alignment after adjusting it. Additionally, an ideal alignment system would be able to measure the absolute alignment between the trap electrodes and magnetic field. The angle between electric and magnetic fields can be measured using a particle, but doing so is a slow process and only returns the amount of misalignment, not its direction.

We therefore have developed a system for quick measurement of the angle between the trap mechanical structure and the magnetic field. This relies on the fact that charged particles follow magnetic field lines. An electron beam, fired from a field emission point, will travel parallel to the magnetic field. Our alignment sensor consists of a field emission point together with a charged particle sensor that encodes the

impact position of the electron beam. This sensor is insensitive to the value of the magnetic field, and is shielded against stray electric fields – it thus directly measures the direction of the magnetic field between the FEP and the target.

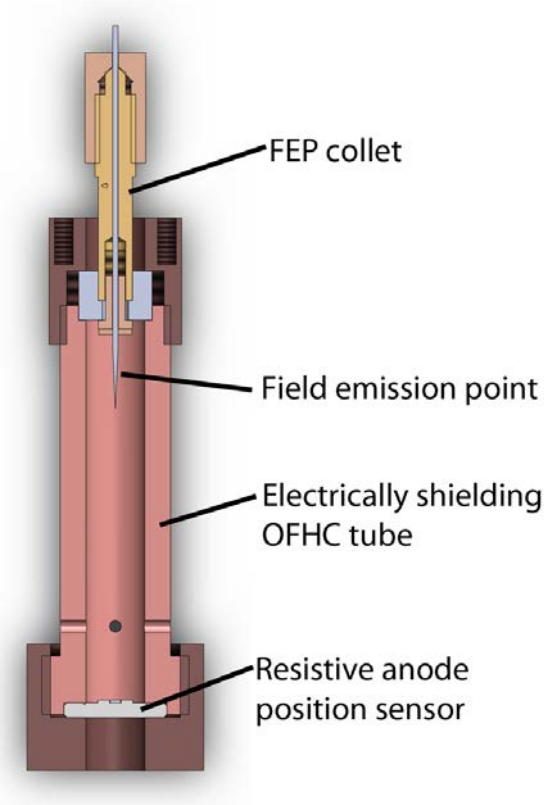

Figure 8.28: FEP-resistive anode alignment sensor. Field emission point and position sensor are held centered by thermal contraction of copper registers around Macor substrate and spacers.

For the target, we use a resistive anode position encoder [124]. The charged particles hit a sheet of resistive material surrounded by electrodes. The current through each electrode is measured; the ratios of these currents gives the position of the particles. Many technologies for position-sensitive charged-particle detectors exist, such as microchannel plates with delay line anodes [125] or silicon position sensitive detectors

[126]; however, most such technologies either don't work at cryogenic temperatures (ie, standard semiconductor-based technologies), or else require optical access (ie, scintillation or phosphorescence-based detectors).

To work at cryogenic temperatures, we had to make some adjustments to a typical resistive-anode position encoder. A semiconductive material is often used for the target, as its resistivity can be adjusted by controlled doping. However, since most charge-carriers freeze out at 4K, we instead use a thin layer of carbon-impregnated resistive Kapton[17] attached with epoxy to a Macor substrate. Electrodes are printed with photolithography onto the Kapton. Additionally, the current into each electrode is typically read out using low-impedance operational amplifier based circuitry [127]. We instead attach a grounding resistor to each pad, through which the current from each electrode flows, and measure the voltage across each resistor. Measuring a voltage across a twisted pair of wires is less sensitive to noise and thermal gradients than using room-temperature circuits to measure a small cryogenic current.

Once the currents in all four electrodes are measured, an algorithm must be applied to determine beam spot position. For the current ratios to linearly correspond to position, the electrodes would have to be hyperbolic and subject to particular ratios of resistivity with the substrate of the pad [128]. However, the boundary conditions for a rectangular resistive anode sensor with electrodes along the sides are analytically solvable, and an effective position estimator can be found using Taylor series – we use the estimator presented in reference [129].

Using grounding resistors to convert the output current to a voltage, rather than measuring current with a low-impedance op-amp, also distorts the position estimate.

[17] DuPont Kapton film 100XC

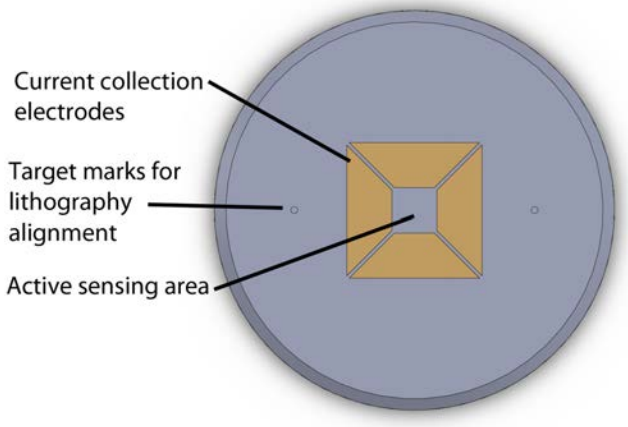

Figure 8.29: Resistive anode position sensor.

As long as the resistor values are small compared to the resistance per square of the Kapton film, this distortion should be small. However, the resistance must be large enough that the induced voltage can be read out over thermal noise in the signal lines. Measurement of the distortion added by this nonzero resistance is ongoing – a prototype sensor has been fabricated and is currently undergoing cryogenic testing in the new Gabrielse lab at Northwestern [130]. Once tested, the device will be attached to the support plate for the trap electrode stack, guaranteeing alignment between the sensor and Penning trap electric field. Two "dummy" sensors have also been fabricated to balance the extra mass due to the sensor and avoid torquing the electrodes. Once the detector is tested and attached to the experiment, together with the cryogenic gearbox it will form a reliable, stable in situ system to control and guarantee alignment between the trap and magnetic field.

Chapter 9

Results and Status of the New Apparatus

The apparatus described in chapter 8 was completed and prepared for cooldown in 2018. In this chapter, we will summarize the initial commissioning results on the apparatus. This work took place over the course of several months between the initial cooldown and the Gabrielse lab's move to Northwestern, culminating in the trapping and characterization of signal from single protons in the precision trap.

9.1 Experiment cryostat

Initial cooldown cycles dealt with the issue of thermoacoustic oscillations, as described in section 8.3.3. After these oscillations were dealt with, we obtained a liquid helium hold time of ∼27 hours. For a helium tank capacity of 3.2 liters, this corresponds to a heat load of 86 mW. This is comparable to the previous-generation

apparatus, where hold times over the last few years ranged from 12 to 30 hours. This confirms that the additional DC and RF connections for the new traps (section 8.4) and the PEEK/G10 connection from the trap can to the alignment gearbox (section 8.6) do not significantly increase the heat load.

Progress during the first few months was hindered not only by the thermoacoustic oscillations, but also by teething troubles with a new apparatus. These included a couple of vacuum leaks (one in a commercial feedthrough flange for the liquid helium level sensor, and one in the re-entrant helium fill port in the experiment hat), electrical shorts on filterboards, and a dull field emission point. During this period, the experiment was cooled down a total of 11 times, with an average duration of two weeks. The final cooldown lasted two months, and was ended so that the experiment and magnet could be shipped to Northwestern.

9.2 Particle loading and alignment

We load particles by firing the field emission point located at the top of the electrode stack while applying a trapping potential to the precision trap. For the commissioning experiments described in this thesis, the loading trap was blocked off by an aluminum foil, which formed the bottom of the trap stack. It was felt that commissioning the loading trap would be more efficient with a well-characterized precision trap already in place, and until the experiment is sent to CERN the loading trap is not required for optimal performance.

When the electron beam hits this foil, it releases adsorbed gas, including water and hydrocarbon molecules. These spread in the vacuum; some were ionized on interaction

with the electron beam. Some of these molecules are ionized by further interaction with the electron beam, and atoms and molecules ionized in the potential well of the precision trap are caught. They can then be resistively cooled by interaction with the resonator if resonant. Successful loading is identified by increased signal in the axial amplifier during this resistive cooling. We can then detect the signal from different ion species, or eject higher-mass ions from the trap to work with only protons.

As the electron beam is fired many times, the spot where it impacts the target foil becomes cleared of adsorbed gas. Temperature cycling the experiment to liquid nitrogen, raising the vapor pressure of some gasses, repopulates the foil with proton-bearing molecules. However, it is much quicker to make small adjustments to the alignment of the experiment in the magnetic field – since the electrons orbit magnetic field lines as they pass through the trap, adjusting the alignment changes the spot where the beam hits the foil. We used this as a signal to demonstrate that the cryogenic gearbox did successfully change the trap alignment. When the target spot became depleted, we were able to load protons again by adjusting the gearbox. This is essentially a binary yes/no signal for adjustment; either an operational alignment-sensor prototype or a careful measurement of anharmonicity with a single proton would be required to fully characterize the gearbox's performance. However, we can confirm cryogenic mechanical adjustment of the experiment alignment.

9.3 Axial signals in the precision trap

We first observed the axial motion of large proton and ion clouds using sideband heating and cooling drives (section 2.6). We used this technique rather than directly

driving the axial motion because the filtering and attenuation added to reduce Johnson noise coupled strong drive signals directly into the detector through the ground plane; this was not an issue for the sideband drive, because we drive and detect at different frequencies. Figure 9.1 shows the response to a sideband cooling drive sweep after loading particles for ions and protons.

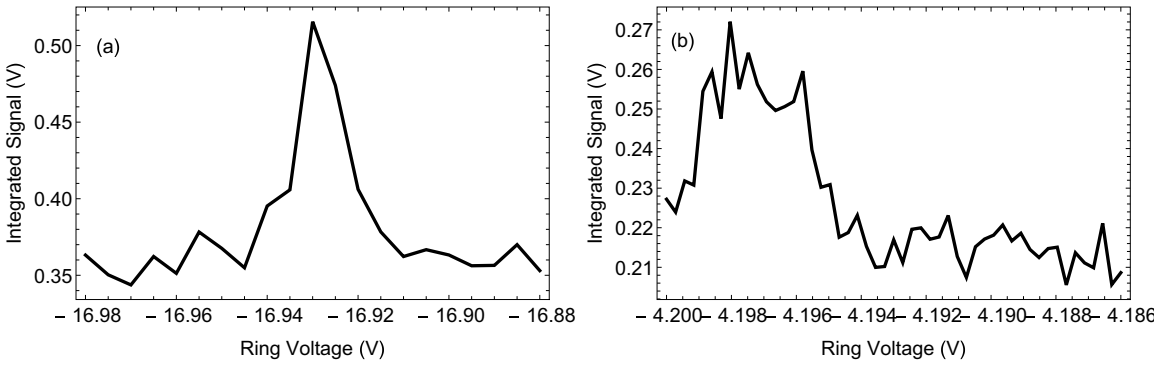

Figure 9.1: Sideband drive scans for (a) Q/m $\approx$ 4 ions and (b) protons.

After characterizing the trapping voltages using sideband drives, we reduced the number of particles in the trap by temporarily reducing the trapping voltage. We then observed dip signals from several protons, using these dips to adjust the anharmonicity compensation ratio. Figure 9.2 shows dips used to adjust the ratio. The narrowest dip observed in this series was 2.2 Hz wide, corresponding to 3 protons.

We then turned to the axial drive. We were able to see strong, driven signals once we had the ratio roughly tuned. This let us tune the anharmonicity compensation further, and demonstrated that the drive lines worked as intended. Figure 9.3 shows axial drives used for anharmonicity tuning, as described in section 2.4.1.

Finally, we were able to observe the undriven signal from a single proton. Figure

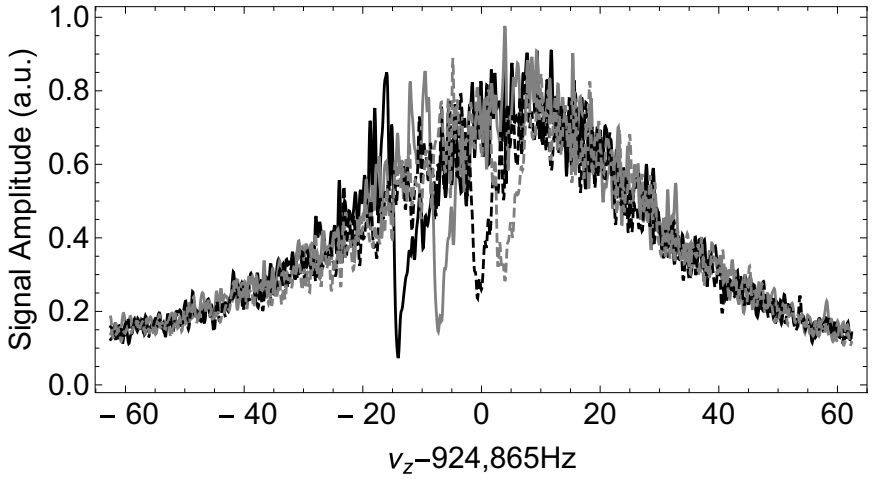

Figure 9.2: Precision trap dips used to tune the anharmonicity compensation in the new apparatus. Figure used as an example in section 2.4.2.

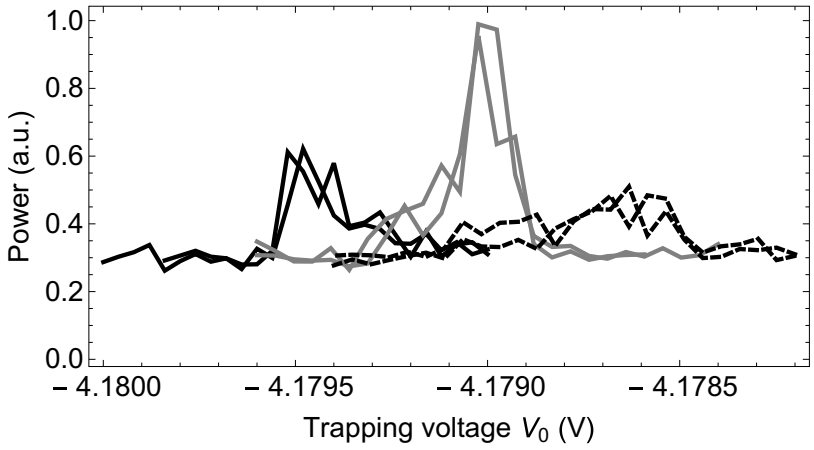

Figure 9.3: Axial drive scans used to tune the anharmonicity compensation in the new apparatus. Figure used as an example in section 2.4.1.

9.4 shows a single-particle dip. This was confirmed to be a single particle using cyclotron signals (see following section).

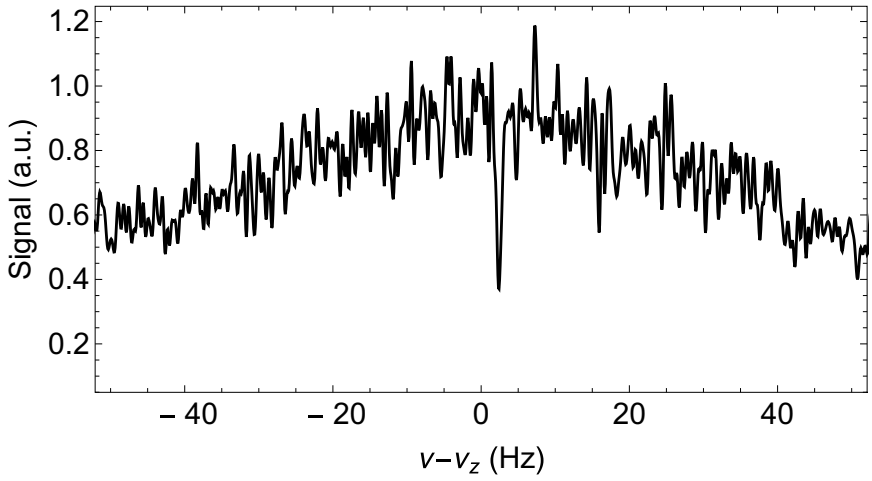

Figure 9.4: One-particle dip in the new apparatus. Figure used as an example in section 2.4.2.

The single-particle signal can also be used to characterize the detector. Measuring the width of the dip allows us to confirm the amplifier's effective resistance on resonance; equation 2.21, $\gamma_z = \left(\frac{q\kappa}{2z_0}\right)^2 \frac{R}{m}$, gives R = 90 MΩ. The measured Q and frequency give an inductance of 1.55 mH via the equation $R = Q\omega L$, in good agreement with the design goal of 1.6 mH.

With a single particle, we could measure the stability of the trap parameters, including the power supply. Figure 9.5 shows an overnight series of measurements of the dip center frequency, with an observed Allan deviation of 70 mHz. While this doesn't address some of the concerns with stability in the analysis trap, including background cyclotron state transitions, it indicates that the voltage supplies, trap geometry, and detection chain are sufficiently stable to identify the ∼135 mHz frequency shift from

a spin flip.[1]

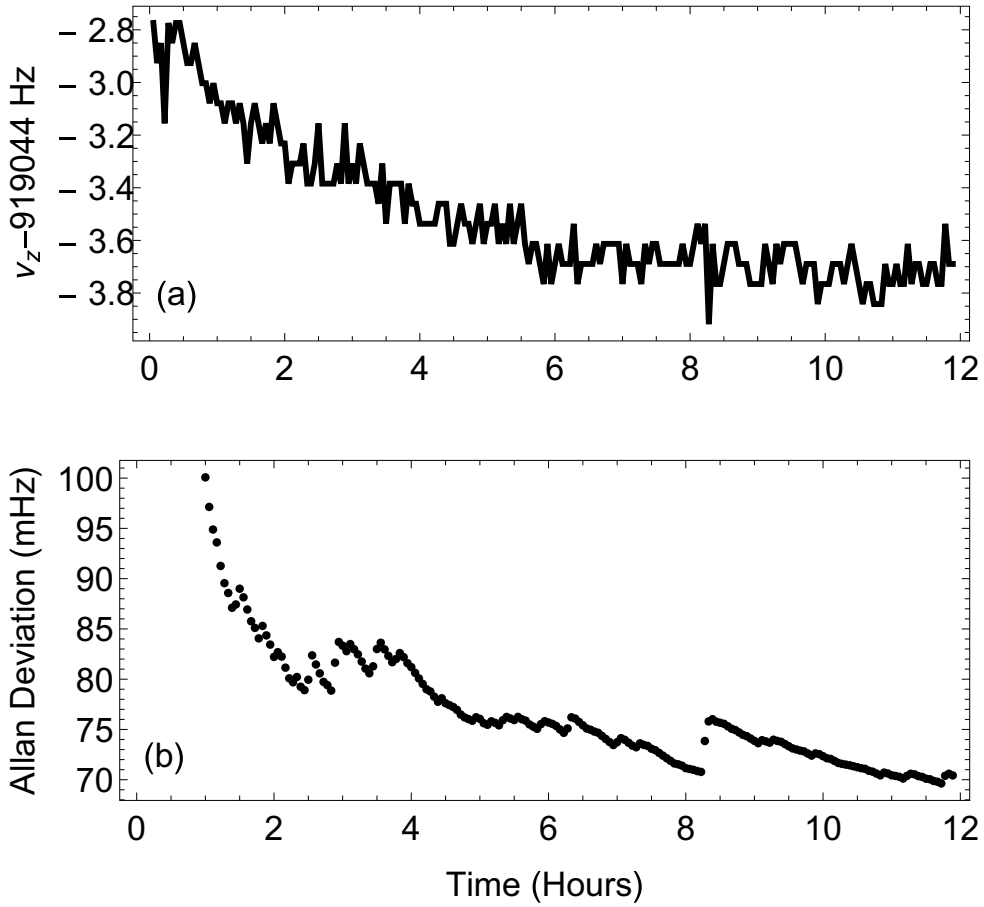

Figure 9.5: (a) Measured frequency and (b) Allan deviation for one night of axial dip measurements.

9.4 Cyclotron signals in the precision trap

We found the cyclotron frequency, and observed the distinct peaks for multiple protons which we reduced to a single particle by iteratively reducing the trap depth, as discussed in section 2.5.1.

[1]This stability was achieved without the SEO and without optimizing averaging time; achieving Allen deviation below 50 mHz in the precision trap at Northwestern should be straightforward.

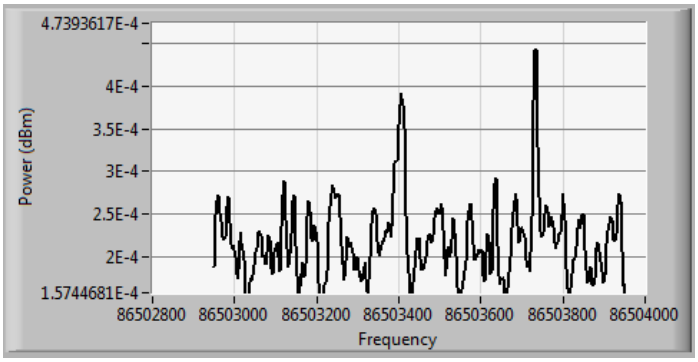

Figure 9.6: Relativistic cyclotron signal from two protons at different radii.

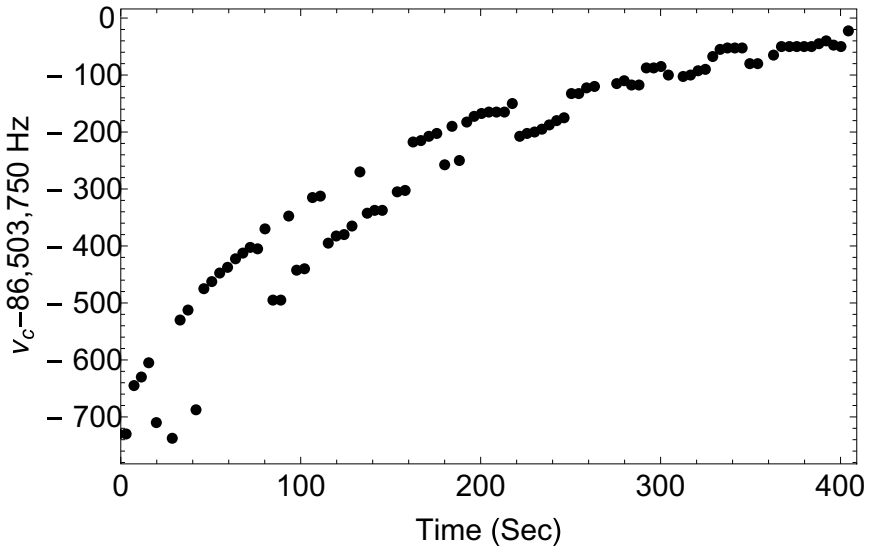

Figure 9.7: Cyclotron decay track in the new apparatus with two excited protons. Figure used as an example in section 2.5.

With a single particle, we observed a decay time of ∼5 minutes, which would correspond to an effective resistance of 10-20 kΩ across the resonator. However, the decays varied from a simple exponential decay. Figure 9.8 shows a single-particle decay fit to an exponential, together with the residuals from the fit. The divergence between fit and data at the start and end of the decay time was characteristic of several observed single-proton decays.

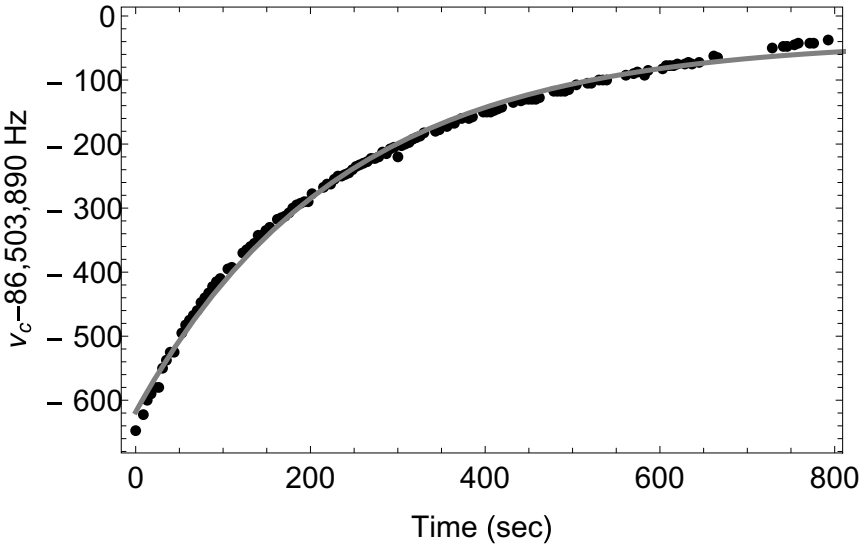

Figure 9.8: Cyclotron decay in the new apparatus, with a simple exponential fit

This indicates some additional effect changing the cyclotron frequency during the decay. An additional term for linear field drift improves the fit accuracy, but the frequency difference between consecutive decays is much smaller than the drift required to account for this curvature. The magnetic field could be less homogeneous than predicted – this curvature prevented us from reliably measuring B_2 by adjusting the particle's equilibrium position – but field inhomogeneity would also not explain this observation; like the relativstic shift, the B_2 shift depends on ρ^2, and would thus

not change the curvature of the decay, only its amplitude. The same is true for an axial frequency shift due to the magnetic bottle – this would shift the trap-modified cyclotron frequency via the invariance theorem $\nu_c = \sqrt{\nu_+^2 + \nu_z^2 + \nu_m^2}$, but this does not introduce a new time dependence.

One possible explanation would be heating in the electrodes due to the strong cyclotron drive. This could change the trap characteristics, and thus the trap-modified cyclotron frequency, with a different timescale from the relativistic decay. Another possible explanation would be some nonlinear interaction with the resonator at large cyclotron radii. A third possibility is that a second ion remained in the trap at a large magnetron or cyclotron radius, distorting the potential at large radii but leaving the behavior at trap center relatively unchanged. Any of these explanations is supported by the fact that excluding the first few minutes of data allows the decays to fit well to a simple exponential. Figure 9.9 shows fits corresponding to different scenarios for this nonlinearity.

9.5 Commissioning the analysis trap

Having characterized the precision trap, we transferred particles into the analysis trap to characterize it. Adiabatic particle transfer was successful, reliably returning particles to the precision trap without excessive magnetron heating. Additionally, we were able to keep a cloud of protons sitting in the analysis trap for >12 hours and return them to the precision trap afterwards.

However, at the point where the experiment had to be warmed up for the move to Northwestern, we had not yet managed to determine the axial trapping potential

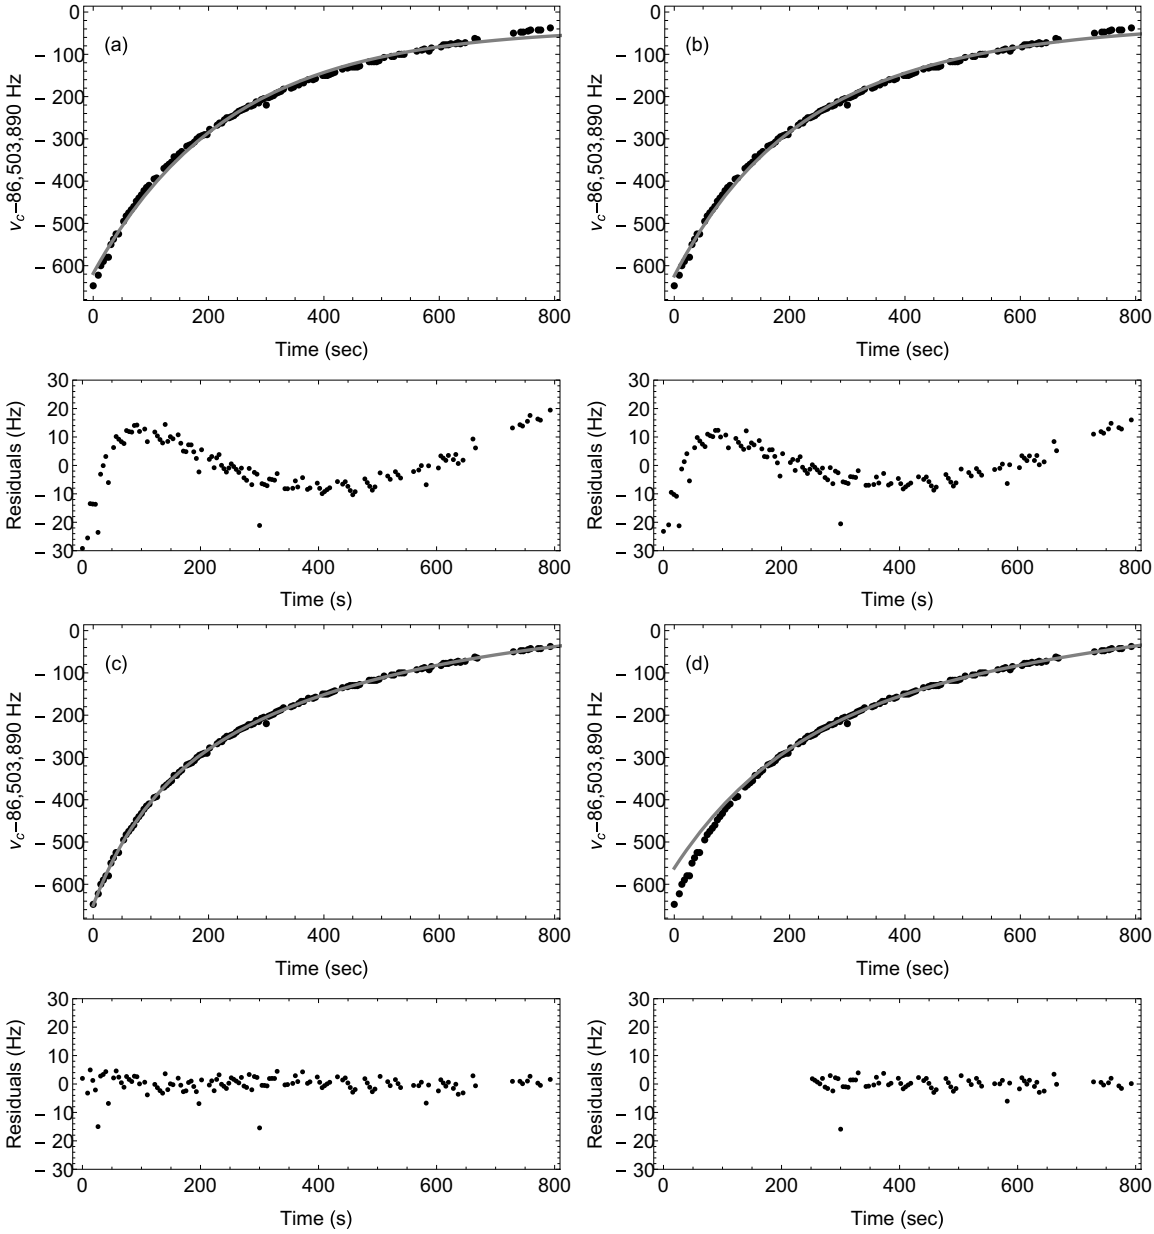

Figure 9.9: Cyclotron decay fits and residuals in the new apparatus.
(a) Decay fit to the expected simple exponential
(b) Incorporating the shifts in ν_z and ν_c from lowest-order even magnetic field inhomogeneity (B_2 and B_4) does not improve the fit
(c) The data fits well to two exponential decays with independent time constants, representing some second process as well as the relativistic shift
(d) Ignoring the first 50 points allows the data to fit well to a simple exponential decay, indicating the unknown process only is active at large radii or shorter times after the drive is applied

in the analysis trap. Voltage scans with the axial drive did not return a signal, likely because of the additional filtering similar to the precision trap. Scanning with a sideband drive has a somewhat larger parameter space, because of the need to identify both the magnetron and axial frequencies. We did observe particle loss on interaction with the sideband drive on some scans, which should help narrow down a range of possible trapping voltages. Next steps for the analysis trap are discussed in section 10.1.1.

9.6 Conclusion

In the new apparatus, we have demonstrated single-proton trapping, detection, and control by observing axial, cyclotron, and magnetron-sideband signals in the precision trap. We have used these signals to characterize the precision trap electrodes and resonators, as well as detection and drive electronics. Particles have been stored overnight in the analysis trap, and ejected using the sideband drive; determining the axial frequency in that trap is the next step for the experiment. All of this was enabled by the new cryostat and vacuum system. Once the analysis trap has been characterized after the move to Northwestern, the road will be open to a sub-ppb measurement of the proton magnetic moment.

Chapter 10

Next Steps Towards a Sub-ppb Measuremetn

While some commissioning work remains to be done, the experiment is nearly ready to make a sub-ppb measurement of the proton magnetic moment. In this section, as a complement to the results presented in chapter 9, we will discuss the remaining steps to be taken before the measurements proposed in chapters 5 and 6 can be implemented. We will also discuss some future improvements to the apparatus.

10.1 Steps to a sub-ppb measurement

The proximate goal for the experiment should be to study the feasibility of implementing the sub-ppb measurement methods proposed in chapters 5 and 6. In this section, we outline a series of investigations to determine which method represents the optimal path for the experiment.

10.1.1 Characterizing the analysis trap

Once the apparatus is established at Northwestern and single protons are again trapped, the first goal should be finding axial signals in the analysis trap. This is necessary not only for eventual precision measurements, but also for several of the diagnostics and demonstrations discussed later in this section.

Single-proton storage times in the analysis trap of 12+ hours were demonstrated during the work described in chapter 9. Additionally, we observed particle loss from the analysis trap due to the application of a strong sideband drive. However, the particle's axial frequency has not yet been observed.

Ideally, we would apply a strong axial drive and find the voltage where the proton responds, as in figure 2.6 for the precision trap. However, because of machining tolerances and thermal contraction, the trap coefficients are initially uncertain. The applied potential is also uncertain due to non-reversing voltages (see sec. 3.2.4 of [35]), possibly including a chemical potential between the copper leads and the iron ring. There is thus a large range of possible resonant voltages to cover. During the final week of operation at Harvard, we did not see results from analysis trap drive sweeps.

Other tests could yield faster results than drive scans with a single proton. Ions with higher charge states could be identified in the precision trap and then sent to the analysis trap, as they have a larger response to an axial drive (see section 5.1.2 of [34]). Alternatively, many protons could be sent to the analysis trap, broadening and increasing the magnitude of the driven response.[1]

[1] A large proton cloud in the analysis trap would have to be transferred back to the precision trap at intervals and measured, as the magnetic bottle and the anharmonicity of the smaller trap could contribute to particle loss out of a large plasma.

Possibly the most promising method would focus on finding the heating and cooling responses to the sideband drive. This method first let us characterize the precision trap when the axial drives were over-filtered. The sideband drive can be made strong enough to give a broad response without interfering with detection. Additionally, interaction between protons and the sideband drive has already been demonstrated via particle loss from the analysis trap. This loss could be used as an initial signal to reduce the possible range of resonant voltages, before resolving ν_z and ν_m with narrower scans at different drive frequencies. Once the trapping voltage for a given ν_z has been found, the frequency stability and trap coefficients can be measured.

10.1.2 The self-excited oscillator

Although single-proton signals with good stability have been observed in the precision trap, we have not yet implemented the self-excited oscillator (section 2.4.3). A digital signal processor (DSP) with a new microchip was constructed by the Harvard Electronic Instrument Design Lab. The new DSP programming has to be characterized and compared to the previous generation. With single-particle dips in the precision trap, it should be possible to follow established procedures for self-excited oscillation [34, 35, 43] using the previous-generation DSP. Once a stable self-excited oscillator is established using that DSP, the new DSP programming can be modified if necessary to give equivalent amplitude control. In both the precision and analysis traps, the SEO should significantly increase SNR over the dip method, and thus improve the other measurements proposed in this chapter.

10.1.3 Feedback and axial temperature

Once the axial frequency is measured in the analysis trap, we can implement the improvement suggested in section 7.2.6. The axial temperature should be measured using the method detailed in section 5.3.1 of [35] at different FET transconductances. If feedback through the FET has a similar effect to feedback applied externally, the transconductance and matching network tuning can be used to lower the temperature. A lower axial temperature would imply a narrower magnetron distribution after sideband cooling, improving cyclotron energy measurements and axial stability in the analysis trap; it would also reduce the linewidth in the precision trap from residual magnetic field gradients.

10.1.4 Cyclotron quantum number measurement

Measurements of ν_z in the analysis trap can also be used to measure the cyclotron energy via the magnetic bottle shift. This is an essential part of the pulsed Ramsey measurement of ν_c (section 5.2) as well as the anomaly drive g-factor measurement (chapter 6). Developing this technique would imply characterizing and minimizing transfer-related shifts in the trapping potential, as well as optimizing the search for ν_z after transfer to the analysis trap.

10.1.5 Characterization of the FET switch and cyclotron energy background noise

The FET switch as implemented is sufficient to short the cyclotron resonator, to the point where it can no longer be detected even with a strong drive (fig. 7.25). To

evaluate whether it is sufficiently decoupled to allow undisturbed cyclotron evolution for the pulsed Ramsey and anomaly-drive techniques, its effect on the particle must be investigated. This could be done as a set of two tests. First, partway through a cyclotron decay the switch should be closed for some time, and the damping rate during that time measured from the change in relativistic shift when the switch is re-opened[2]. Second, a series of transfers should be performed to the analysis trap and the cyclotron energy measured after each transfer, as described in sections 5.2 and 6.4.2. The spread of energies with the FET switch open and closed, at different times spent in the precision trap, gives a measure of remaining coupling to the resonator.

A histogram of the observed cyclotron energies with the FET switch open could be fit to a Boltzmann distribution, to measure the cyclotron amplifier's electronic temperature. This can be compared to a histogram for the cooling cyclotron amplifier once the cyclotron frequency is found in that trap.

If the FET switch successfully decouples the particle from the resonator, measurements of the cyclotron energy after transfer with the switch closed would characterize the cyclotron state evolution, to evaluate sub-PPB measurement prospects.

1. The scatter in cyclotron energy with repeated transfers would characterize the background for an anomaly-drive measurement (section 6.4.2). Measuring this at different starting ρ_c values would give limits on acceptable ρ_c for that method.

2. The scatter at different starting ρ_c can be Fourier transformed to find its scaling with different powers of ρ_c. Any constant (ρ_c^0) component can be attributed to

[2]This would have to be carefully disentangled from electrostatic or frequency-pulling effects caused by the amplifier detuning – see for example figure 3.12 of [45]. Measuring damping rate vs time with the amp shorted should separate these effects.

voltage instability, ν_z measurement resolution, etc, as couplings to the cyclotron motion should scale with ρ_c. Once characterized, these scatter sources can be identified and mitigated if necessary.

3. The particle can be left in either the precision or analysis trap for different amounts of time between measurements to find the scaling of transition rate vs trap size (estimated in sec. 6.4.2). The scatter in the analysis trap can be compared to that observed in [35], and the scatter per transfer can be compared to [18], as indicators of whether we have sufficiently filtered noise near ν_c.

4. Similar data can be taken for different transfer parameters – well depth, potential difference between adjacent electrodes, wait time per step, and rate of voltage change during transfer steps – to confirm that our transfer procedures are sufficiently adiabatic to leave the cyclotron state undisturbed.

5. Finally, we should compare scatter in cyclotron energies measured with the same initial ρ_c and transfer parameters, but with the experiment alignment shifted using the gearbox. The cyclotron state should be unchanged during transfers, even with slightly misaligned fields; misalignment should express itself as magnetron radius growth. However, this should be tested experimentally.

10.1.6 Developing cyclotron frequency measurement techniques

In section 9.4 we discussed the state of cyclotron frequency measurement when the experiment was shipped to Northwestern. Cyclotron decays have been observed,

but the deviation from an exponential must be understood before decay fitting can be relied on to the sub-Hz level. Possible explanations were discussed in section 9.4. Trap electrode heating can be separated from ρ_c-dependent nonlinearity using the FET switch – closing the switch after the particle is excited would let the electrodes cool back to thermal equilibrium. If the deviation from the exponential is unchanged by this cooling time, electrode thermal expansion can be ruled out.

As discussed in chapter 5, cyclotron decays are not well-suited for a g-factor measurement because of the cooling time and relativistic shift. A ν_c measurement using a cyclotron-to-axial sideband drive and split dip [51], as used in [18, 19] and discussed in section 5.1, can be considered in our apparatus. However, the precision of this method is limited by instability in the axial dip measurement.

This split dip method also requires a cyclotron-to-axial sideband drive, whose geometry is precluded in the precision trap by the current loops for the anomaly drive (section 2.7). Simultaneous application of a drive to a compensation electrode and half of the ring also gives the correct gradient, as demonstrated for the magnetron sideband drive in chapter 9; but phase delays on the separate paths to those electrodes may have a larger effect at the cyclotron frequency. Demonstrating this method would require successfully applying a coupling drive and observing a split dip.

The pulsed Ramsey method detailed in section 5.2 is more promising. An initial proof of concept would involve transferring to the analysis trap after one or two drive pulses and measuring the cyclotron energy. This would attempt to show (a) excitation from a thermal state to the same well-defined initial amplitude many times using an identical drive pulse; (b) a consistent effect of the second drive pulse, for an evolution

time short enough that it would not be sensitive to magnetic field drift[3]; and (c) a fringe pattern when scanning that short evolution time. Once those have been demonstrated, the evolution time can be increased for a precision measurement.

10.1.7 Evaluate the magnetic field

Once any of the above precision cyclotron frequency measurements has been demonstrated, it can be used to evaluate magnetic field drift and stability. This will likely take multiple weeks, as the magnet settles after charging, and should be started as soon as reasonable (but can be done simultaneously to other investigations). The pressure regulation system will be the main control over field drift, but the association of magnetic field stability with various temperatures (e.g. magnet chassis, hat, ambient room temperature etc) and conditions in the experiment hall (e.g. other solenoids, operation of other experiments) should be monitored.

The cyclotron frequency can also be used to measure the magnetic field profile in the trap. Magnetic field gradients in the precision trap can be measured by shifting the particle's equilibrium position, with an asymmetric offset voltage applied to both endcaps, and measuring ν_c. (This requires an offset be applied to the ring and comps as well, to keep the axial frequency constant; see section 2.3 of [45]). This was attempted before the move to Northwestern, but the resulting ν_c shifts were smaller than the uncertainty due to the non-exponential cyclotron decay profile. However, we did produce an initial measurement of the axial frequency shift as a function of applied asymmetric endcap offset voltage, which could be used as a starting point for

[3]for example, an evolution time representing only a .1 or .01 ppm cyclotron frequency resolution

this measurement.

Uncertainties in the magnetism of materials at 4K and 5.6 T – especially the copper trap electrodes and the drive line tuning capacitors – could result in a measured B_2 value larger than predicted for the precision trap design in section 8.1.1. In this case, several options are available for B_2 optimization, which are discussed in section 10.2.4.

10.1.8 Commissioning the cooling trap

To evaluate the single-spin-flip and anomaly-drive measurement techniques, we will need to commission the cooling trap. Without the cooling trap, the single-spin-flip method will be extremely slow because of the need for sub-thermal particle selection; the anomaly-drive method will also be hindered to some degree.

Ideally, the cooling trap will operate as a full trap, with its own axial amplifier and drives. The electronics for this are already constructed, but the cooling axial amplifier still needs to be installed. However, the cooling trap is an uncompensated three-electrode trap (to maximize the coupling between the resonator and particle). Design principles for the cooling trap are detailed in [62]. Finding a driven axial response in this trap should be possible, but depending on the trap coefficients after machining imperfections and thermal contraction it may not be possible to find a dip.

The cooling trap does not require high axial frequency resolution, and could be operated without an axial amplifier – it is only required for diagnostics and magnetron cooling. As long as particle storage is stable for longer than the cyclotron damping time, magnetron cooling could also be performed in the analysis trap. Finding ν_c will therefore be the essential step in commissioning the cooling trap.

The cooling trap is located in the magnetic field gradient of the iron analysis ring. The particle thus experiences a different average magnetic field depending on the amplitude and center of the axial motion. This may allow tuning of ν_c with asymmetric voltages, which would avoid the RF loss from a varactor. However, it also introduces uncertainty in the predicted ν_c, which could make finding the frequency difficult. We propose two methods to search for ν_c, depending on axial signals.

1. If we are able to measure an axial dip in the cooling trap, then we can monitor it as we sweep a cyclotron drive. When the drive passes through resonance ρ_c should increase dramatically. This will couple to B_2 from the iron ring, causing the axial dip to shift or vanish. Even though B_2 in the cooling trap will be much smaller than in the analysis trap, we can still observe large changes in ρ_c.

2. If we are unable to find a dip, we can periodically measure ρ_c by transferring to the analysis trap and finding ν_z. Other than, this procedure is similar to (1) – finding ν_c by observing the frequency where a drive excites ρ_c.

These methods share the advantage of not requiring a priori tuning of the cyclotron amplifier; they are thus better suited to finding an uncertain cyclotron frequency than decays. Once we have found ν_c, the cyclotron amplifier can be tuned onto resonance. The cyclotron frequency as a function of axial position (with asymmetric voltages applied to the endcaps) can be measured to determine the viability of tuning ν_c to match the resonator; and the damping time and temperature can be measured to determine how quickly we can select a cold proton for g-factor measurements.

10.1.9 Selecting a sub-ppb measurement method

Once the precision, cooling, and analysis traps have been characterized, we should have all the necessary information to compare the measurement methods described in chapters 5 and 6. Criteria for this comparison were described in section 6.7.

At this point, we should have access to all the necessary techniques to make a proof-of-concept anomaly-drive measurement of the g-factor, using large Rabi frequencies and power-broadening the transitions. This would be the proposed next step under conditions favorable to the anomaly method.

If conditions in the experiment favor the simultaneous SOF method, then the next step would be a demonstration of adiabatic fast passage (AFP) as described in section 4.3. Our estimates of the achievable precision and time in chapter 5 relied on AFP – without it, the comparison to the anomaly-drive method is much less favorable. Demonstrating AFP would imply measuring the Allen deviation with and without adiabatic drive sweeps, and showing improved spin-state identification over [7].

10.2 Further apparatus improvements

In this section we will discuss near-term upgrades to the apparatus which could assist in the measurement program described above, or which may be necessitated by some of those measurements.

10.2.1 Improvements to detection

The axial detectors implemented in this work already represent a significant improvement over past generations of the experiment. However, we have not yet controlled for the effects of feedback on the lower-loss resonators. As described in section 7.2.8, a pair of tuning varactors should be added to the output matching networks on the axial amplifiers. This would allow individual control over Q, gain, and impedance matching. This would be especially beneficial if it gives an additional lever to control the particle's axial temperature; even without that advantage, *in situ* optimization of damping and signal-to-noise would be a significant new capability.

The precision cyclotron amplifier is likely limited by anomaly-drive wiring along its ground path. However, the cooling cyclotron amplifier could be improved. The cooling amplifier on the CERN apparatus has a Q of 1200 [62], while the coil on this apparatus has a Q of only 800. Attaching a new coil (or the existing Q=1200 coil) would reduce the damping time. This coil could also be measured while attached to each trap, to show to what extent the precision trap Q is limited by the drive lines.

10.2.2 Cyclotron amplifier decoupling

The FET switch was shown in section 7.4.1 to sufficiently short the amplifier that no signal was observed from either Johnson noise or an external drive. However, if the measurement in section 10.1.5 indicates residual coupling, there are two more options for decoupling the amplifier. Reference [18] reported successful decoupling using a tuning varactor to shift the resonance frequency by several times the linewidth. This kept particle-amplifier interaction low enough to allow largely free evolution.

A more drastic approach would be to implement a piezo-based cryogenic mechanical switch, which could either disconnect the amplifier from the trap or short it to ground. This was partially developed as an undergraduate project, and space for it exists in the tripod. Testing demonstrated successful cryogenic operation, but the ceramic material of the piezo frequently broke due to water absorption; reliable operation over multiple thermal cycles was only observed with the switch under vacuum. The next step would be to cool down the switch attached to a resonator and measure its effect on Q. A similar switch was described in [87] and demonstrated in [131].

If none of these methods proves sufficient, the cyclotron amplifier could also be disconnected entirely from the precision trap and moved to the loading trap. Potential issues with this approach are discussed in [62].

10.2.3 Improvements to wiring

There are several potential improvements which could be made to the experiment wiring, both at DC and at RF. Reference [62] discusses a scheme to speed up transfers using diode pairs to bypass the filtering time constants when applying large voltages, previously demonstrated in [87]; this should be added to the experiment.

Additionally, the DC filterboards as currently constructed have suffered occasional component breakages. These were assembled with 60-40 lead-tin solder, using successful techniques from past experiments. However, the new apparatus is more vulnerable to component failures because of the number of electrical connections. We should consider replacing the 60-40 solder on our DC filterboards with tin-lead-antimony solder. This has been recommended [132, 133] for cryogenic use because it is more ductile

than 60-40 after application, putting less strain on components during cooldown.

There are also a couple of areas for improvement of the RF wiring. The RF drive lines do not currently connect to the cryogenic 50Ω tee attenuators at the 77K and 18K stages. Those attenuators exist to reduce noise picked up at room temperature and to replace room-temperature with cryogenic Johnson noise. Currently, the floating grounds of each attenuator shunt the current down to the low side of the twisted pair or coaxial cable, rather than attenuating it. The ground pads should be shorted to overall experiment ground. They would then function as intended, although care should be taken about potential ground loops added from this coupling.

Additional filtering may be required to reduce noise at ν_c (as discussed in section 10.1.5), as free cyclotron evolution is essential for both proposed sub-PPB measurement techniques. The capacitive divider ratio could be increased on the cyclotron drive line, and additional filtering could be added to other drive lines.

10.2.4 Controlling the magnetic bottle

If the magnetic field inhomogeneity in the precision trap is significantly worse than calculated in section 8.1.1, several modifications can be made. The most likely source of inhomogeneity is the presence near the trap of tuning capacitors for the drive lines. Their position was chosen to account for their magnetization as measured at Northwestern using a SQUID magnetometer. However, the SQUID measurements had a high uncertainty, and the actual magnetization may be larger than measured.

Moving the capacitors further from the precision trap would reduce their B_2 contribution while requiring a change to their tuning. They are currently connected to

the trap with a length of coax, long enough to suppress the expected B_2 but short enough that the capacitor-electrode distance is in the near-field limit at the drive frequency. As long as the trap electrodes and capacitors remain close compared to the wavelength, this near-field limit holds; deviations from this approximation can be tuned out by adjusting the value of the capacitor. If the capacitors need to be moved a significant fraction of the wavelength away from the trap, a full half-wavelength of coax can be added without affecting the drive tuning.

A second source of a greater-than-predicted bottle field could be the electrodes themselves. The shape of the copper trap electrodes is chosen (sec. 8.1.1) such that their B_2 contribution cancels the predicted residual B_2 from the analysis trap iron ring. However, this cancellation relies on a measured value of copper's magnetization, which is subject to a large uncertainty. If the real and design values diverge, we could be left with a substantial B_2. The precision trap has been designed to allow correction of B_2 without constructing new electrodes. Rings of material with known magnetization can be added around the endcaps to adjust B_2, as shown in figure 10.1. When fixed in place, these would stably adjust the field shape. This requires accurate measurements of the residual B_2 and the magnetization of the material added.

The precision trap also allows active control over the magnetic field shape. The location marked in figure 10.1 can hold field-correction coils – many turns of thin wire potted in epoxy. Currents through these wires would modify the field shape in the precision trap. Assuming we use .008" diameter lacquered magnet wire, this geometry would contain ∼90 turns of wire in each coil. The $\hat{z}$ magnetic field from a

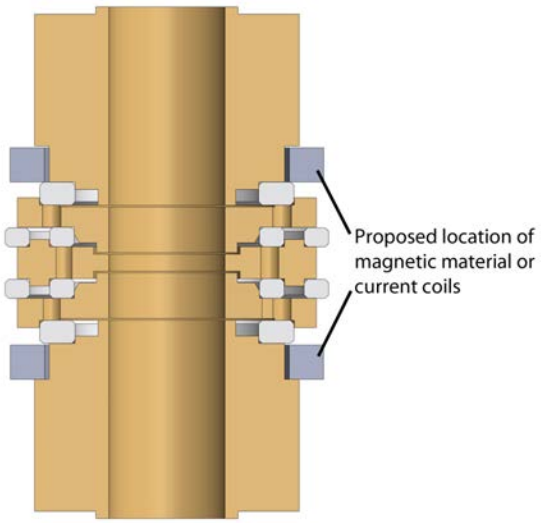

Figure 10.1: Precision trap with proposed geometry for magnetic bottle modification.

loop with radius R and current I at an axial distance Z, near the z-axis, is [134]

$$B_z(z,\rho) = \frac{\mu_0 I R^2}{2}\left(\frac{1}{\left(R^2+(Z+z)^2\right)^{3/2}} + \frac{3\left(R^2-4(Z+z)^2\right)}{4\left(R^2+(Z+z)^2\right)^{7/2}}\rho^2\right) \tag{10.1}$$

Expanding this around $z = 0$ gives the first and second order axial gradients:

$$B_z = \frac{\mu_0 I R^2}{2\left(R^2+Z_0^2\right)^{3/2}} - \frac{3\mu_0 I R^2 Z_0}{2\left(R^2+Z_0^2\right)^{5/2}}z - \frac{3\mu_0 I R^2(R^2-4Z_0^2)}{4\left(R^2+Z_0^2\right)^{7/2}}\left(z^2-\frac{\rho^2}{2}\right)+\ldots \tag{10.2}$$

where the second-order axial gradient has the same form as the axial inhomogeneity B_2. The first-order gradient averages to zero over the course of an axial oscillation, while the zeroeth-order term shifts the overall field value. The radial field can be expanded near the origin as well:

$$B_\rho = \frac{3\mu_0 I R^2 Z_0}{4\left(R^2+Z_0^2\right)^{5/2}}\boldsymbol{\rho} + \frac{3\mu_0 I R^2(R^2-4Z_0^2)}{4\left(R^2+Z_0^2\right)^{7/2}}z\boldsymbol{\rho}+\ldots \tag{10.3}$$

The first term here is analogous to the B_1 radial term from magnetized material; it

adds to zero for current loops with reflection symmetry across the xy-plane.[4] The second term corresponds to the radial term of B_2 from equation 2.15, giving a total B_2-cancelling shift from the current loops of

$$B_2\left(\left(z^2-\frac{\rho^2}{2}\right)\hat{B}-\left(\hat{B}\cdot\mathbf{z}\right)\boldsymbol{\rho}\right),\quad B_2=\frac{3\mu_0 I R^2(R^2-4Z_0^2)}{4\left(R^2+Z_0^2\right)^{7/2}} \tag{10.4}$$

A sum over 90 loops by position in the coil gives B_2 of 0.11 Tesla/m^2 per milliampere of applied current. Less than one mA should be sufficient to cancel the residual bottle field from a discrepancy between the copper magnetization used to design the trap and the true value. The entirety of the residual bottle field due to the analysis trap iron ring could be canceled with less than 3 mA of current.

The loops of figure 10.1 using .008" magnet wire would correspond to ∼320 total inches of wire, with a room temperature resistance of ∼4.8 Ω per foot. Conservatively estimating a fourfold reduction in resistivity at 4K, the heat load would be 30 μW per mA; this should dissipate easily through the bias lines. (The heat load could be reduced with more turns and less current, or eliminated with superconducting wire.)

Using methods developed for the antiproton q/m measurement [46, 45], we should be able to measure B_2 with a precision of .01 Tesla/m^2. Cancelling the bottle with this precision would require a current source with stability better than 0.1 mA. B0 stability imposes a stricter requirement – equation 10.2 gives a shift of 8.5 μT per mA. Modern precision current supplies[5] can provide short-term stability of ∼5 nA and long-term stability of order 20-30 PPM out of 1mA; this corresponds to a short-term B0 stability of $7*10^{-12}$ and long-term stability of $4.5*10^{-11}$, respectively well

[4] Even in the case of imperfect symmetry, this can be neglected, since it adds in quadrature with the homogeneous field. The radial contribution from 90 turns and 1 mA shifts B_0 by less than 10^{-15}.

[5] e.g., Rhode and Schwarz model 6166, or equivalent models from Keysight etc.

below the measurement linewidth and expected background magnetic field drift.

A final estimate deals with the effect of the external magnetic field on these current coils. The Lorentz force would apply an inward pressure against the epoxy of

$$I * N * B_0 * \frac{2\pi R}{A} \tag{10.5}$$

where I is current, N is the number of turns, R is coil radius and A is coil inner surface area. For 90 turns per coil and the geometry of figure 10.1, the pressure would be .085 PSI per mA. This pressure will be applied in against the epoxy; if the current loop axis is not perfectly parallel to the magnetic field, some of this pressure could be applied as torque on the electrode stack. The magnitude of this force is small compared to thermal contraction and eddy currents during insertion into magnet. Still, before the current coils are attached to the experiment they should be tested on dummy electrodes to confirm no effect from the Lorentz force.

Chapter 11

Conclusion

This thesis reports work which concludes the first and begins the second generation of a proton-antiproton magnetic moment experiment. The first generation concluded with a 680-fold improved comparison of the proton and antiproton magnetic moments, as well as the first identification of individual spin-flips of a single proton. Methods were next developed for the second generation of the experiment, with the goal of a measurement with a sub-ppb lineshape. Finally, the required new apparatus was designed, constructed, and commissioned with single trapped protons.

11.1 Results in the first-generation apparatus

The first stage of work reported in this thesis consisted of contributions to two reports published in Physical Review Letters [6, 7]. Reference [6] reported the first single-particle measurement of the antiproton magnetic moment:

$$\frac{\mu_{\bar{p}}}{\mu_N} = \frac{g_{\bar{p}}}{2} = 2.\ 792\ 845\ (12) \tag{11.1}$$

To accomplish this, we transported and established the experiment at the CERN Antiproton Decelerator, where we successfully trapped and cooled antiprotons despite challenges from our small trap diameter. A single antiproton was stored for several months in the trap. We achieved sufficient frequency stability, despite noise from the accelerator and other experiments, to resolve the frequency shift due to spin-flips, and used that to measure the spin excitation fraction and the spin and cyclotron lineshapes. The reported result represented a factor of 680 improvement over previous measurements, and (combined with the previous proton measurement in the same apparatus) provided a test of CPT invariance with a precision of 5 ppm.

Reference [7] reported the first observation of individual spin flips of a single proton (published simultaneously with [8]). This demonstrated the ability to identify the spin state of a proton at a particular point in time, a necessary prerequisite for high-precision measurements using quantum jump spectroscopy.

11.2 Two methods for sub-ppb precision

To improve the CPT test in the baryon sector represented by the proton-antiproton magnetic moment comparison, we intend to measure a linewidth narrower than one part per billion. This requires a new method which simultaneously interrogates the spin and cyclotron frequencies. We therefore evaluated in detail two possible methods to achieve a sub-ppb linewidth. The simultaneous separated oscillatory fields method would measure both frequencies using simultaneous pulses and a shared evolution time. The quantum walk method, proposed here, involves driving simultaneously on the spin and anomaly transitions, inducing scatter in the cyclotron state when

both drives are resonant. Each of these measures the magnetic moment directly by measuring a ratio of frequencies. While the separated oscillatory fields method has the advantage of measuring an interference pattern whose width is determined by the evolution time, the quantum walk method can accumulate statistics faster by driving many cyclotron transitions per trial. The detailed estimates indicate that the quantum walk method may be significantly more efficient; however, we also have provided a framework to evaluate the two methods in the context of measured experimental parameters in the new environment at Northwestern.

11.3 A new, improved apparatus

In order to implement these proposed methods, we designed and constructed an entirely new apparatus. This apparatus features a redesigned set of Penning traps. The precision trap will reduce magnetic field inhomogeneity. The analysis trap will reduce the rate of cyclotron transitions due to RF noise. The loading trap will catch and store more antiprotons in a reservoir. The cooling trap will quickly damp the particle's cyclotron motion. These electrodes were polished with high geometric precision and low surface roughness, and gold-plated with a robust electroplating procedure.

Besides the electrodes, other significant upgrades to the new apparatus include improved detection electronics and RF drive lines, as well as the implementation of a cryogenic alignment and positioning system. The detection electronics benefitted from low-loss materials, rigorous coil-winding techniques, and an improved understanding of effects such as magnetoresistance and feedback. The new RF drive lines reported include two designs for drive-line resonators which enable efficient coupling of the

drive current to the electrodes and particle. The cryogenic positioning gearbox enables in situ adjustment of the angle between trap electric and magnetic fields, while an alignment sensor in development will allow reliable correction of that angle. The new apparatus also featured a new titanium cryostat, a precisely aligned G10 support structure, and denser DC wiring with improved filtering.

Finally, we report initial results from the commissioning of the new apparatus. Protons were loaded with the field emission point, and their number reduced to a single particle. The axial, cyclotron, and magnetron motions of a single proton were each measured in the precision trap. Transfers to the analysis trap were successfully demonstrated, together with long storage times in both traps. Some commissioning work remains to be done after the move to Northwestern, especially finding an axial signal in the analysis trap and implementing techniques for efficient, precise measurement of the cyclotron radius and cyclotron frequency. Once those have been demonstrated, the path will be open for application of the new methods to a measurement of the proton magnetic moment with a sub-ppb linewidth.

Bibliography

[1] L. S. Brown and G. Gabrielse, Rev. Mod. Phys. **58**, 233 (1986).

[2] K. Blaum, Yu. N. Novikov, and G. Werth, Contemporary Physics **51**, 149 (2010).

[3] E. Myers, International Journal of Mass Spectrometry **349-350**, 107 (2013), 100 years of Mass Spectrometry.

[4] B. Odom, D. Hanneke, B. D'Urso, and G. Gabrielse, Phys. Rev. Lett. **97**, 030801 (2006).

[5] D. Hanneke, S. Fogwell, and G. Gabrielse, Phys. Rev. Lett. **100**, 120801 (2008).

[6] 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).

[7] J. DiSciacca, M. Marshall, K. Marable, and G. Gabrielse, Phys. Rev. Lett. **110**, 140406 (2013).

[8] A. Mooser, H. Kracke, K. Blaum, S. A. Bräuninger, K. Franke, C. Leiteritz, W. Quint, C. C. Rodegheri, S. Ulmer, and J. Walz, Phys. Rev. Lett. **110**, 140405 (2013).

[9] R. Frisch and O. Stern, Z. Phys. A: Hadrons Nucl. **85**, 17 (1933).

[10] I. Estermann and O. Stern, Z. Phys. A: Hadrons Nucl. **85**, 4 (1933).

[11] F. Bloch, W. W. Hansen, and M. Packard, Phys. Rev. **70**, 474 (1946).

[12] H. M. Goldenberg, D. Kleppner, and N. F. Ramsey, Phys. Rev. Lett. **5**, 361 (1960).

[13] P.F. Winkler, D. Kleppner, T. Myint, and F.G. Walther, Phys. Rev. A **5**, 83 (1972).

[14] B. L. Roberts, Phys. Rev. D **17**, 358 (1978).

[15] A. Kreissl, A.D. Hancock, H. Koch, Th. Köehler, 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, Z. Phys. C: Part. Fields **37**, 557 (1988).

[16] T. Pask, D. Barna, A. Dax, R.S. Hayano, M. Hori, D. Horváth, S. Friedrich, B. Juhász, O. Massiczek, N. Ono, A. Sótér, and E. Widmann, Phys. Lett. B **678**, 55 (2009).

[17] J. DiSciacca and G. Gabrielse, Phys. Rev. Lett. **108**, 153001 (2012).

[18] 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 (2017).

[19] G. Schneider, A. Mooser, M. Bohman, N. Schon, J. Harrington, T. Higuchi, H. Nagahama, S. Sellner, C. Smorra, K. Blaum, Y. Matsuda, W. Quint, J. Walz, and S. Ulmer, **358**, 1081 (2017).

[20] T. D. Lee and C. N. Yang, Phys. Rev. **104**, 254 (1956).

[21] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, Phys. Rev. **105**, 1413 (1957).

[22] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, Phys. Rev. Lett. **13**, 138 (1964).

[23] G. Lüders, Ann. Phys. **2**, 1 (1957).

[24] M. D. Schwartz, *Quantum field theory and the standard model* (Cambridge University Press, ADDRESS, 2014).

[25] V. Alan Kostelecký and R. Potting, Nucl. Phys. **B359**, 545 (1991).

[26] V. Alan Kostelecký and R. Potting, Phys. Rev. D **51**, 3923 (1995).

[27] D. Colladay and V. A. Kostelecký, Phys. Rev. D **55**, 6760 (1997).

[28] C. Patrignani et al (Particle Data Group), Chin. Phys. C **40**, 100001 (2016).

[29] A. D. Sakharov, Sov. J. of Exp. and Th. Phys. Lett. **5**, 24 (1967).

[30] A. D. Dolgov and Ya. B. Zeldovich, Rev. Mod. Phys. **53**, 1 (1981).

[31] G. Gabrielse, D. Hanneke, T. Kinoshita, M. Nio, and B. Odom, Phys. Rev. Lett. **97**, 030802 (2006), *ibid.* **99**, 039902 (2007).

[32] Emmanuel Chang, William Detmold, Kostas Orginos, Assumpta Parreño, Martin J. Savage, Brian C. Tiburzi, and Silas R. Beane, Phys. Rev. D **92**, 114502 (2015).

[33] Assumpta Parreño, Martin J. Savage, Brian C. Tiburzi, Jonas Wilhelm, Emmanuel Chang, William Detmold, and Kostas Orginos, Phys. Rev. D **95**, 114513 (2017).

[34] N. Guise, Ph.D. thesis, Harvard University, 2010.

[35] J. DiSciacca, Ph.D. thesis, Harvard University, 2013.

[36] J. Goldman, Ph.D. thesis, Harvard University, 2004.

[37] G. Gabrielse and F. Colin MacKintosh, Intl. J. of Mass Spec. and Ion Proc. **57**, 1 (1984).

[38] G. Gabrielse, L. Haarsma, and S. L. Rolston, Intl. J. of Mass Spec. and Ion Proc. **88**, 319 (1989), ibid. 93:, 121 (1989).

[39] D. Wineland, P. Ekstrom, and H. Dehmelt, Phys. Rev. Lett. **31**, 1279 (1973).

[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] N. Guise, J. DiSciacca, and G. Gabrielse, Phys. Rev. Lett. **104**, 143001 (2010).

[44] B. D'Urso, Ph.D. thesis, Harvard University, 2003.

[45] David Forrest Phillips, Ph.D. thesis, Harvard University, 1996.

[46] G. Gabrielse, A. Khabbaz, D. S. Hall, C. Heimann, H. Kalinowsky, and W. Jhe, Phys. Rev. Lett. **82**, 3198 (1999).

[47] G. Gabrielse, D. Phillips, W. Quint, H. Kalinowsky, G. Rouleau, and W. Jhe, Phys. Rev. Lett. **74**, 3544 (1995).

[48] 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).

[49] E. A. Cornell, R. M. Weisskoff, K. R. Boyce, and D. E. Pritchard, Phys. Rev. A **41**, 312 (1990).

[50] H. Häffner, T. Beier, S. Djekić, N. Hermanspahn, H.-J. Kluge, W. Quint, S. Stahl, J. Verdú, T. Valenzuela, and G. Werth, Eur. Phys. J. D **22**, 163 (2003).

[51] S. Ulmer, K. Blaum, H. Kracke, A. Mooser, W. Quint, C. C. Rodegheri, and J. Walz, Phys. Rev. Lett. **107**, 103002 (2011).

[52] L. S. Brown, Ann. Phys. **159**, 62 (1985).

[53] G. Gabrielse, X. Fei, L. A. Orozco, S. L. Rolston, R. L. Tjoelker, T. A. Trainor, J. Haas, H. Kalinowsky, and W. Kells, Phys. Rev. A **40**, 481 (1989).

[54] 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).

[55] G. Gabrielse, Adv. At. Mol. Opt. Phys. **50**, 155 (2005).

[56] James F. Ziegler, M.D. Ziegler, and J.P. Biersack, Nuc. Inst. and Meth. B **268**, 1818 (2010).

[57] P. Yesley, Ph.D. thesis, Harvard University, 2001.

[58] G. Gabrielse, N. S. Bowden, P. Oxley, A. Speck, C. H. Storry, J. N. Tan, M. Wessels, D. Grzonka, W. Oelert, G. Schepers, T. Sefzick, J. Walz, H. Pittner, T. W. Hänsch, and E. A. Hessels, Phys. Rev. Lett. **89**, 213401 (2002).

[59] X. Fei, Ph.D. thesis, Harvard University, 1990.

[60] J. Estrada, Ph.D. thesis, Massachusetts Institute of Technology, 2002.

[61] H. Häffner, T. Beier, N. Hermanspahn, H.-J. Kluge, W. Quint, S. Stahl, J. Verdú, and G. Werth, Phys. Rev. Lett. **85**, 5308 (2000).

[62] K. Marable, Ph.D. thesis, Harvard University, 2019.

[63] G. Gabrielse, H. Dehmelt, and W. Kells, Phys. Rev. Lett. **54**, 537 (1985).

[64] C. H. Tseng, D. Enzer, G. Gabrielse, and F. L. Walls, Phys. Rev. A **59**, 2094 (1999).

[65] C.J Hardy, W.A Edelstein, and D Vatis, J. Magn. Reson. **66**, 470 (1986).

[66] B. Odom, Ph.D. thesis, Harvard University, 2011.

[67] N. F. Ramsey, Phys. Rev. **78**, 695 (1950).

[68] E.A. Cornell, R. M. Weisskoff, K. R. Boyce, R. W. Flanagan Jr., G. P. Lafyatis, and D. E. Pritchard, Phys. Rev. Lett. **63**, 1674 (1989).

[69] Brianna J. Mount, Matthew Redshaw, and Edmund G. Myers, Phys. Rev. A **82**, 042513 (2010).

[70] S. Sturm, A. Wagner, B. Schabinger, and K. Blaum, Phys. Rev. Lett. **107**, 143003 (2011).

[71] F. Heisse, F. Kohler-Langes, S. Rau, J. Hou, S. Junck, A. Kracke, A. Mooser, W. Quint, S. Ulmer, G. Werth, and K. Blaum, Phys. Rev. Lett. **119**, 033001 (2017).

[72] V. Natarajan, K. R. Boyce, F. Difilippo, and D. E. Pritchard, Phys. Rev. Lett. **71**, 1998 (1993).

[73] V. Natarajan, Ph.D. thesis, MIT, 1993.

[74] D. Hanneke, Ph.D. thesis, Harvard University, 2007.

[75] J. A. Sedlacek, J. Stuart, W. Loh, R. McConnell, C. D. Bruzewicz, J. Sage, and J. Chiaverini, J. App. Phys. **124**, 214904 (2018).

[76] M. J. Borchert, P. E. Blessing, J. A. Devlin, J. A. Harrington, T. Higuchi, J. Morgner, C. Smorra, E. Wursten, M. Bohman, M. Wiesinger, A. Mooser, K. Blaum, Y. Matsuda, C. Ospelkaus, W. Quint, J. Walz, Y. Yamazaki, and S. Ulmer, Phys. Rev. Lett. **122**, 043201 (2019).

[77] Jaroslaw Labaziewicz, Yufei Ge, Paul Antohi, David Leibrandt, Kenneth R. Brown, and Isaac L. Chuang, Phys. Rev. Lett. **100**, 013001 (2008).

[78] Guang Hao Low, Peter F. Herskind, and Isaac L. Chuang, Phys. Rev. A **84**, 053425 (2011).

[79] N. Daniilidis, S. Narayanan, S. A. Möller, R. Clark, T. E. Lee, P. J. Leek, A. Wallraff, St. Schulz, F. Schmidt-Kaler, and H. Häffner, New Journal of Physics **13**, 013032 (2011).

[80] R. McConnell, C. D. Bruzewicz, J. Chiaverini, and J. Sage, Phys. Rev. A **92**, 020302 (2015).

[81] J. A. Sedlacek, J. Stuart, D. H. Slichter, C. D. Bruzewicz, R. McConnell, J. Sage, and J. Chiaverini, Phys. Rev. A **98**, 063430 (2018).

[82] J. A. Sedlacek, A. Greene, J. Stuart, R. McConnell, C. D. Bruzewicz, J. Sage, and J. Chiaverini, Phys. Rev. A **97**, 020302 (2018).

[83] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland, Phys. Rev. Lett. **109**, 103001 (2012).

[84] J. Chiaverini and J. Sage, Phys. Rev. A **89**, 012318 (2014).

[85] E. Myers, Private communication.

[86] L. D. Brown, T. T. Cai, and A. DasGupta, Stat. Sci. **16**, 101 (2001).

[87] T. Kaltenbacher, Ph.D. thesis, Graz University of Technology, 2013.

[88] G. Gabrielse, S. E. Fayer, T. G. Myers, and X. Fan, Atoms **7**, (2019).

[89] W. Quint and G. Gabrielse, Hyperfine Interact. **76**, 379 (1993).

[90] L. S. Brown, Phys. Rev. Lett. **52**, 2013 (1984).

[91] F. Bloch and A. Siegert, Phys. Rev. **57**, 522 (1940).

[92] Andrew M. Childs, David Gosset, and Zak Webb, Science **339**, 791 (2013).

[93] Salvador Elías Venegas-Andraca, Quantum Information Processing **11**, 1015 (2012).

[94] Y. Aharanov, L. Davidovich, and N. Zagury, Phys. Rev. A .

[95] A. Ambainis, E. Bach, A. Nayak, A. Vishwanath, and J. Watrous, in *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, *STOC '01* (ACM, New York, NY, USA, 2001), pp. 37–49.

[96] Hilary A Carteret, Mourad E H Ismail, and Bruce Richmond, Journal of Physics A: Mathematical and General **36**, 8775 (2003).

[97] B. C. Travaglione and G. J. Milburn, Phys. Rev. A **65**, 032310 (2002).

[98] Antoni Wójcik, Tomasz Łuczak, Paweł Kurzyński, Andrzej Grudka, and Małgorzata Bednarska, Phys. Rev. Lett. **93**, 180601 (2004).

[99] T. Pertsch, P. Dannberg, W. Elflein, A. Bräuer, and F. Lederer, Phys. Rev. Lett. **83**, 4752 (1999).

[100] 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).

[101] D. W. Knight.

[102] R. A. Woode, E. N. Ivanov, M. E. Tobar, and D. G. Blair, Electronics Letters **30**, 2120 (1994).

[103] Khalid Z. Rajab, Mira Naftaly, Edmund H. Linfield, Juan C. Nino, Daniel Arenas, David Tanner, Raj Mittra, and Michael Lanagan, Journal of Microelectronics and Electronic Packaging **5**, 2 (2008).

[104] S. Ulmer, H. Kracke, K. Blaum, S. Kreim, A. Mooser, W. Quint, C. C. Rodegheri, and J. Walz, Rev. Sci. Instrum. **80**, 123302 (2009).

[105] M. W. Pospieszalski, S. Weinreb, R. D. Norrod, and R. Harris, IEEE T. Microw. Theory **36**, 552 (1988).

[106] R. M. Weisskoff, G. P. Lafyatis, K. R. Boyce, E. A. Cornell, R. W. Flanagan Jr., and D. E. Pritchard, J. Appl. Phys. **63**, 4599 (1988).

[107] E. Novitski, Ph.D. thesis, Harvard University, 2017.

[108] W. W. Macalpine and R. O. Schildknecht, Proc. IRE **47**, 2099 (1959).

[109] X. Fan, Ph.D. thesis, Harvard University, Forthcoming.

[110] C. Smorra, A. Mooser, K. Francke, H. Nagahama, G. Schneider, T. Higuchi, S.V. Gorp, K. Blaum, Y. Matsuda, W. Quint, J. Walz, Y. Yamazaki, and S. Ulmer, **358**, 1081 (2017).

[111] Jae Bon Koo and Duk Yong Yoon, Metallurgical and Materials Transactions A **32**, 1911 (2001).

[112] 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).

[113] J. W. Ekin, *"Experimental Techniques for Low-Temperature Measurements"* (Oxford, New York, 2006).

[114] C. Y. Ho, R. W. Powell, and P. E. Liley, *Thermal Conductivity of Selected Materials Part 2* (National Bureau of Standards, ADDRESS, 1968).

[115] G. K. White and S. B. Woods, Phil. Trans. Roy. Soc. Lond. Ser. A **21**, 273 (1959).

[116] P. K. Gupta and R. Rabehl, Applied Thermal Engineering **84**, 104 (2015).

[117] MTM Inc.

[118] L. S. Brown, G. Gabrielse, J. N. Tan, and K. C. D. Chan, Phys. Rev. A **37**, 4163 (1988).

[119] D. Hanneke, S. Fogwell Hoogerheide, and G. Gabrielse, Phys. Rev. A **83**, 052122 (2011).

[120] S. Ulmer, C. C. Rodegheri, K. Blaum, H. Kracke, A. Mooser, W. Quint, and J. Walz, Phys. Rev. Lett. **106**, 253001 (2011).

[121] A. Mooser, S. Ulmer, K. Blaum, K. Francke, H. Kracke, C. Leiteritz, W. Quint, C. C. Rodegheri, C. Smorra, and J. Walz, Nature **509**, 596 (2014).

[122] Christopher Coleman, *An Introduction to Radio Frequency Engineering* (Cambridge University Press, Cambridge, 2004).

[123] D. B. Pinegar, Ph.D. thesis, University of Washington, 2007.

[124] M. Lampton and F. Paresce, Rev. Sci. Instrum. **45**, 1098 (1974).

[125] E. Liénard, M. Herbane, G. Ban, G. Darius, P. Delahaye, D. Durand, X. Fléchard, M. Labalme, F. Mauger, A. Mery, O. Naviliat-Cuncic, and D. Rodríguez, Nucl. Instrum. Methods Phys. Res., Sect. A **555**, 375 (1974).

[126] J. T. Walton, H. A. Sommer, G. J. Wozniak, G. F. Peaslee, D. R. Bowman, W. L. Kehoe, and A. Moroni, IEEE T. on Nucl. Sci. **37**, 1578 (1990).

[127] H. J. Woltring, IEEE T. on Electron Devices **22**, 581 (1975).

[128] M. Lampton and C. W. Carlson, Rev. Sci. Instrum. **50**, 1093 (1979).

[129] S. Cui and Y. C. Soh, IEEE T. on Electron Devices **57**, 2310 (2010).

[130] A. Ionescu, Ph.D. thesis, Northwestern University, Forthcoming.

[131] T. Kaltenbacher, F. Caspers, M. Doser, A. Kellerbauer, and W. Pribyl, Nucl. Instrum. Methods Phys. Res., Sect. A **729**, 762 (2013).

[132] G. C. Firth and V. E. Watkins, An Interim Report on Investigation of Low-Temperature Solders for Cryogenic Wind Tunnel Models, https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19860001763.pdf, 1986.

[133] S. Stahl, Private communication.

[134] James Simpson, John Lane, Christopher Immer, and Robert Youngquist, https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20010038494.pdf, 2001.