Developing the Precision Magnetic Field for the E989 Muon g–2 Experiment

Matthias W. Smith

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2017

Reading Committee:

David Hertzog, Chair

Blayne Heckel

Steven Sharpe

Program Authorized to Offer Degree: Physics

University of Washington

Abstract

Developing the Precision Magnetic Field for the E989 Muon g–2 Experiment

Matthias W. Smith

Chair of the Supervisory Committee:
David Hertzog
Physics

The experimental value of $(g–2)_\mu$ historically has been and contemporarily remains an important probe into the Standard Model and proposed extensions. Previous measurements of $(g–2)_\mu$ exhibit a persistent statistical tension with calculations using the Standard Model implying that the theory may be incomplete and constraining possible extensions. The Fermilab Muon g-2 experiment, E989, endeavors to increase the precision over previous experiments by a factor of four and probe more deeply into the tension with the Standard Model. The $(g–2)_\mu$ experimental implementation measures two spin precession frequencies defined by the magnetic field, proton precession and muon precession. The value of $(g–2)_\mu$ is derived from a relationship between the two frequencies. The precision of magnetic field measurements and the overall magnetic field uniformity achieved over the muon storage volume are then two undeniably important aspects of the experiment in minimizing uncertainty. The current thesis details the methods employed to achieve magnetic field goals and results of the effort.

TABLE OF CONTENTS

Page

ACKNOWLEDGMENTS

I thank my parents for stoking my scientific curiosity and consistent love and support. Many thanks to my teachers, professors, mentors, and advisors for funneling that scientific curiosity toward the tangible, lofty goals which help push the boundaries of human knowledge. And, I would like to thank all of my friends for the special blend of support and reprieve that they afforded.

Chapter 1

INTRODUCTION

In particle physics, g–2 of the muon is an important and deep probe into the Standard Model (SM), because it can be both measured and calculated with very high precision. The Fermilab g–2 experiment, E989, will measure the anomalous magnetic moment of the muon to an overall precision of 140 ppb. Parallel efforts in the particle physics theory community aim to calculate $(g\text{–}2)_\mu$ to a comparable precision.

What is g–2? The quantity g–2, well $(g\text{–}2)/2$, is the anomalous magnetic moment of a lepton,

$$a_\ell = \left(\frac{g\text{–}2}{2}\right)_\ell . \tag{1.1}$$

The qualifier "anomalous" denotes a deviation from the value, $g \equiv 2$, in Dirac's theory of relativistic quantum mechanics. The anomaly arises from off-shell interactions with particles in the quantum vacuum. Such interactions are included in the modern theory of particle physics, which is based on quantum field theory. The anomaly manifests as a small enhancement in the coupling of a leptons's spin vector to an external magnetic field. The spin vector of a particle in a magnetic field precesses according to

$$\omega_s = g\frac{eB}{2m} \tag{1.2}$$

where g is the g-factor, e is the magnitude of electron charge, B is the magnetic field strength, and m is the mass of the particle. The g-factor is a particle-dependent constant that enters in the proportionality between the magnetic dipole moment, μ, and the spin vector, $\vec{S}$:

$$\vec{\mu} = g\frac{e}{2m}\vec{S}. \tag{1.3}$$

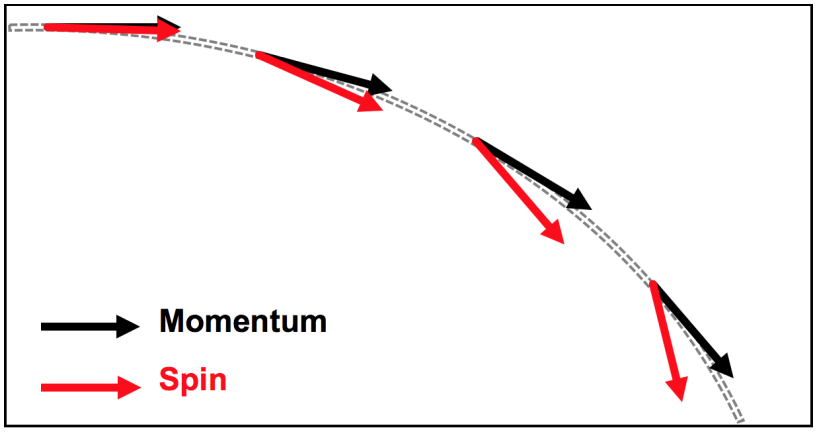

Figure 1.1: An illustration of the spin vector phase advance for muons trapped in a cyclotron. The momentum vector undergoes perfect cyclotron motion and the spin vector edges forward due specifically to the anomalous magnetic moment. The effect is exaggerated in the diagram.

The Fermilab g–2 experiment measures a_μ as the difference between two frequencies of muons in a magnetic storage field. Inside of the dipole storage magnet, muon spin vectors precess in a manner consistent with equation 1.2 and muon momentum vectors rotate at the cyclotron frequency of the magnetic field as given by

$$\omega_c = \frac{eB}{m_\mu}. \tag{1.4}$$

As muons circle around the storage magnet, the spin vector phase advances relative to the momentum vector at a rate exactly proportional to g–2, a concept illustrated in figure 1.1. E989 measures the difference between the precession frequency and the cyclotron frequency, the anomalous spin precession frequency, ω_a. As stated above and shown in equation 1.5, ω_a is directly proportional to the anomalous magnetic moment of the muon.

$$\omega_a = \omega_s - \omega_c = \left(\frac{g_\mu - 2}{2}\right)\frac{eB}{m_\mu} = a_\mu \frac{eB}{m} \tag{1.5}$$

Note: the two labels, a and g–2, are used somewhat interchangeably to denote the anomaly. Keep in mind that they technically differ by a factor of two.

1.1 Magnetic Dipole Moments

1.1.1 Definition

Classically, an entity attains a magnetic moment, $\vec{\mu}$, when the object has a specific correlation between the motion of its electric charge and the physical space that it occupies. Equation 1.6 illustrates how to calculate such a correlation for a body [17]. The analogous macro-scale quantity uses the term spin polarization or magnetization, $\vec{M}$, defined in equation 1.7.

$$\vec{\mu} = \frac{1}{2} \int \vec{r} \times \vec{j}\, dV \tag{1.6}$$

$$\vec{M} = \sum_i \vec{\mu}_i \tag{1.7}$$

Equation 1.6 can be re-framed in terms of angular momentum, $\vec{l}$, and charge density, ρ, by noting that $\vec{j} = \rho\, \vec{p}/m$ and $\vec{l} = \vec{r} \times \vec{p}$.

$$\vec{\mu} = \frac{g}{2m} \int \rho\, \vec{l}\, dV \tag{1.8}$$

Equation 1.8 works for classical angular momentum and orbital angular momentum $\vec{L}$, however, in the case of intrinsic spin angular momentum, $\vec{S}$, there is a caveat. A proportionality constant, the g-factor, must be included in the spin calculation of the equation giving rise to equation 1.9 for particle magnetic moments that arise due to spin.

$$\vec{\mu} = g \frac{e}{2m} \vec{S} \tag{1.9}$$

The g-factor, g, represents the coupling strength of a particle to the type of interactions mediated by magnetic fields. This quantity can vary extensively between different particles for instance, the proton g-factor is 5.5 [28], and the neutron has a g-factor of -3.8 [28]. It is worth noting that naively the neutron should not even have a magnetic moment due to a lack of net charge. For a point-like particle, the g-factor predicted by Dirac Equation is exactly 2. However, for charged leptons, the g-factor is slightly larger than 2, hence, the anomalous magnetic moment.

1.1.2 Properties

Alignment with External Fields

Intuitively, a magnetic moment is similar to a macroscopic magnet that nearly everyone is familiar with. The most prominent behavior of magnets is a tendency to align with external magnetic fields. Magnets align to minimize potential energy in the external field as given by equation 1.10.

$$U = -\vec{\mu} \cdot \vec{B} \tag{1.10}$$

Such behavior is exploited in compasses to align with the Earth's magnetic North Pole and seen with bar magnets where opposite ends attract. Though the examples are macro-scale, this behavior persists statistically in the micro-scale as well.

Spin Precession

Precession is another important but lesser known behavior of magnets. Precession actually arises from the same force which aligns a magnet in an external magnetic field. A more familiar analogue comes from the spinning top example. The child's toy receives a torque when the normal force on the top's tip and the ground is misaligned with the center of mass as illustrated in figure 1.2.

$$\vec{\tau} = \vec{r} \times \vec{F}_N \tag{1.11}$$

While the top is stabilized by rotational inertia, the induced torque causes the body to rotate about the top's tip. The magnetic moment case is realized by allowing the normal force to be replaced by the Lorentz force from the external magnetic field on the effective currents in an object with a non-zero magnetic moment, equation 1.12.

$$\vec{\tau} = \vec{\mu} \times \vec{B} \tag{1.12}$$

The induced torque causes the magnetic moment vector to rotate in plane orthogonal as prescribed by equation 1.13. The behavior is known as Larmor precession.

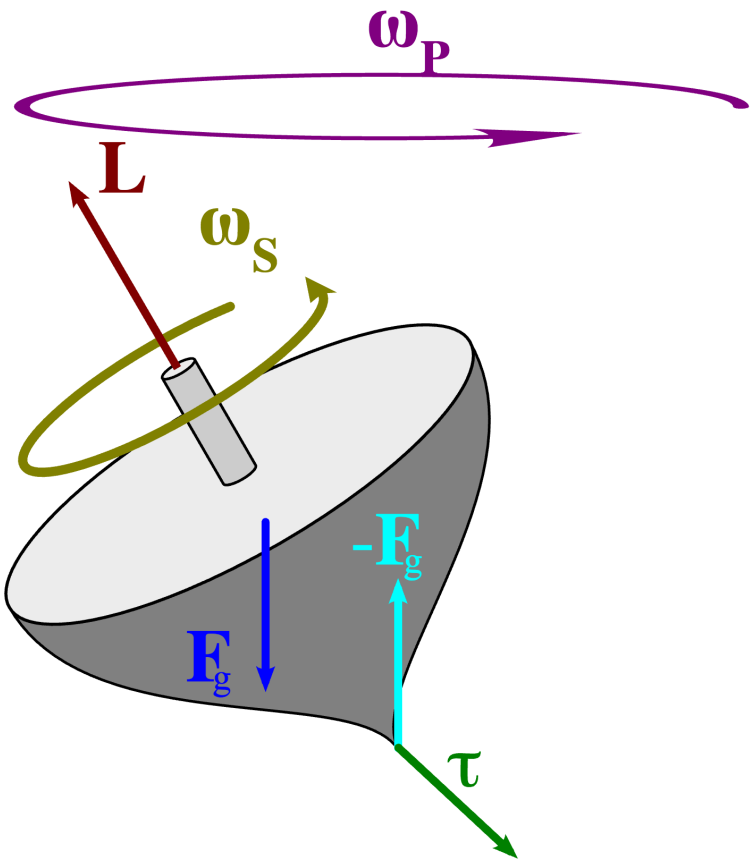

Figure 1.2: The spinning top example of precession as a force diagram. A torque arises from the off-center normal force, and that torque causes the top to precess at ω_p. The spinning top diagram was created by Xavier Snelgrove and found in the Wikipedia article for precession.

$$\omega_L = \frac{\mu B}{2m} \tag{1.13}$$

1.2 Muon Attributes

Muons are identical to electrons in most properties. Both particles are charged leptons and lepton universality posits that they are identical except for mass and the fact that muons are intrinsically unstable. Together with the tauon, they complete the charged lepton particle family. For a bit more context, examine table 1.2 which enumerates lepton properties compiled by the Particle Data Group [30].

Table 1.1: Lepton Properties

Property	e	μ	τ
charge	-1	-1	-1
spin	1/2	1/2	1/2
mass	$0.511\,\mathrm{MeV/c}$	$105.7\,\mathrm{MeV/c}$	$1776.9\,\mathrm{MeV/c}$
lifetime	$>2.0 \times 10^{36}\,\mathrm{s}$	$2.2 \times 10^{-6}\,\mathrm{s}$	$290.3 \times 10^{-15}\,\mathrm{s}$
primary decay	None	$\mu \to e\,\bar{\nu}_e\,\nu_\mu$	$\tau \to \mu\,\bar{\nu}_\mu\,\nu_\tau$
branch fraction	None	$\approx 100\%$	17.39%

1.2.1 Decay Chain

The pion decay chain is commonly used to produce muons. The principle branch from pion decay produces muons with very high probability, more than 99.9 % [30]. Similarly, the muon decays into an electron and two neutrinos with nearly 100 % probability [30]. The decay formulas are shown in equation 1.14.

$$\begin{aligned} \pi^+ &\to \mu^+ \bar{\nu}_\mu & \qquad \pi^- &\to \mu^- \nu_\mu \\ \mu^+ &\to e^+ \bar{\nu}_e \nu_\mu & \qquad \mu^- &\to e^- \nu_e \bar{\nu}_\mu \end{aligned} \tag{1.14}$$

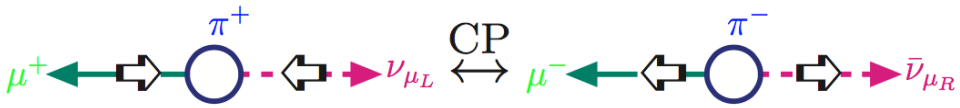

Figure 1.3: A diagram depicting pion decay in the rest frame. The spin vectors are shown in white and momentum vectors in color for the decay. Both the muon and the neutrino receive half the rest mass of the pion. The neutrino needs to have left-handed chirality, so both the neutrino and the muon are emitted with spin and momentum vectors anti-aligned. For the π^+ decay, the particles are emitted with right-handed chirality. Figure reproduced from reference [18].

Pion Decay

The two-body decay is isotropic in the pion rest frame and imparts 59.4 GeV of the pion rest energy to the (anti-)neutrino and 75.6 GeV to the muon. In the SM, neutrinos are always produced with left-handed chirality meaning that the direction of the spin vector and the momentum vector are anti-aligned. Conservation of angular momentum then forces the muon produced via pion decay to have the same chirality as the neutrino in the pion rest frame.

The decay lends itself to production of spin polarized muons. Here, polarization means alignment of the spin vector with the lab frame momentum vector rather than with the magnetic field as in the previous section. When the isotropic decay is boosted into the lab frame for a pion beam, both the highest energy (forward decay) or lowest energy (backward decay) muons exhibit strong spin polarization. As such momentum selection can be used to achieve a highly polarized muon beam.

Muon Decay

The muon decay is a bit more complicated than pion decay, since the process is three-body instead of two-body. A common framing for discussions of muon decay is the differential decay rate, $n(y)$, and asymmetry function, $a(y)$, over the range of possible muon energies, $y = E/E_{max}$ [14]. The differential decay rate, also called the Michel Spectrum, is a proxy for the overall probability of a decay electron with the energy y. The asymmetry function

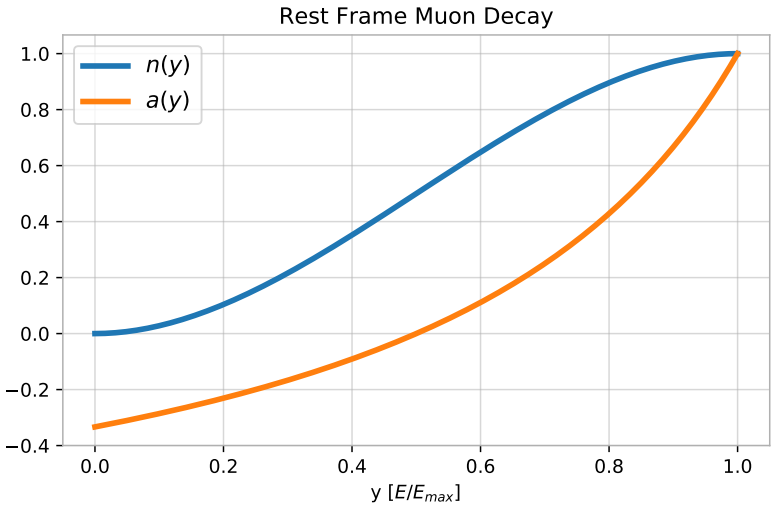

Figure 1.4: The plot depicts characteristics of the muon decay in the rest frame. The decay rate, $n(y)$, and the decay asymmetry, $a(y)$, are shown as a function of fractional energy, $y = E/E_{max}$.

represents the likelihood of the decay electron momentum being in the direction of the muon spin. Both functions are normalized to have a maximum value of one. The decay functions for the muon rest frame in equation 1.15, and the corresponding plots are shown in figure 1.4.

$$n_{rest}(y) = y^2(3 - 2y) \qquad a_{rest}(y) = \frac{2y - 1}{3 - 2y} \tag{1.15}$$

It is instructive to examine the extreme cases where both neutrinos are emitted in the same direction as shown in figure 1.5, and the case where they have opposing momentum directions.

The first case is the highest energy decay electron, the two neutrinos receive slight less than half the decay energy and the electron receives slightly more than half. The neutrino/anti-neutrino pair carry away no angular momentum, because the neutrino must have left-handed chirality and the anti-neutrino must have right-handed chirality, opposing spins. The decay electron then must carry the spin of the parent muon.

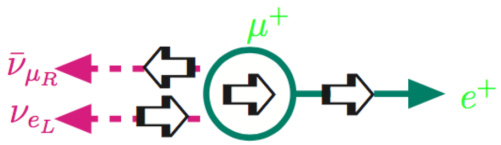

Figure 1.5: A diagram depicting the spin vectors in white and momentum vectors in color for muon decay. Together, the two neutrinos emitted in parallel carry away half the muon's rest mass in energy and no angular momentum. The electron then receives the other half of the muon's rest mass energy, and momentum in the direction of the muon's spin. Figure reproduced from reference [18].

The second case is the lowest energy decay electron, the neutrino and the anti-neutrino each receive half the rest mass of the muon. The electron is produced nearly at rest in this case. The angular momentum carried by the neutrinos in this case is no longer zero, but rather one. The electron then will be produced with spin in the opposite direction to the parent muon.

1.3 Motivation for $(g–2)_\mu$

There are two major questions that need answers to specifically motivate the muon g–2 experiment. The first question is: "Why measure the anomalous magnetic moment?", and the second question is: "Why use muons?".

Why g–2?

To the first question, there are many appropriate replies. One answer is that it is a simple measurement in principle. The experimenters need a spin polarized volume of leptons, an external magnetic field, and a technique to determine the spin direction. Another, more phenomenological answer is that contributions to g–2 come from the full gamut of fundamental interactions and constants of the universe in the SM: Quantum Electrodynamics (QED), Electro-Weak Interactions (EW), and Quantum Chromodynamics (QCD) all playing important roles. It traverses the entirety of Standard Model interactions aside from gravity. Precision measurements of g–2 are somewhat unique in this regard. Couple that with the fact that g–2 is measurable and calculable to high precision, and it's no wonder

that g–2 has fascinated experimentalists and theorists for so many years.

Why muons?

And, to the second question, there are physics considerations and practical considerations. The physics considerations weigh in with the desire for sensitivity to small effects predicted by the Standard Model. A more precise measurement enhances the visibility of small effects much like a microscope enhances optically small characteristics. The electron g–2 has already been measured with around 2,000 times more precision [15], so why not just measure g_e even more precisely? Well, simply put the physics contributions for g_e and g_μ are just different. Higher order interactions with particles of larger masses contribute with mass suppression terms, $\propto (\frac{m_\ell}{M})^2$, or possibly an even higher power! These interactions mostly manifest in the EW and QCD sector. The relative mass ratio between the electron and the muon enhances the sensitivity to these terms by a factor of $(\frac{105.66}{0.511})^2 \approx 43,000$ [19]! Even though g_e can be measured to much higher precision, g_μ is still around 2,000 times more sensitive to large mass contributions. To that end, the electron g–2 is a deeper probe of QED and the muon g–2 is a more rounded probe of the SM.

The preceding argument raises a further question though: "Why not use the τ?". The answer is really a practical consideration; the ephemeral τ has a lifetime of only 0.29 fs compared to 2.2 µs for the μ [28]. The τ also has many decay modes rather than a single dominant one like the muon. While the τ particle would be appealing if a significant enough Lorentz boost could be achieved and a clean method to understand the decay modes was presented, an experiment using τ particles is simply not practical with current technology.

1.4 The History of Experiments

1.4.1 Basic Principles

What then, one might ask, are the bare necessities to perform the g–2 experiment? Number one is a reliable source of spin polarized muons, the higher the net polarization the better. Next, there must be a method to store the muons in an external magnetic field while they precess. Recall that the experiment measures a frequency, so longer storage time increases

the overall precision. Muon storage can be achieved through a cyclotron magnet as in older experiments or a magnetic storage ring as in many newer experiments. The third necessity is a method by which to measure the direction of the muon spin vector over time. Measuring the spin direction relies on the fact that spin information is transfered to the momentum of the decay electron. The detector could be as simple as a forward/backward electron sensitive detector that measures decay electrons. The last requirement is a method to characterize the magnetic field experienced by the muons as they precess. Keep these points in mind as the historical experiments are discussed in the following sections.

1.4.2 Anomaly Detection and Parity Violation

The notion that $g \neq 2$ had to be reconciled with experiment as measurement precision progressed and statistical tensions arose with the values predicted by theory. Deviations from a pure Dirac particle for leptons were first observed in hyperfine splitting measurements of several different isotopes by Kusch and Foley in 1947 [22]. The theory community quickly resolved the discrepancy when Julian Schwinger released his seminal paper introducing corrections from quantum field fluctuations in 1948 [32]. The paper describes a now standard QED calculation of the lowest order self-interaction term for leptons emitting and reabsorbing photons, see figure 1.6. The Schwinger term brought experiment and theory back into good agreement, an early triumph of QED:

$$a_{Schwinger} = \frac{\alpha_{QED}}{2\pi}. \tag{1.16}$$

The first detection of a magnetic moment anomaly from measurements of free electrons came later, and in muons much later. The first measurement of g_e was done in 1953 by H. R. Crane, et al. [26]. Subsequent theoretical calculations and supporting experimental measurements established g_e as the most precisely predicted and measured quantity in QED [15]. As for measuring g_μ, it wasn't clear how to properly control the polarization of the muons, until parity violation of the weak interaction was understood [24, 36]. With parity violation inherent in the weak interaction, researchers quickly realized that a pion beam would decay into muons polarized along the beam direction (see section 1.2 for more detail

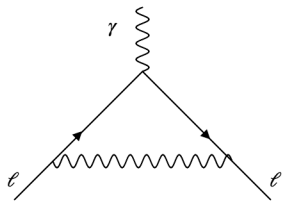

Figure 1.6: The Feynman diagram for the so-called Schwinger term. The propagating lepton radiates a photon, interacts with the magnetic field via a disconnected photon and reabsorbs the radiated photon. The contribution to a_μ from the interaction is $\frac{\alpha}{2\pi}$.

on pion decay).

1.4.3 CERN-I (1958-1962)

Building off the weak parity revelation, researchers measured g_μ for the first time. In 1960 Columbia personnel measured the quantity to 5% uncertainty, and not long after in 1961 a more precise measurement was performed at CERN. The experiment worked by injecting relatively low energy, 83 MeV, muons into a long, narrow magnetic trap where the polarized muons would undergo both cyclotron motion and lateral drift, see figure 1.7. At the opposite end of the magnet, the muons were ejected, tagged for storage time, and stopped where the momentum direction of the decay electron was recorded as either forward or backward. Using this technique CERN scientists were able to achieve a precision of 0.4% [7]. At the achieved precision, a deviation of $1.6\,\sigma$ was present between the measured value and the value predicted by the Schwinger term [10].

$$a_\mu = 116\,500\,000(500\,000) \times 10^{-11} \tag{1.17}$$

1.4.4 CERN-II (1962-1968)

The second iteration of $(g\text{–}2)_\mu$ at CERN improved by a factor of 15 over the first. CERN-II was the first $(g\text{–}2)_\mu$ experiment to use the now familiar magnetic storage ring as shown in figure 1.8. The storage ring design allowed the experimenters to measure a_μ directly as

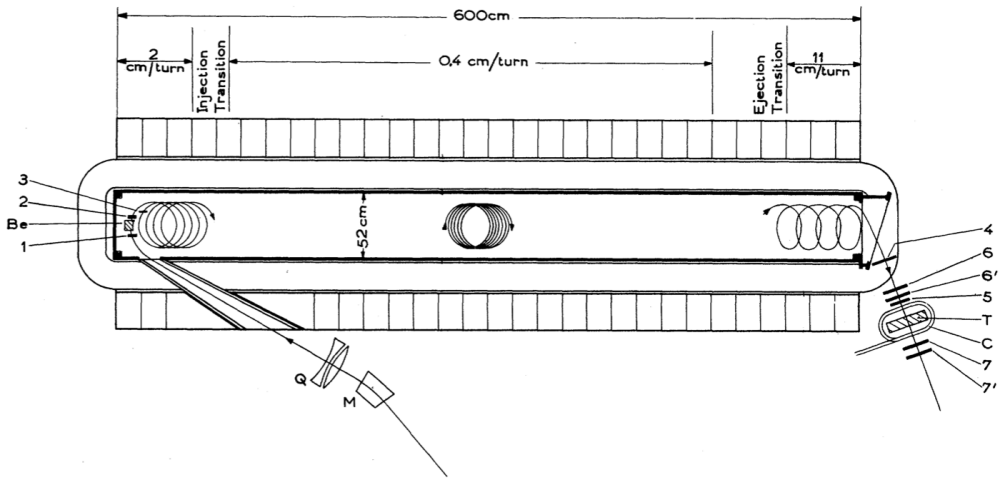

Figure 1.7: A diagram of the experimental setup in the first muon g–2 experiment at CERN from the original paper [7]. The muons enter from the lower left, go through the energy moderator to put the cyclotron radius at 19 cm, drift and circle toward the ejection side of the magnetic field, escape from the magnet, stop in the fiducial block, and decay into an electron with momentum correlated to the spin direction.

opposed to g_μ as explained near the beginning of the chapter. In order to store muons, the experiment injected a beam of protons which hit a pion production target. A slice of the pion production phase space matched the momentum acceptance of the ring well enough to remain for several revolutions. The stored muons were mostly forward decays that lost a bit of energy and matched the ring's momentum acceptance. The decay electrons curled inward to produce signals in electron counting detectors at a rate that modulated at the ω_a. The injected muons had a relativistic γ of 12 which allowed the researchers to increase the observation time of muon spin precession to more than 130 μs. The increased duration led to improved determination of a_μ to 270 ppm [10].

$$a_\mu = 116\,616\,000(31\,000) \times 10^{-11} \tag{1.18}$$

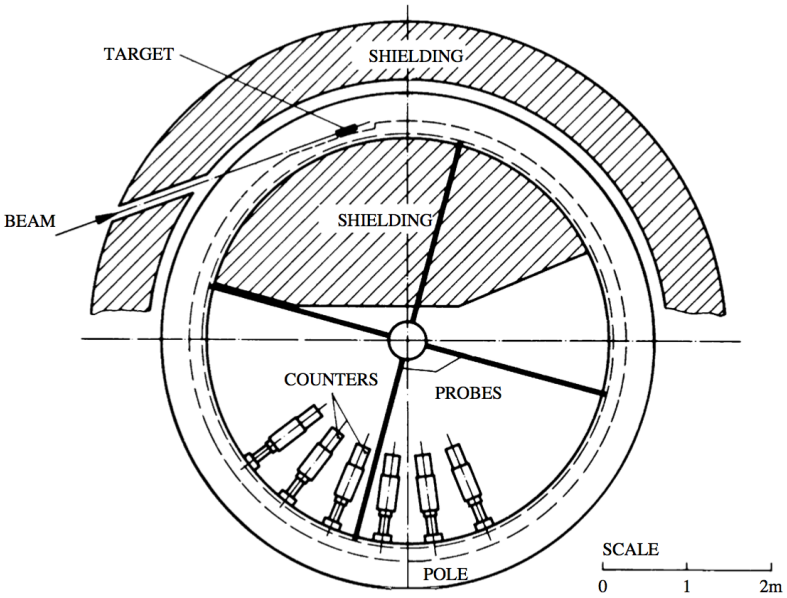

Figure 1.8: The first muon storage ring used in the CERN-II experiment as depicted in the muon g-2 review [10]. Pions were injected and a fraction of the forward decay muons matched the ring's momentum acceptance. Those polarized muons orbited the ring until decaying into electrons which activated the detectors. The electron rate signal was used to extract a_μ.

1.4.5 CERN-III (1969-1976)

The third CERN muon g–2 experiment was also the last. A major innovation introduced in CERN-III was use of the so-called "magic" momentum. Observe equation 1.19, the expression for spin precession in electromagnetic fields for a relativistic muon.

$$\vec{\omega}_a = \frac{e}{m}\left[a_\mu \vec{B} - \left(a_\mu - \frac{1}{\gamma^2 - 1}\right)\frac{\vec{\beta} \times \vec{E}}{c}\right] \tag{1.19}$$

A muon beam at a very specific momentum, 3.094 GeV/c, cancels the effects of radial electric fields which allows electrostatic focusing to be used on the muon beam instead of magnetic gradient focusing. The perturbations from magnetic focusing were a large source of uncertainty in the previous experiment. Another major innovation for the third CERN experiment was the use of pion injection instead of proton injection. The measurement techniques of CERN-III were similar to CERN-II. The new magnetic storage ring constructed for the experiment is shown in figure 1.9. With the achieved improvements, the CERN team was able to drive down the uncertainty on a_μ to 7 ppm [10], nearly a 40-fold improvement!

$$a_\mu = 116\,592\,300(800) \times 10^{-11} \tag{1.20}$$

1.4.6 E821 at BNL (1984-2003)

The most precise $(g\text{–}2)_\mu$ experiment to date took place at Brookhaven National Laboratory (BNL). The experiment, E821 as it is labeled in high energy physics ledgers, pushed precision muon physics to a new level. The precision goals for E821 demanded a 400-fold increase in statistics over its predecessor. The magnetic storage ring was redesigned once again as is shown in figure 1.10. One critical improvement in the experiment was that E821 injected muons rather than pions, or protons. Muon injection provided cleaner data earlier in each injection, and with an exponentially decaying number of muons, an early start time could lead to a huge rate improvement. The experiment design also focused on improving the homogeneity of the magnetic storage field. The aperture of the storage region was increased to facilitate a more uniform magnetic field across the muon storage volume. Field measurement was also improved by implementing both a suite of stationary magnetometers

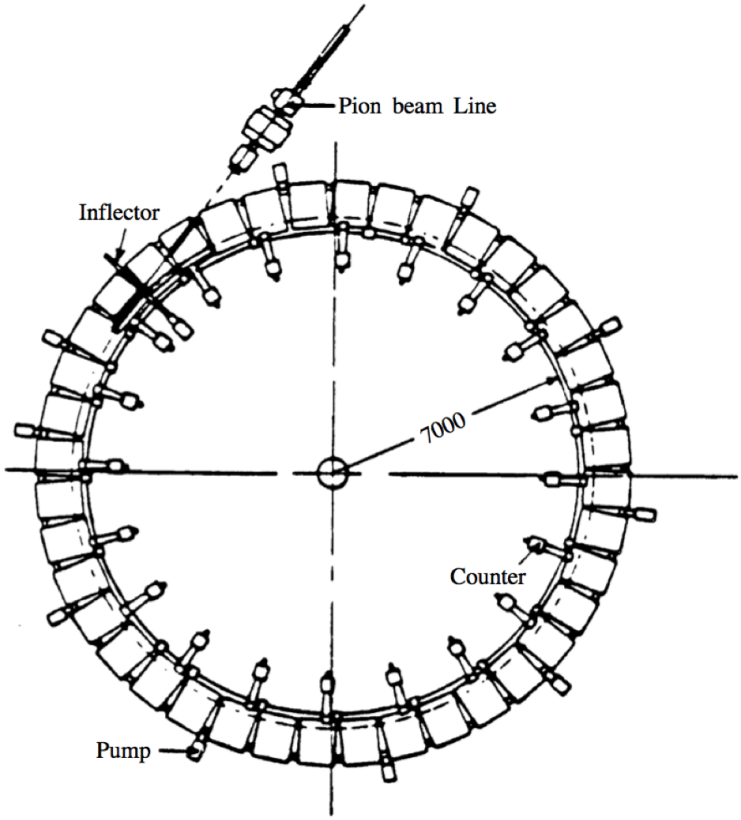

Figure 1.9: The second muon storage ring used in the CERN-III experiment. A radial line of 7000 mm is draw for scale. Pions were again injected, though finer momentum selection was used which yielded higher spin polarization for the muons. The electron counting detection method was similar to the technique used in CERN-II. Figure reproduced from reference [10].

outside of the storage region to continuously monitor magnetic field drift and a trolley outfitted with an array of 17 magnetometers to measure the field in the storage volume periodically.

In the end, the experiment nearly reached the initial goal of 350 ppb uncertainty on a_μ, actually achieving 540 ppb uncertainty [3]. The E821 g–2 result was at odds with theoretical calculations. Depending on which theoretical models were employed, the measurement was somewhere around $3.3\sigma - 3.6\sigma$ away from theoretical prediction, a statistical tension!

$$a_\mu = 116\,592\,082(55) \times 10^{-11} \tag{1.21}$$

1.5 The State of Theory

The theoretical contributions to $(g\text{–}2)_\mu$ come from all corners of particle physics. The typical division of contributions breaks them into QED, EW, QCD hadronic vacuum polarization (HVP), and QCD hadronic light-by-light (HLbL) as laid out in equation 1.22. The QED correction is by far the largest. It is well understood though, so the uncertainty on the correction is small. The next largest contributions to a_μ come from the QCD sector. The QCD terms also dominate the uncertainty on the calculation. The EW terms contribute a small, but well understood correction to a_μ. Figure 1.6.3 in section 1.6 is a visual overview of the corrections. Theoretical progress for QED and EW corrections involves calculating increasingly intricate Feynman diagrams. The values are well understood though, so there is no need for a major theoretical effort. The QCD contributions are trickier and represent a challenge to the theory community.

$$a_\mu = a_{QED} + a_{EW} + a_{QCD-HVP} + a_{QCD-HLbL} \tag{1.22}$$

1.5.1 QED Effects

The correction from QED interactions is the largest contribution to a_μ. Feynman diagrams provide a convenient way to intuit some of the relevant effects. All diagrams for QED are of the radiative correction type, shown in figure 1.11 (1–6). These diagrams involve

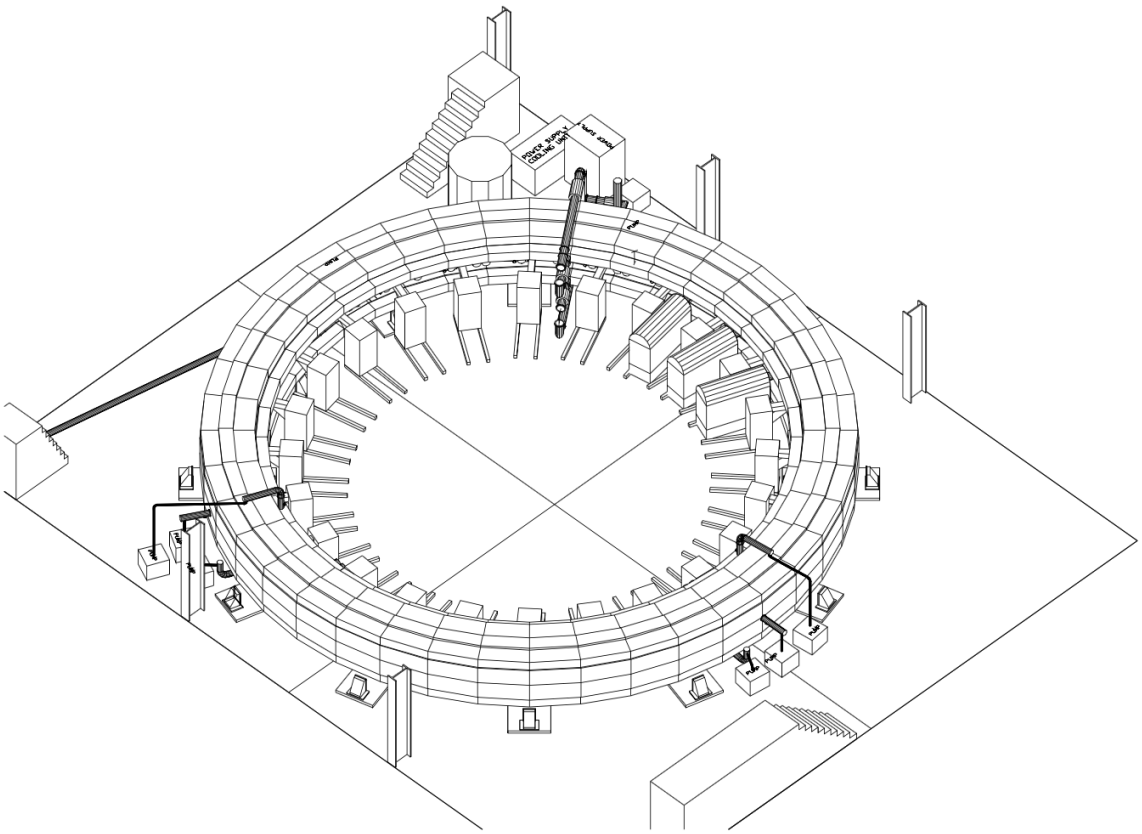

Figure 1.10: The magnetic storage ring used in the BNL experiment. The size of the ring was very similar to the CERN-II ring with a "magic" radius of 7112 mm for muon orbits. In this storage ring, muons were injected directly into the storage ring. Figure reproduced from reference [3].

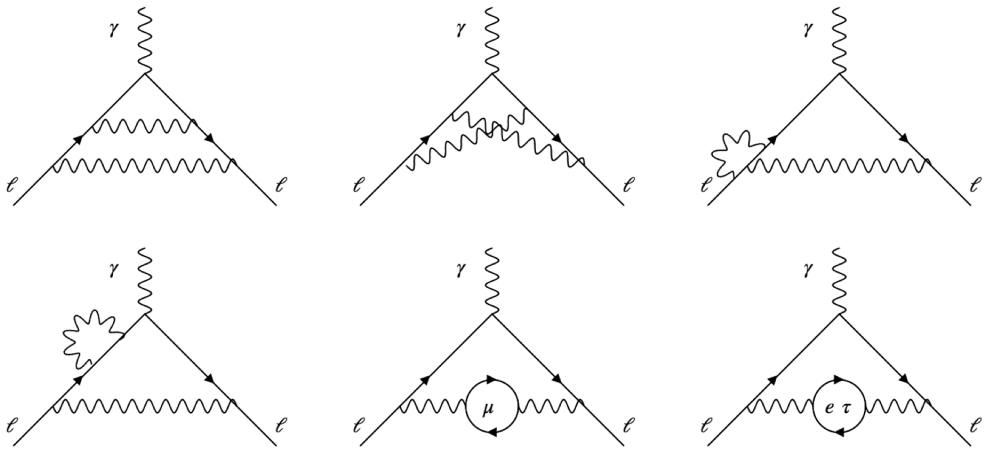

Figure 1.11: Several examples of QED interactions illustrated as Feynman diagrams. All six constitute radiative corrections which involve emission and absorption of photons. The last two are termed vacuum polarization effects, since the off-shell photon undergoes pair production.

emission and recapture of a photon(s). Another common type of diagram illustrates vacuum polarization interactions, shown in figure 1.11 (5–6), with the off-shell photon going through pair production and annihilation along its path.

The QED corrections constitute over 99 % of the total anomalous magnetic moment. The QED corrections lend themselves to a perturbative approach represented by a power series in the QED coupling constant, α_{QED}.

$$a_{QED} = \sum_n A_n \left(\frac{\alpha_{QED}}{\pi}\right)^n \tag{1.23}$$

The first order term in the series is the leptonically universal Schwinger term, $a_{Schwinger} = \frac{\alpha_{QED}}{2\pi}$. In terms of the measured value of a_e, the Schwinger term is 100.134% of the measured value. And for g_μ, the lowest order term is a slightly larger fraction of the total at 99.59% [28]. For the electron case, the higher order contributions are net negative which reduces the anomalous magnetic moment while further contributions are net positive for a_μ. In both cases though, the Schwinger term dominates the corrections.

$$a_e = 115\ 965\ 218.091(26) \times 10^{-11}$$

$$a_\mu = 116\ 592\ 080(63) \times 10^{-11}$$

$$a_{Schwinger} = 116\ 120\ 635.555(27) \times 10^{-11}$$

For higher order terms in the QED correction, one can gain some intuition by investigating a vertex expression in the different mass limits. First, consider the case of radiative corrections in which the radiated particles are light compared to the muon, $m_p \ll m_\mu$. Equation 1.24 is a general expression for 2nd order terms in the small mass limit [19] where n_p is the coupling order and $k_p < n_p$ is a possible enhancement to the logarithm term.

$$\delta a_\mu^p \sim \left(\frac{\alpha}{\pi}\right)^{n_p} \ln^{k_p}\left(\frac{m_\mu}{m_p}\right) \tag{1.24}$$

And, now consider the opposite case where $m_p \gg m$. The expression changes to equation 1.25.

$$\delta_\mu^p \sim \left(\frac{\alpha}{\pi}\right)^{n_p} \left(\frac{m_\mu}{m_p}\right)^2 \ln^{k_p}\left(\frac{m_\mu}{m_p}\right) \tag{1.25}$$

The mass suppression term can be understood as a result of the fact that these types of muon interactions with heavier particles requires helicity flips for the muon which add a $\frac{m_\mu}{m_p}$ factor and that vertex in the end is squared to yield an amplitude [1].

The two equations, 1.25 and 1.24, offer insight into the difference between muon and electron g–2. The electron mass is much smaller than the muon, so these are the type of higher order interactions that are enhanced in the muon g–2. Another major insight follows from the fact that electron interactions with large mass particles are suppressed. The electron g–2 can be used to calculate a very precise value of α_{QED} [15].

The contribution of QED effects have been calculated to 5 loops [2]. That includes more than 10,000 diagrams! The correction from QED is well known and does not limit the certainty on the theoretical value. The resulting contribution to $(g\text{–}2)_\mu$ in total is

$$a_\mu^{QED} = 116\ 584\ 718.845(0.037) \times 10^{-11}. \tag{1.26}$$

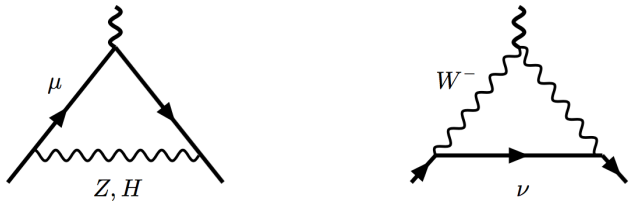

Figure 1.12: The largest contributing diagrams from the weak interaction. The muon in the left diagram goes off shell by emitting and re-absorbing a neutral Z or Higgs boson. And, the muon in the right diagram converts to a W^- (W^+ for μ^+) and emits and re-absorbs a muon neutrino to conserve fermion number. Figure reproduced from reference [1].

1.5.2 EW Effects

In general corrections due to the weak force are mass suppressed compared to QED corrections. The lowest order and largest contribution to the weak corrections comes two diagrams, see figure 1.12. One is similar to the Schwinger Diagram, the difference being that the photon propagator has been substituted for the Z boson. The second entails emission of a muon neutrino, conversion to a W boson of the appropriate charge, and recapture of the neutrino. The expression, equation 1.27 for the diagrams is calculated in [19] where the first term in brackets is derived for the W boson interaction and the second term for the Z boson. The total contribution for the lowest order weak corrections is then 194.9×10^{-11}.

$$a_\mu^{EW} = \frac{G_F m^2}{8\sqrt{2}\pi^2}\left[\frac{10}{3} + \frac{1}{3}\left(-5 + (1 - 4\sin\theta_W{}^2)^2\right)\right] \tag{1.27}$$

The next order in EW interactions might be expected to be nearly negligible. Naively, they would be suppressed by a coupling factor of $\frac{\alpha}{\pi} \approx 0.002$, and so would amount to a small effect. However, the contribution also gets an enhancement due to the large logarithm of $\log(m_Z/m) \approx 6.8$, and coincidentally many terms add coherently. More difficulty arises as QCD effects arise within the weak boson propagators, which will not receive more than a mention here. The total contribution from the second order weak diagrams ends up being -40×10^{-11} [19]. The total contribution to a_μ is given in equation 1.28.

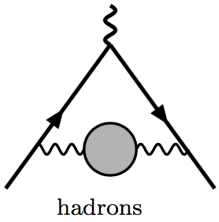

Figure 1.13: The basic diagram for hadronic vacuum polarization. Figure reproduced from reference [1].

$$a_\mu^{EW} = 154(2) \times 10^{-11} \tag{1.28}$$

1.5.3 QCD Effects

The QCD sector is undoubtedly the most difficult domain to calculate accurately and precisely in determining g–2 of the muon. In general QCD calculations can be extremely difficult owing to the non-perturbative nature of many QCD problems, and $(g\text{–}2)_\mu$ calculations are no exception. Some calculations leverage effective low-energy perturbation theories, others chiral perturbation approaches, and others still utilize lattice calculations [19]. The comprehensive approach involves both model-dependent calculations and model-independent contributions. Sorting the bevy of contributions is a daunting task, fortunately muon g–2 theory reviews have already done just that for experimentalists [19, 8, 6, 18].

Hadronic Vacuum Polarization

The general form of HVP is quite similar to the QED vacuum polarization described in section 1.5.1. The muon radiates a photon or emits another boson, and the newly created particle pair produces then annihilates before recapture with the muon, illustrated in figure 1.13. The difference being that in this case the pair production pulls from the hadronic sector instead of the lepton sector. Such possibilities include: π_0, π^+, π^-, ρ_0, etc. The size of the HVP correction is second only to the QED correction coming in at around 7000×10^{-11}.

The hadronic vacuum polarization can be anchored to measurements of the hadronic cross-section. The general method employed uses the integral relationship given as

$$a_\mu^{hvp} = \frac{\alpha}{3\pi}\int_{s_0}^{\infty}\frac{ds}{s}\frac{\sigma_{hadr}(s)}{\sigma_{point}(s)}a_\mu^{(1)}(s) \tag{1.29}$$

where $a_\mu^{(1)}(s)$ is the 1-loop contribution to a_μ from a neutral vector boson with mass $\sqrt{s}$, σ_{hadr} is the e^+e^- hadronic annihilation cross-section, and σ_{point} and further detail is given in reference [1]. With the data-driven method employed it is technically possible to extract the value exactly. The reality of the calculation turns out to be difficult, since data from many different experiment needs to be well understood and combined to cover the whole spectrum.

In a fairly recent theory review by Thomas Blum, et al [6], a recent value for the leading order and next order contributions combine to

$$a_\mu^{QCD-HVP} = 6824(42)\times 10^{-11}. \tag{1.30}$$

Hadronic Light-by-Light

The general form for hadronic light-by-light (HLbL) contains more interaction vertices than HVP and, as is to be expected, therefore a smaller contribution to the total muon anomaly. The core idea of hadronic light-by-light scattering is that the propagating the muon interacts with three photons and those photons interact with some sort of QCD loop which interacts with the external field. The HLbL scattering is one of the most difficult a_μ contributions to calculate, because it cannot be anchored to experimental data using a dispersion relation like the HVP contribution [19].

The value is largely model dependent for $a_\mu^{QCD-HLbL}$. The most discussed model in reference [1] uses a color extension allowing a large number of colors and includes constraints from chiral and short-distance QCD. Other model calculations arrive at values up to 50% different, but most are closer. Many groups working in the field convened to establish the "Glasgow Consensus" as an agreed upon value for the correction [14]. The consensus value is

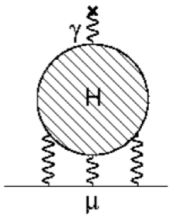

Figure 1.14: The basic diagram for hadronic light-by-light scattering. The propagating muon interacts with a QCD black box which interacts with the disconnected photon diagram, i.e., the external field. Figure reproduced from reference [6].

$$a_\mu^{QCD-HLbL} = 105(26) \times 10^{-11}. \tag{1.31}$$

1.5.4 Beyond the Standard Model

Consider the simplest extensions of the Standard Model where a new particle comes in analogous to a current particle. A general argument can be made for an approximation of the size of the correction from such diagrams [19]. For light particles where the new particle mass is less than the muon mass, $m_p < m_\mu$, the contribution is expected to have the following form:

$$\delta a_\mu \sim \left(\frac{\alpha}{\pi}\right)^{n_p} \left(\ln \frac{m_\mu}{m_p}\right)^{k_p} \tag{1.32}$$

where the exponent n_p represents the loop order of the correction and the exponent k_p is the possible logarithmic enhancement at the order. Note that $k_p < n_p$, but not otherwise predictable.

The situation changes a bit when more massive particles are considered, because helicity flips must be accounted for. In general those helicity flips add a unitless mass suppression term to the vertex function which is squared in the amplitude, so we get the following relation:

$$\delta a_\mu \sim \left(\frac{\alpha}{\pi}\right)^{n_p} \left(\frac{m_\mu}{m_p}\right)^2 \left(\ln \frac{m_\mu}{m_p}\right)^{k_p} \tag{1.33}$$

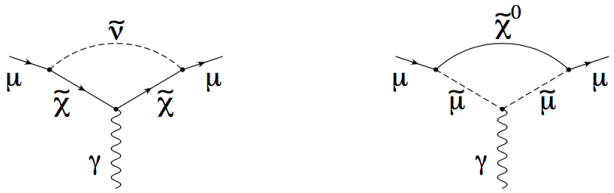

Figure 1.15: Two example diagrams of SUSY which would contribute to the anomalous magnetic moment. In the left diagram, the muon converts into the supersymmetric particles, the sneutrino, $\tilde{\nu}$ and chargino $\tilde{\chi}$, and back to the muon. In the right diagram the muon undergoes a different supersymmetric interaction vertex to become a smuon, $\tilde{\mu}$, and a neutralino, $\tilde{\chi}^0$. These one-loop diagrams are expected to be the largest contributors from minimal SUSY extensions of the SM. Figure reproduced from [8].

Supersymmetric Extensions

With these two general estimates established, it is instructive to talk about a few specific types of corrections. Supersymmetry (SUSY) is one proposed theory that can account for the deviation in a_μ with some tuning. Supersymmetric contributions arise through smuon-neutralino and sneutrino-chargino conversions in loop diagrams such as figure 1.15. A general expression for contributions from SUSY is given in equation 1.34 [8]. In equation 1.34, $\tilde{m}$ is the mass of the SUSY particle

$$\left|a_\mu^{SUSY}\right| \approx \left(\frac{\alpha(M_z)}{8\pi \sin^2\theta_W}\right)\left(\frac{m_\mu}{\tilde{m}}\right)^2 \tan\beta\left(1-\frac{4\alpha}{\pi}\ln\frac{\tilde{m}}{m_\mu}\right) \approx 130\times 10^{-11}\left(\frac{100GeV}{\tilde{m}}\right)^2 \tan\beta \tag{1.34}$$

The $\tan\beta$ enhancement can naturally be around 40 or 50. The difference $a_{expt} - a_{theory} \approx 250\times 10^{-11}$ puts the mass scale at $\tilde{m} \approx 500$GeV/c for a SUSY effect the size of the muon anomaly tension.

Radiative Extensions

Another proposed theory, radiative mass models, can explain both deviations from g–2 and the "unnaturally" light mass of the leptons. In this model the mass of the muon is generated by emission and absorption of an as yet unknown particle while it propagates, and the bare

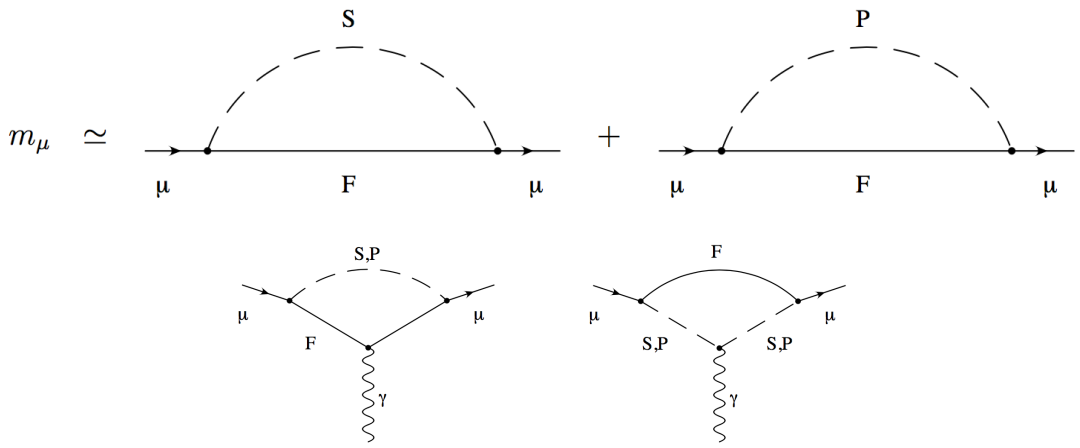

Figure 1.16: Example diagrams which would cause the mass of the muon and corrections to $(g–2)_\mu$. The letters represent possible new scalar (S), pseudoscalar (P), and fermion (F) particles. Figures reproduced from reference [8].

mass of the muon is zero. The same new particle that provides mass would come in to the anomalous magnetic moment as an unaccounted for Schwinger-like term. The size of the contributions in such a model are given in equation 1.32, and the diagrams in figure 1.16 [8].

Other Extensions

Many other SM extensions have been proposed throughout the years. See references [8, 19, 14] for deeper discussions, but a few are listed here. Dark photons are light, but very weakly coupling particles that could account for $(g–2)_\mu$. Anomalous properties of the W boson, such as an anomalous magnetic moment or electric quadrupole, have been proposed as possible explanations. Also, new gauge bosons models illustrate a possible avenue for anomalous g–2.

Much of the phase space for the proposed standard model extensions has been ruled out by other measurements. The LHC for instance has constrained the parameters in many SUSY models and other experiments have strongly restricted the possibility of dark photons. However, most models are not completely ruled out, and the assay of possible solutions will

need to be careful pruned and tuned by the theory community in the coming years.

1.6 Experiment Versus Theory

From the experimental side, the combined average value from all measurements comes to

$$a_{\mu,expt} = 116\,592\,082(55) \times 10^{-11}. \tag{1.35}$$

Similarly, all individual theory contributions combine to the value given in equation 1.36. **Note:** this value varies with different models, and the sum here is over the specific values listed in section 1.5.

$$a_{\mu,theory} = 116\,591\,802(49) \times 10^{-11} \tag{1.36}$$

The difference between the two values is then

$$a_{\mu,expt} - a_{\mu,theory} = 280(74). \tag{1.37}$$

The discrepancy between theory and experiment given above is $3.8\,\sigma$, and in the majority of models a discrepancy greater than $3\,\sigma$ persists. The next step is to push the precision even further on both fronts, and turn a statistical tension into a $>5\,\sigma$ discovery (assuming constant central values). The persistent discrepancy is a possible indication of new physics from beyond the Standard Model.

1.6.1 Improvements in Theory

All necessary theory work is in the QCD domain. Currently, the QCD-HVP uncertainty dominates the theory uncertainty. The value is expected to improve by adding additional cross-section data from Novosibirsk and BESIII experiments, and possibly through a lattice based approach [6].

The next largest source of uncertainty comes from the QCD-HLbL correction. There are two possibilities for reducing the uncertainty on the QCD-HLbL contribution. One method is to make measurements of hard photon physics which the value can be anchored to similar

to the QCD-HVP term [6]. Another possibility is calculation on the lattice which is already underway [4, 5].

The theory community aims to reduce the uncertainty on each contribution by a factor of two. If these precision goals are met, the uncertainty on the theoretical value would be 205 ppb.

1.6.2 Improvements in Experiment

The tension between past experiments and theory inspired a new wave of $(g–2)_\mu$ measurements. A new iteration of the $(g–2)_\mu$ experiment is underway at Fermi National Accelerator Laboratory (FNAL). The experiment, E989, reuses the storage ring, superconducting coils, and various components from the BNL experiment. The precision goal of E989 is an overall uncertainty of 140 ppb, likely pushing the tension to discovery levels. Additionally, a sister experiment has been proposed at J-PARC as a completely independent measurement of $(g–2)_\mu$ with similar sensitivity to the BNL experiment.

1.6.3 Summary

The major values related to $(g–2)_\mu$ are brought together for comparison in figure 1.6.3. The three most recent experimental values are shown as vertical bars with the projected value from E989 also included as a dashed vertical line. The theory contributions are grouped into QCD, EW, and QED sectors and split by order with the value shown in blue and the uncertainty shown in blue.

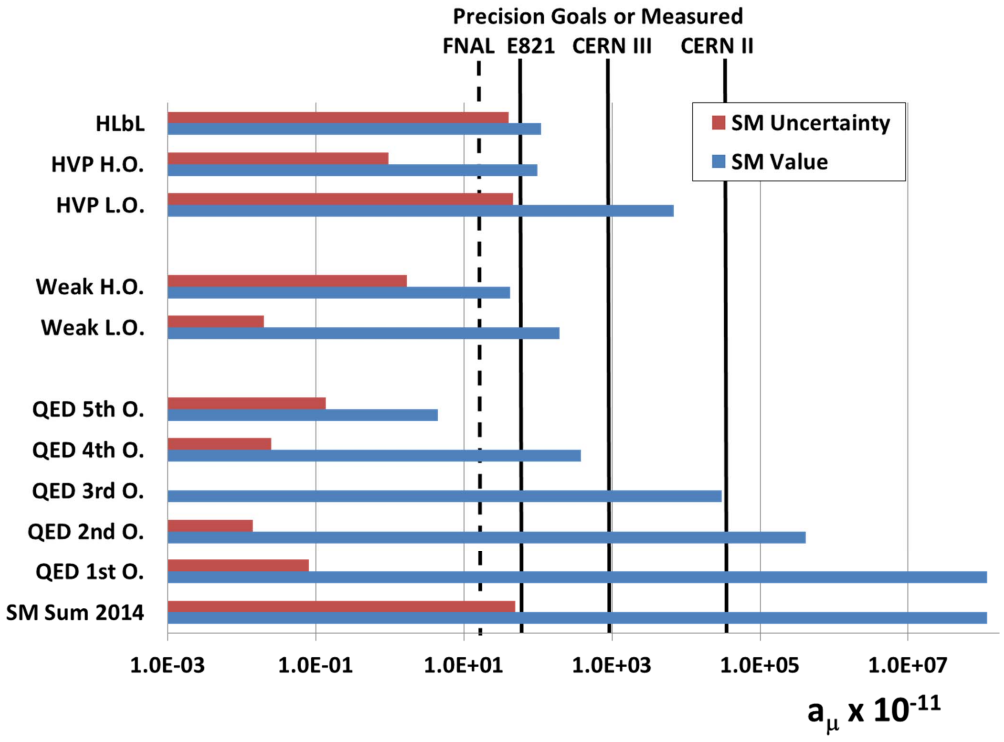

Figure 1.17: The visual depicts all different theoretical contributions to $(g–2)_\mu$ with experimental values for comparison. The value from each storage ring experiment is represented as a vertical line with the projected precision for E989 as a dashed line. The size and uncertainty of various SM corrections grouped by interaction type are depicted as bars extending onto the vertical axis. The most important point highlighted by the figure is that theory only needs to improve in the hadronic light-by-light and hadronic vacuum polarization sectors. Figure reproduced from reference [16].

Chapter 2

E989: THE MUON g–2 EXPERIMENT

2.1 Overview

2.1.1 The Muon Anomaly

The principal result from the E989 experiment can be represented by a single expression in which the two frequency parameters are quantities measured very precisely by the $(g–2)_\mu$ experiment and all other parameters are fundamental quantities precisely measured in other experiments.

$$a_\mu = \frac{\omega_a}{\omega_p}\frac{\mu_p}{\mu_e}\frac{m_\mu}{m_e}\frac{g_e}{2} \tag{2.1}$$

The parameter ω_a is the anomalous spin precession frequency for the muon, and ω_p is the Larmor frequency for a free proton. Both frequencies are induced by the same magnetic field and thereby exactly correlated. The analysis technique for ω_a is described in more detail in section 2.6. The ω_p analysis methods are further discussed in section 2.5 and further still in chapter 3. The CODATA values [28] for the magnetic moment and mass ratios and reference [15] for the g-factor of the electron are given below.

$$\mu_p/\mu_e = -658.210\ 684\ 8(54) \tag{2.2}$$

$$m_\mu/m_e = 206.768\ 284\ 3(52) \tag{2.3}$$

$$g_e/2 = 1.001\ 159\ 652\ 180\ 73(28) \tag{2.4}$$

$$\tag{2.5}$$

An alternative expression for a_μ can be derived by using the relation $m_\mu/m_e = (g_\mu\mu_\mu)/(g_e/\mu_e)$.

$$a_\mu = \frac{\mathcal{R}}{\lambda^+ - \mathcal{R}} \tag{2.6}$$

where $\mathcal{R} = \omega_a/\omega_p$ and $\lambda^+ = \mu_{\mu^+}/\mu_p$ which is measured in a muonium hyperfine splitting experiment [25]. Using the CODATA values [28], the ratio computes to

$$\lambda^+ = 3.183\ 345\ 139(10). \tag{2.7}$$

2.1.2 Precision Goals

The overall precision goal on a_μ is 140 ppb. 100 ppb of uncertainty is allotted for statistics. The systematic uncertainty limit on ω_a is 70 ppb, and the systematic uncertainty limit on ω_p is equally 70 ppb. The individual terms summed in quadrature make up the entire uncertainty budget for the FNAL $(g\text{–}2)_\mu$ experiment, nearly a four-fold improvement upon the precision of the previous experiment at BNL. The statistical uncertainty goal can be achieved through recording 1.8×10^{11} physics events [14]. The plan to reach systematic uncertainty limits includes broad experimental improvements over E821 in hardware systems, data quality, and analysis techniques. A subset of the improvement plans for the ω_p uncertainties is discussed in chapter 3.

2.1.3 Stages of the Experiment

The $(g\text{–}2)_\mu$ experiment can be broken into several stages. In the subsequent sections the following stages are discussed:

- Muon Production
- The Magnetic Storage Ring
- Muon Injection
- The Magnetic Field
- Spin Precession

In this document, a brief overview of information in the E989 Technical Design Report (TDR) is given. See the original document, reference [14], for further detail.

2.2 Muon Production

The first stage in an experiment that measures a property of muons is of course producing muons. The full chain of muon production goes from protons to pions to muons (to electrons). Expert personnel at Fermilab have developed a new beamline to deliver muons to the g–2 experiment using much of the existing particle production and accelerator infrastructure. The structure of the beamline is illustrated in figure 2.1.

Much of the FNAL accelerator infrastructure is reused for protons in new beamlines for the Muon Campus. The protons begin in the Linac and accelerate through Booster. From there, the protons continue into the Recycler Ring where they are manipulated into the high intensity bunches with short timing structure for g–2. Each proton bunch contains $\mathcal{O}(10^{12})$ protons with 8 GeV kinetic energy in time windows less than 90 ns.

The protons then propagate from the Recycler to the AP0 target where they collide with the target. The collision produces $\mathcal{O}(10^{9})$ positive secondary particles of which many are pions. The secondary particles are focused via an electrostatic lithium lens into a secondary beam which goes through a momentum filter shortly after focusing. Momentum selection yields a beam of 3.1 GeV/c with a momentum spread of $\pm 0.10\,\frac{\delta \mathrm{p}}{\mathrm{p}}$. The secondary beam then proceeds through P1, P2, M1, M2, and M3 beamlines into the Delivery Ring.

The goals in the Delivery Ring are twofold. First, the beam cycles around the Delivery Ring to create a spatial separation between the pions/muons and the more massive protons (slightly lower velocity for the same momentum), so that the protons can be removed. Secondly, essentially all pions undergo in-flight weak decay into muons, so the outgoing beam is a very pure muon beam. Four orbits around the Delivery Ring are enough to achieve both goals.

After the Delivery Ring, the now muon beam is extracted onto the path toward the $(g\text{–}2)_\mu$ storage ring. Through the pion decay process the high and low energy muons have a net spin polarization (as discussed in section 1.2), and the beamline design acceptance is narrow around the filtered secondary energy of 3.1 GeV/c. The muons produced at 3.094 GeV/c by the pion beam are forward decays and thereby achieve a net spin polarization of around 95 %. A bunch of muons produced in the beamline is referred to as a "fill". These

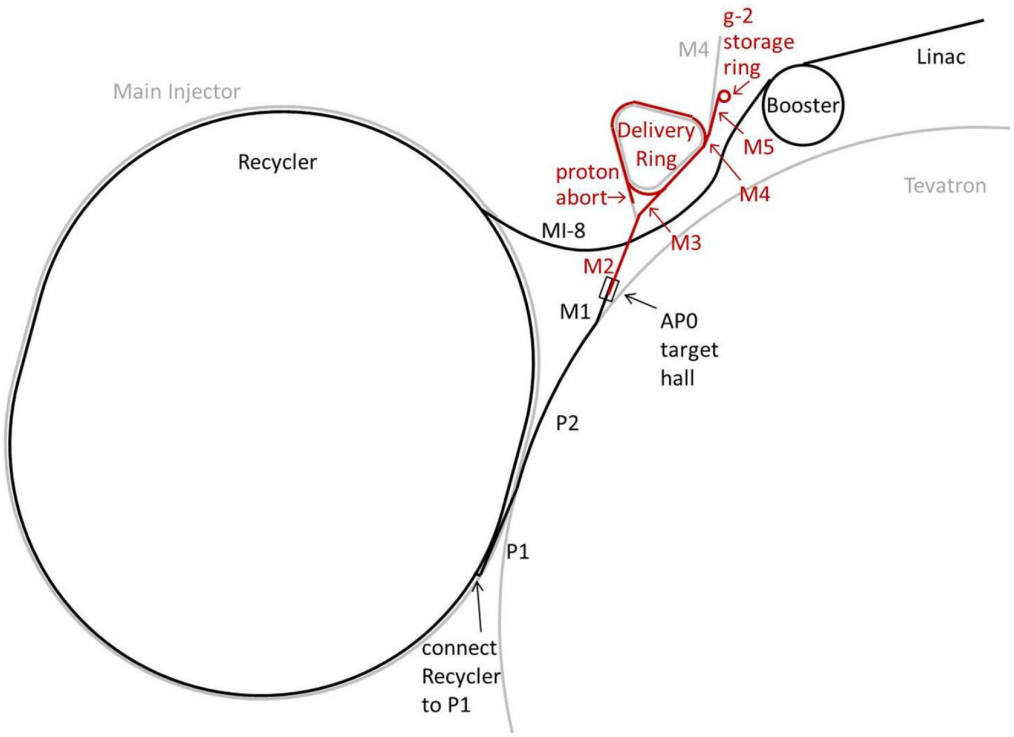

Figure 2.1: The diagram depicts all relevant FNAL beamlines for g–2. Protons begin accelerating in the Linac, continue in the Booster, and enter the Recycler Ring. In the Recycler the protons are bunched into high intensity, small time windowed groups. The protons exit the Recycler and propagate down the P1, P2, and M1 beamlines to the secondary production target. Positive secondaries (including pions which later yield muons) at 3.1 GeV/c are focused and transported down the M2 and M3 beamlines to the Delivery Ring. In the Delivery Ring, the bunch propagates long enough to develop a timing separation between protons and the muons now populating the beam. With the timing separation, the protons can be dumped and the muons can be extracted to continue along M4 and M5 to the g–2 storage ring. Figure reproduced from reference [14].

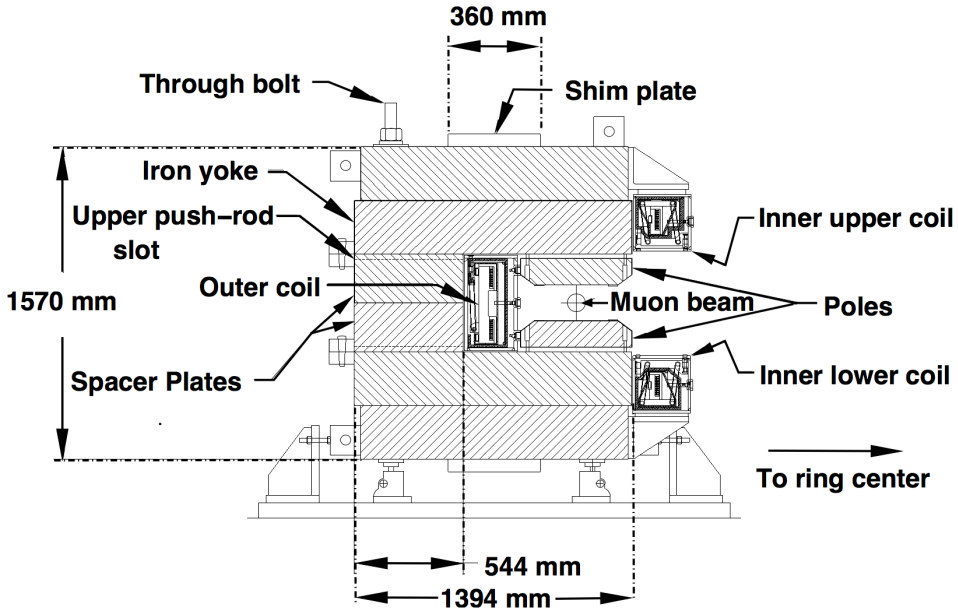

Figure 2.2: The cross-sectional view of the storage ring. The magnetic field is produced by currents in the three superconducting coils. The field strength is primarily determined by the geometry for the flux capture in the "C" yoke and pole surfaces. The field can be adjusted locally in azimuthal using a built-in shim kit that includes top hat shims, edge shims, and wedge shims. Figure reproduced from reference [3].

fills deliver $\mathcal{O}(10^4)$ muons to the storage ring at an average rate of $12\,\mathrm{s}^{-1}$ [14].

2.3 The Magnetic Storage Ring

The magnetic storage ring is the hardware for creating the main magnetic dipole field for the experiment. The major components include three superconducting NbTi/Cu coils, a dozen "C"-shaped flux return yokes in the form of steel blocks and plates, 72 high-purity steel poles, and a built-in adjustment kit with more than 1,000 tunable knobs. Figure 2.3 depicts a cross-sectional view of the magnet. The full assembly is around 3 m tall, 15 m wide, and weighs more than 700 tons [14]. The storage ring was designed with the objective of creating a magnetic field of 1.4513 T that is as uniform as possible over a toroidal storage volume with a major radius of 7.112 m and a minor radius of 4.5 cm.

The magnetic field itself is generated by driving a current of a current of 5176 A through

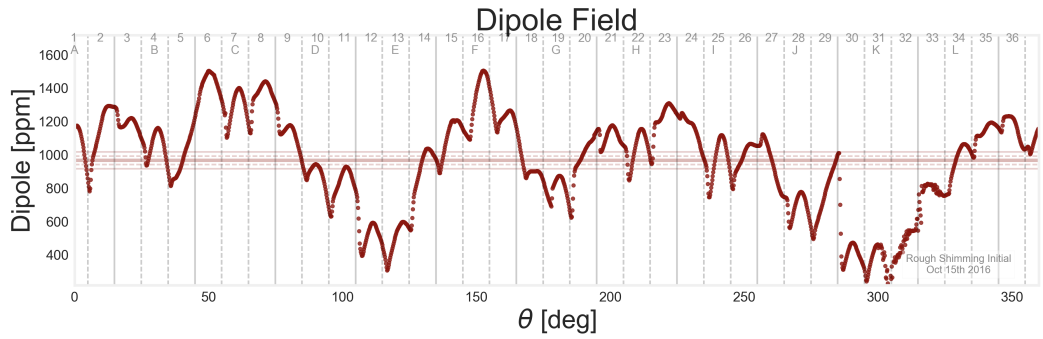

Figure 2.3: The initial dipole (average at azimuthal point) magnetic field when the storage field was first fully measured at Fermilab. The peak-to-peak variation was around 1400 ppm. The umbrella-like sub-structure occurs every 10 degrees and corresponds to the changing gap between the slightly curved pole pieces. The horizontal bands represent the average field value and the uniformity goals set forth by the experiment. The band formed by the two dotted lines around the average value indicates the experimental goal of ±25 ppm.

all three superconducting coils. Two of the coils reside at $R_{inner} = 6677\,\text{mm}$, inside the storage radius. And, one of the coils resides at a larger radius, $R_{outer} = 7512\,\text{mm}$. The outer coil contains twice as much superconducting material as the two inner coils to generate approximately equal magnetic fields from inner and outer coils. The inner and outer coils run with opposing currents to create a nearly vertical B-field in the space between them. The overall strength of the magnetic field depends on the amount of ferric material around the coils to contain magnetic flux and concentrate the magnetic field. By design, the return yokes and pole pieces draw in magnetic flux at the operational current to produce a magnetic field close to 1.4513 T in the opening of the "C". To first approximation the magnetic field uniformity depends on the geometric uniformity of the gap between the two pole surfaces. The precision machined, high-purity pole pieces provide two extremely flat parallel surfaces for optimal field uniformity. The final dipole magnetic field uniformity at BNL was around ±50 ppm [3]. An effort was made to recreate the final state of the magnet assembly after the storage ring was transported to FNAL and re-assembled. When initially powered though, the magnetic field had variations of around 1400 ppm as shown in figure 2.3.

The last essential piece of the magnetic storage ring is the intrinsic shim kit. Shims in this context are objects made of ferric material with full azimuthal coverage of the storage

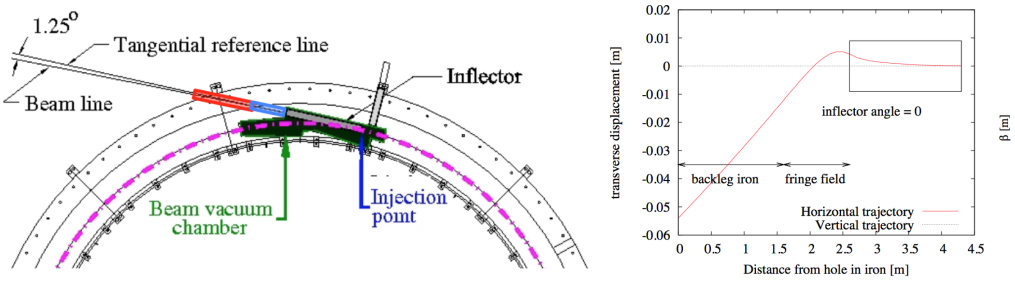

Figure 2.4: The inflector diagram on the left illustrates the relative position of the inflector as the bridge from essentially outside of the storage field through to within the storage volume. The trajectory of the beamline through the inflector area is depicted in the right plot. Figure reproduced from reference [14].

ring which can be adjusted to improve the magnetic field uniformity locally. A collection of edge shims, wedge shims, and top hats provides over 1,000 knobs to make azimuthally localized changes to the magnetic field. Magnetic field shimming is discussed in further detail in chapter 5.

2.4 Muon Injection

To store muons in the strong vertical magnetic field, they must be transported through the back leg, fringe field region of the magnet. The non-linear natural path through this region is non-trivial to predict or manufacture in the storage ring. The solution to simplify muon injection was to create a homogeneous magnetic field volume to provide a predictable path into the storage region. This field region is created with a cleverly designed magnet, called the inflector, which is able to achieve a strong, nearly uniform vertical field over a small volume. It is designed to contain much of its flux from leaking and perturbing the nearby precision magnetic field region. The flux capture in the initial design was insufficient and it was later improved by adding a superconducting shield to the outside. The inflector was designed to be adjustable over small angles which allowed the apparatus to optimize the nearly straight injection of muons. See figure 2.4 for an image of the inflector and the beam trajectory through the inflector. [14, 3]

Another problem arises after the injected muons complete their first orbit around the

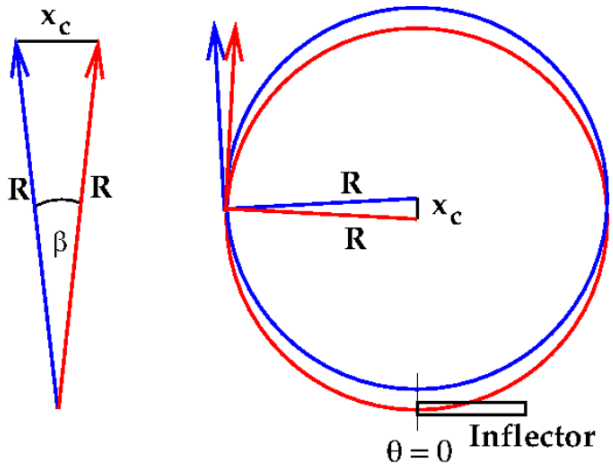

Figure 2.5: The figure illustrates the orbital mismatch problem with muon injection. The red track illustrates the problem where the muons collide with the injection point, and the blue track illustrates the good orbit achieved after a perfect kick. Figure reproduced from reference [14].

storage ring. Moving on circular orbits in a nearly uniform vertical magnetic field, the muons would hit the downstream end of the inflector as they return to the injection point. In figure 2.5, the red track illustrates the problem and the blue track illustrates a optimally shifted orbit achieved after a fast kicker magnet imparts the necessary angular shift onto the muon bunch. The kicker would ideally produce a flat magnetic field around 275 G with very sharp rise and fall times, $\mathcal{O}(10\ ns)$. The real pulse shape deviates from a perfect square pulse and re-kicks muons with the residual magnetic field, but a newly designed kicker will help to prevent double kicking muons as they return to the kicker region on their second orbit of the ring [14].

2.5 Magnetic Field

The magnetic field is of critical importance to the muon g–2 experiment. A magnetic field of 1.4513 T puts muons with "magic" momentum, 3.094 GeV/c, into uniform circular motion at the "magic" radius of 7.112 m. The prescribed parameters lock a fraction, $\mathcal{O}(0.03)$, of the injected muons in circular motion until they decay. To first order, the value of the

magnetic field directly affects the rate of muon spin precession and cyclotron frequency. In this light, the average magnetic field must be very well measured, since it folds directly into the determination of ω_a. To second order the magnetic field couples to the symmetries in the muon beam influencing the beam dynamics of the stored muons. These deviations in beam dynamics make the analysis which matches muon trajectories with the magnetic fields along them more difficult, and therefore add uncertainty to the determination of the expectation value for ω_p.

2.5.1 The Field Expansion

The ideal magnetic field for the experiment is entirely vertical at 1.4513 T with no deviations. However, Maxwell's equations do not permit such a perfect field within a finite amount of space. Reality requires that the experiment consider the quality of the central field and the effects of perturbations to the field which is conveniently done using a decomposition. The expansion in equation 2.8 introduces a common set of functions used in discussing magnetic field perturbations. The expansion assumes the domain where $B_r \ll B_z$ and $B_\phi \ll B_z$ effectively reducing the 3D problem to a 2D problem. The idea is to then expand inside the two-dimensional plane defined by a 4.5 cm circle in r and z at a single value of ϕ.

$$\begin{aligned} B_z &= \sum_{n=0}^{\infty} \rho^n [a_n \sin n\phi + b_n \cos n\phi] \\ B_r &= \sum_{n=1}^{\infty} \rho^n [c_n \sin n\phi + d_n \cos n\phi] \end{aligned} \tag{2.8}$$

A visualization of the terms in the multipole expansion is given in figure 2.6. Each multipole highlights a possible symmetry of the field over the 2D azimuthal slice of the storage region. The magnetic field is characterized in terms of multipoles where the dipole should average to 1.4513 T and all other terms should be minimized.

2.5.2 Determining $\langle\omega_p\rangle$

The magnetic field is measured and monitored with a suite of magnetometers, custom pNMR probes. Chapter 3 discusses the analysis in depth, but a cursory overview is still provided

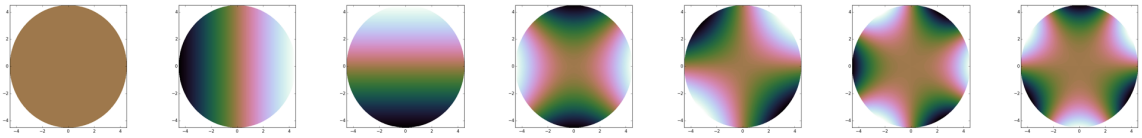

Figure 2.6: The first seven multipoles in the field expansion, equation 2.8. The first on the left is the dipole term which simply averages each point equally over the domain. The next two are the normal and skew quadrupole which represent an inner to outer or top to bottom asymmetry in the field respectively. The next four terms similarly represent further symmetries of the field. They are termed: normal sextupole, skew sextupole, normal octupole, and skew octupole.

here for completeness. The field analysis combines pNMR measurements from three different major subsystems. The trolley system is a cylindrical aluminum shell equipped to travel through the muon storage volume and an array of pNMR probes to measure the magnetic field in 2D azimuthal slices. A full set of measurements from the trolley system is combined to produce a magnetic field map as a function of position over the entire muon storage volume. While the trolley resides in the storage volume, the ring cannot accept muons, so a set of 378 stationary pNMR devices outside the storage volume called the fixed probe system monitor the magnetic field drift between trolley runs. The final subsystem is the absolute calibration probe system which includes several, essentially independent devices from the rest of the pNMR probes. The absolute calibration probes are used to correct for systematic effects that shift the proton precession frequency in the trolley and fixed probes.

All the field measurements combine to produce the magnetic storage field as a function of time and position, the target quantity from the field measurements. Expression 2.1 and 2.6 actually need $\langle \omega_p \rangle$ which represents the average field experienced by muons. The muon distribution, $M(\vec{r}, t)$, and the field values need to be integrated over the all muons that contribute to determination of ω_a. An appropriate expression for $\langle \omega_p \rangle$ is then given in equation 2.9.

$$\langle \omega_p \rangle = \int M(\vec{r}, t)\, \omega_p(\vec{r}, t)\; dt\; dV \tag{2.9}$$

2.6 Spin Precession

It is essential that the storage ring contain muons in stable orbits until they decay into electrons. The muons come in polarized in the injection direction. While propagating around the storage ring, the polarization vector of the muons undergoes spin precession.

2.6.1 Decay Characteristics

Eventually the muons decay into electrons and neutrinos as discussed in section 1.2. Due to maximal parity violation in weak decays, the spin direction is correlated to the energy and momentum direction of the decay electrons. In the rest frame (see figure 1.4) it can be seen that accumulation of events with energy either less than half or greater than half the total possible energy would result in a signal (anti-)correlated with the spin direction. Applying similar logic to the boosted spectrum (see equation 2.10 and figure 2.7) is a key insight into the way g–2 works.

$$n_{lab}(y) = \frac{-8y^2 + y + 1}{4y^2 - 5y - 5} \qquad\qquad a_{lab}(y) = \tfrac{1}{5}(y-1)(4y^2 - 5y - 5) \tag{2.10}$$

The point is still to choose a domain which has an integral maximizing the spin-momentum correlation for the events. The expression for the statistical figure of merit for an event counting analysis is

$$NA^2 = \int_{y_{thresh}}^{1} n(y)a^2(y)\, dy. \tag{2.11}$$

The ideal energy threshold for this figure of merit is $0.4{\times}3.094\,\mathrm{GeV/c} = 1.24\,\mathrm{GeV/c}$ if all electrons could be effectively captured. When detector acceptance is folded in, the optimal energy threshold becomes $0.6{\times}3.094\,\mathrm{GeV/c} = 1.86\,\mathrm{GeV/c}$ [14]. The connection of muon spin to the decay electron energy is now established, so the last step in measuring the spin precession requires a technique to measure the energy and the birth time of emitted electrons. [3]

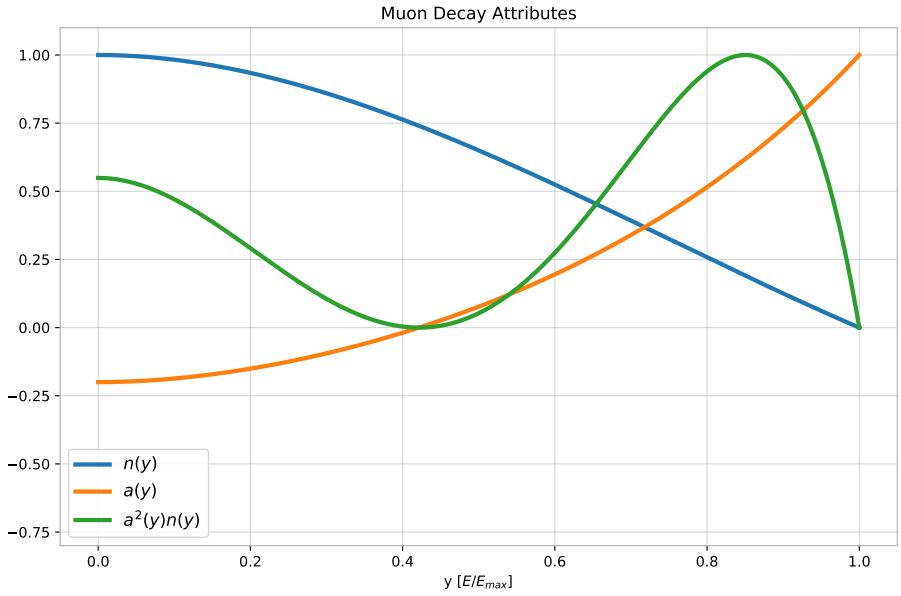

Figure 2.7: The plot depicts an unnormalized probability distribution for the number of decay electrons in the boosted lab frame, and similarly the fractional asymmetry as a function of the electron energy where the energy is represented as a fraction of the maximum possible electron energy from a muon momentum at 3.094 GeV/c. Alongside the distributions, there is also a plot of NA^2 which is the statistical figure of merit for measuring decay events and maximizing the signal from the spin correlation.

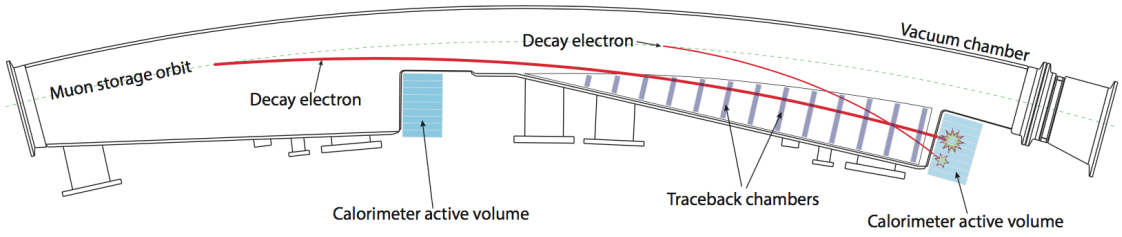

Figure 2.8: The figure shows a few typical decay electron trajectories. The muon decays into an electron with lower momentum which by necessity takes a path with a smaller radius of curvature. The electron curls inward into one of 24 calorimeter blocks around the inner radius of the storage ring. Figure reproduced from reference [14].

2.6.2 Electron Calorimetry

A suite of calorimeters arranged around the inner radius of the storage ring measure properties of the decay electrons. The decay electrons always have less energy than the magic momentum, and so must curl inward on a smaller orbit radius than the storage orbit radius of the initial muon. The trajectory of a typical decay electron is shown in figure 2.8. The trajectory of the electrons intersects with a calorimeter, at least the higher energy electrons anyway. The higher energy electrons have a wider trajectory similar to the parent muon and hit the calorimeters head on. Lower energy electrons have a sharp trajectory which can curl inward between calorimeter stations.

The calorimeters establish a time and energy for each detected electron. The device consists of a 6x9 array of PbF_2 crystals each with a physically smaller array of Geiger-like photon counting hardware called a silicon photomultiplier (SiPM). Segmenting the detector decreases the likelihood of two decay events occurring simultaneously in the same measurement channel. These double events are referred to as pileup and were a major source of uncertainty in E821. The incoming electron produces thousands of Čerenkov photons per GeV. The photons collect at the opposite end of the crystal where they activate SiPM channels. The subsequent pulse undergoes signal shaping, digitization, and a pulse fitting analysis routine to determine the energy and the time of the electron decay event.

2.6.3 Determining ω_a

The signal for the anomalous muon precession frequency manifests in histogramming electron detection events with energy above a cutoff threshold into time bins. The so-called "wiggle" plot exhibits clear oscillations which through careful systematic analysis result in a precise value for the anomalous precession frequency, ω_a. The muons undergo decay into electrons at an exponential rate, and the energy distribution of those electrons changes depending on the direction of the spin vector. As the spin vector precesses, the number of decay electrons measured above threshold also fluctuates as shown in figure 2.9. The basic model for the precession frequency signal is then given as

$$N(t, E) = N_0(E)e^{-t/\tau_\mu}\left[1 + A(E)\cos(\omega_a t + \phi_0(E))\right] \tag{2.12}$$

where N_0 is the total number of muons above threshold at time zero, τ_μ is the effective muon lifetime, $A(E)$ is the total asymmetry of muons above threshold, and $\phi_0(E)$ is the phase of the spin vectors at t_0 for the fit. A five parameter fit using equation 2.12 yields the anomalous spin precession for muons. The distribution and fit is depicted in figure 2.9.

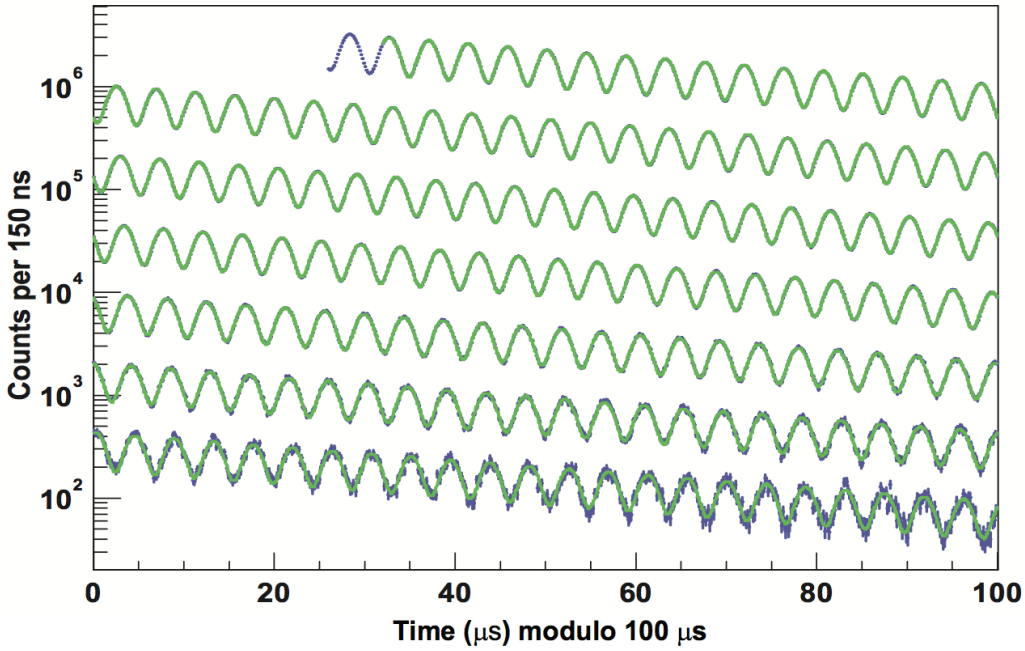

Figure 2.9: The histogram is from the E989 TDR, reference[14]. It shows the number of electron decay events above threshold as a function of time. The oscillations seen on top of the exponential decay correlate to the spin precession of the muons as they propagate around the storage ring. The number of high energy decays is enhanced as the spin aligns with the momentum and decreased when the spin anti-aligns with the momentum vector.

Chapter 3

MAGNETIC FIELD MEASUREMENT

3.1 Magnetic Field Uncertainty Budget

The $(g–2)_\mu$ experiment allots 70 ppb total uncertainty for ω_p which predominantly comes from magnetic field measurements. The uncertainty budget contributions are enumerated in source [14] and summarized in table 3.1. Uncertainty constraints highlight relevant details for the next sections which walk through the logistics of magnetic field measurements.

Table 3.1: Magnetic Field Uncertainty Budget

Source	Uncertainty [ppb]
Calibration of fixed probes	35
Calibration of trolley probes	30
Trolley measurements of B_0	30
Interpolation with fixed probes	30
Muon distribution	10
Inflector fringe field	-
Transient external B-fields	5
Other sources	30
Total $\delta\omega_p$	70

3.2 Measurement Technique

The uncertainty allocation for magnetic field determination is entirely systematic. Therefore, a typical measurement needs to be precise enough that the statistical error on the overall magnetic field determination has negligible influence. This motivates a stated uncertainty goal of 10 ppb for each standard field measurement [14]. Commercial magnetome-

ters with the required precision and geometric interface were not available within the E989 financial budget, so researchers at the University of Washington developed a set of custom magnetometers. The magnetometers use pulsed nuclear magnetic resonance (pNMR), a reprise of the measurement technique used in E821. Using pNMR with a sharp resonance at the expected Larmor frequency of the g–2 magnetic field, the magnetometers are able to reach the aforementioned precision goals.

3.2.1 Pulsed Nuclear Magnetic Resonance

A pNMR measurement exploits Larmor precession of protons to produce a signal with magnetic field information embedded in a frequency measurement.

Spin Polarization

Each pNMR device has a volume with a population of quasi-free protons which are manipulated to produce the measurement signal. The protons in the active volume have a small net spin polarization in the direction of the external field due to the spin-field interaction, $U = -\vec{\mu}\cdot\vec{B}$. The statistical state of the proton spins due to the spin-field interaction follows from the Boltzmann distribution [23]. At room temperature (300 K) in the g–2 magnetic field (1.4513 T), the average difference between the fraction of spin up and spin down protons using CODATA values for μ_p and k_B is given by

$$\frac{e^{-\mu_z B_z/k_B T} - e^{\mu_z B_z/k_B T}}{e^{-\mu_z B_z/k_B T} + e^{\mu_z B_z/k_B T}} = 4.94 \times 10^{-6}, \tag{3.1}$$

a surprisingly small, but still very useful net spin polarization.

Radio Frequency Pulse

To initiate a pNMR signal, a secondary radio frequency (RF) magnetic field is applied in the orthogonal plane by driving a current through a solenoid surrounding the proton sample. The RF magnetic field alternates at a frequency approximately equal to the expected spin precession frequency in the given external magnetic field. Define the magnetic fields to be

$$\vec{B}(t) = b_0 \cos(\Omega t)\hat{x} + b_0 \sin(\Omega t)\hat{y} + B_0\hat{z} \tag{3.2}$$

where B_0 is the original magnetic field strength, b_0 is the RF field strength, and Ω is the angular frequency of the RF field.

The process unfolds intuitively in the reference frame rotating at $\Omega = \omega_p$. In the rotating frame, the effective spin precession frequency from B'_z is $\omega_p - \Omega = 0$, and therefore the effective magnetic field $B'_z = 0$. On the x-axis, $B_x = b_0$ in the rotating frame, so the RF field rotates the spin polarization vector around the x-axis into the x-y plane. The RF field is turned off at $T = \pi/(2\Omega)$, so that the spin polarization vectors only rotates $\pi/2$ rad and ends on the y-axis, hence, it is often called a $\pi/2$ pulse. The essential behavior is similar when the RF frequency and the Larmor frequency do not match exactly. In the lab frame the process can be understood in terms of a resonance effect due to the RF frequency matching the spin transition frequency.

Free Induction Decay (FID)

After the $\pi/2$ pulse, the polarized protons precess in the x-y plane at a rate proportional to the original vertical field. The precession signal and the subsequent decay is known as a free induction decay. The same solenoid used to produce the $\pi/2$ pulse now acts as a pickup coil, and the net proton spin vector in the orthogonal plane produces a changing magnetic flux and thereby an induction current as it precesses.

The signal attenuates due to two dominant effects. The first effect is called longitudinal relaxation. It occurs due to spin alignment with B_z, the same spin-field interaction which imparted a net polarization in the first place. The net spin vector returns to equilibrium at an exponential rate with a time constant usually called T_1. The second effect is called transverse relaxation. It is essentially decoherence among the collection of proton spins. Not all spins experience exactly the same magnetic fields, so after the RF pulse they precess at different rates and the net spin vector can decohere. The rate of decoherence is also exponential, and it is given the time constant T_2. The precession frequency determined by analyzing the pNMR signal can serve as high precision magnetic field proxy.

Bloch Equations

The aggregate spins in the proton sample evolves according to the Bloch Equations (equations 3.3), an ordinary differential equation that describes the time evolution of the magnetization vector in external magnetic fields.

$$\begin{aligned}
\frac{dM_x(t)}{dt} &= \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_x - \frac{M_x(t)}{T_2} \\
\frac{dM_y(t)}{dt} &= \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_y - \frac{M_y(t)}{T_2} \\
\frac{dM_z(t)}{dt} &= \gamma(\mathbf{M}(t) \times \mathbf{B}(t))_z - \frac{M_z(t) - M_0}{T_1}
\end{aligned} \tag{3.3}$$

Consider the case of a perfectly uniform magnetic field, $T_2 \to \infty$, and a perfect $\pi/2$ pulse. This case is referred to in this document as an "ideal FID" waveform, and it can be solved exactly from the Bloch equations. The resulting equation, 3.4, exhibits the primary characteristics of an FID, a sinusoid term from proton precession and an exponential decay envelope from the longitudinal relaxation. Figure 3.2.1 exhibits the waveform of an ideal FID.

$$f(t) = e^{-t/T_1} \sin(\omega t - \phi_0) \tag{3.4}$$

In a realistic situation, the FID signal is a summation of signals produced by protons precessing at different frequencies due to gradients in the magnetic field, and the signal can bear strong deviations from the ideal FID form. Nevertheless, the ideal signal is a valid testing ground for frequency extraction algorithms. Examples of measured FIDs are shown in figure 3.2.1. The measured FIDs exhibit distortions from defects in electronics like the asymmetry of the first waveform and the spikes in the second waveform. They also exhibit the effects of magnetic field gradients on the signal. For instance, the second FID plot shows node-like behavior on the waveform which is indicative of an FID of a pNMR probe in magnetic field gradients.

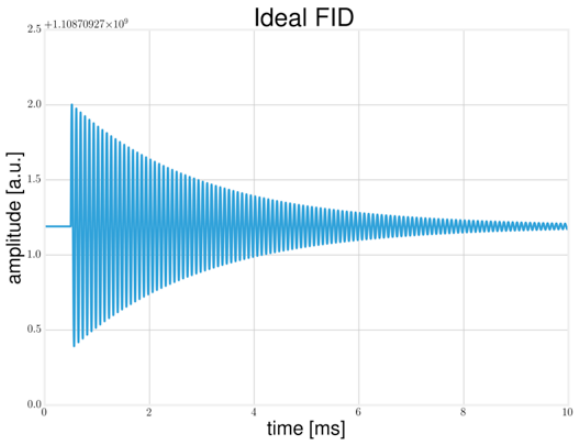

Figure 3.1: The figure shows an idealized FID waveform. Functionally, the waveform is a sinusoid multiplied with an exponential decay between at some initial time. It represents the signal from a perfect $\pi/2$ to a perfect pNMR probe.

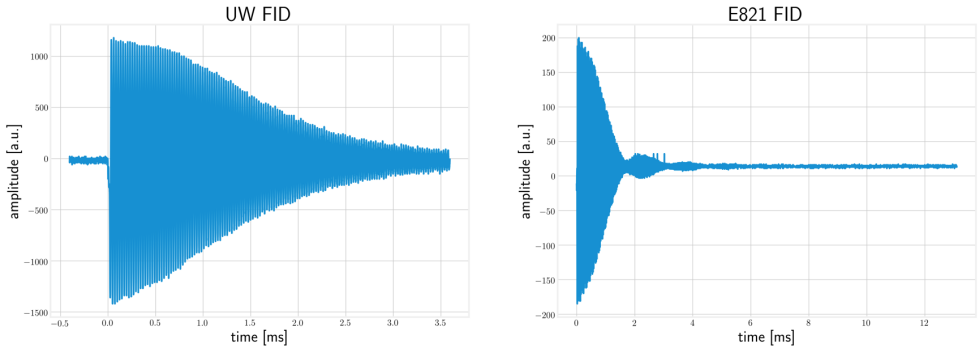

Figure 3.2: Two examples of measured FIDs are shown above. The image on left was measured with a test setup at CENPA (University of Washington), and the waveform on the right is an example from the E821 experiment fixed probe system, probe 1 from multiplexor 1 taken on August 1st 1998 [20]. The measured FIDs show distortions from imperfections in electronics like the asymmetry of the first waveform and the spikes in the second waveform. They also exhibit the effects of magnetic field gradients on the signal. For instance, the second FID plot shows node-like behavior on the waveform which is indicative of a gradient of FID frequencies beating against each other.

Systematic Effects

There are several systematic effects that probe designs can mitigate. The typical set of corrections to pNMR frequency is given in the $(g–2)_\mu$ TDR [14] as

$$\omega_{probe} = (1 - \sigma(T) - \delta_b(T) - \delta_p - \delta_s)\omega_p. \tag{3.5}$$

The first correction is due to diamagnetic shielding effects in the protonated volume. The correction is represented by the σ function. The electron clouds of nearby molecules act as a local dielectric and shield the proton from the full magnetic field. This effect is typically on the order of tens of ppm. It is typically modeled with temperature dependence and corrected for in the resulting frequency [31].

The second term, δ_b, denotes the correction for the bulk susceptibility of the sample. It is also sometimes represented in the form $\left(\epsilon - \frac{4\pi}{3}\right)\chi(T)$. This includes geometric effects which shift the resonance frequency due to magnetization of the sample. They are minimized and understood by using ideal shapes with known factors for instance, $\delta_{sphere} = \frac{4\pi}{3}$ and $\delta_{cylinder} = 2\pi$ [14]. Additionally, cylinders and spheres can be fabricated with high accuracy making them common design choices.

The third term, δ_p, corrects for paramagnetic impurities in the sample. impurities are often added to shorten the signal to the desired time scale. For instance, H_2O has a natural signal relaxation time on the order of 10 s, but H_2O doped with $CuSO_4$ can reduce the relaxation time to around 40 ms [3].

The last term, δ_s, includes the effects of magnetic perturbations incorporated into the material design of the system. The obvious solution is to use materials that are "non-magnetic" even though all materials exhibit some magnetic susceptibility. Another clever solution is the combine materials in a way that cancels paramagnetic and diamagnetic effects such as copper and aluminum [31].

References [33] and [31] are quality resources for more details on NMR fundamentals and systematics respectively.

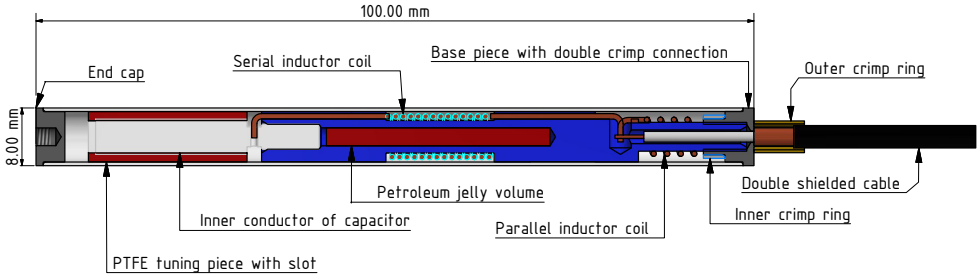

Figure 3.3: The E989 pNMR probe design. In the center of the diagram, a cylindrical volume holds the petroleum jelly which contains protons to polarize. The proton volume is surrounded by the serial inductor coil used to inject a the $\pi/2$ pulse and the pick up the FID signal. The other inductor matches the impedance of the system, and the end of the probe contains a Teflon screw for fine tuning of the resonance. Figure reproduced from reference [12].

3.2.2 Probe Designs

The custom pNMR probes come in three different builds. The standard probes which are used in the fixed probe system and the trolley system for in situ measurements of the magnetic field in the muon storage region. The absolute calibration probes are used in measurements external to the muons and the storage ring, and the plunging probes are the third type of probes which are used for cross-calibration of the absolute calibration probe and the standard probes.

Standard pNMR Probes

The core design of the pNMR probes is the same as the design from the previous experiment. Figure 3.2.2 depicts the design and components of the probe. Each probe consists of two induction coils, a Teflon backbone, a tunable capacitor, an aluminum shell, and a 5 m RF cable with SMA connector. The main induction coil is used to inject the $\pi/2$ pulse, and again to pick up the pNMR precession signal. A second induction coil is used to match the 50 Ω impedance of the line. The capacitor allows the probe resonance frequency to be tuned finely. The aluminum shell provides the outer conductor of a cylindrical capacitor and shields against external effects.

The field team at the University of Washington iterated on the probe design making several improvements. One of the improvements was in the robustness of the connection to the signal cable. The new design used a crimp connection to secure the cable to provide a more secure connection. Another improvement was the design of the tuning capacitor. With the new design, probe tuning can be done without removing the outer shell. Removing the shell has a chance of damaging structures inside the probe, and minimizing removals makes tuning easier, faster and safer overall. Another improvement is a substitution of the material used as a proton sample. The E821 design used H_2O as a proton source with $CuSO_4$ to decrease the T_1 relaxation time constant. The new design replaces the mixture with petroleum jelly as a proton source. The water design was seen to cause corrosion in many of the E821 probes, so the new design should alleviate those concerns and provide a longer, stable lifetime for the device.

The standard pNMR probes are used in two different but complementary systems to measure the magnetic field, the fixed probe system and the trolley system, which are detailed in section 3.3.

Absolute Calibration Probes

The standard pNMR probes can only give a proxy for the magnetic field, not a well calibrated absolute value. For the absolute magnetic field value, the absolute calibration pNMR probes have to come into the fold. The absolute calibration probes work on the same pNMR principles, but the design bears more consideration in minimizing and understanding known systematic effects which shift the precession frequency from that of free protons. For instance, they may use a precisely fabricated spherical proton volume in order to minimize geometric effects. And, all probe materials and probe holder materials are stringently tested for magnetic perturbations.

Plunging Probes

The last species of pNMR probe is the so-called plunging probe. The probe is specifically designed to work with a system that transfers the calibration from the absolute probes to

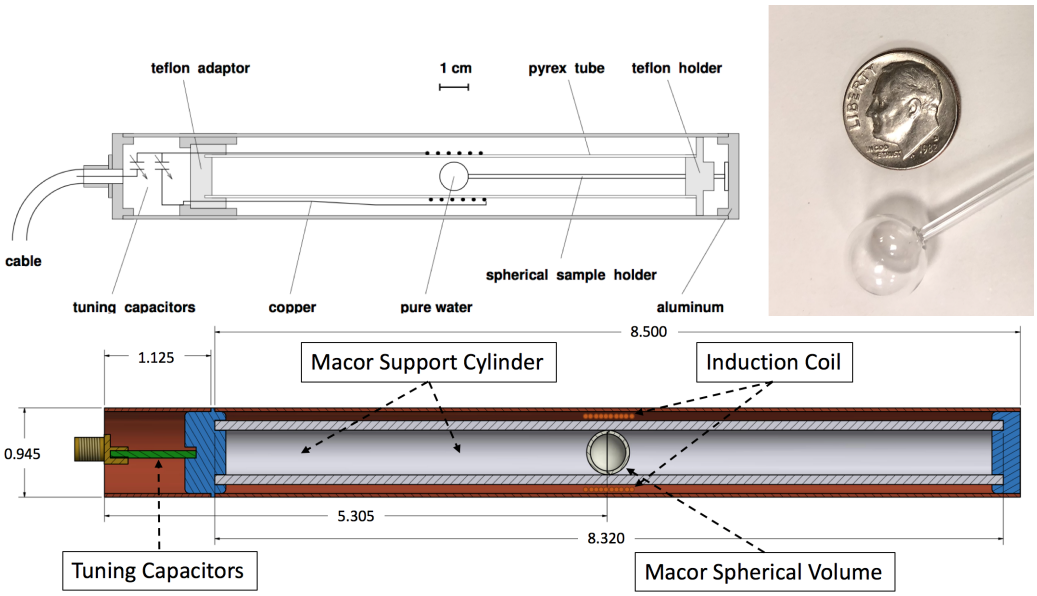

Figure 3.4: The top image is a design diagram for the E821 absolute calibration probe, and the bottom design diagram is for the new E989 calibration probe. The E821 probe is still in use alongside the new Macor probe designed by UMass and an entirely different type of probe using ^{3}He instead of water as a proton source. The top left figure is reproduced from reference [3], the top right figure was reproduced from reference [21], and the bottom figure was reproduced from [13].

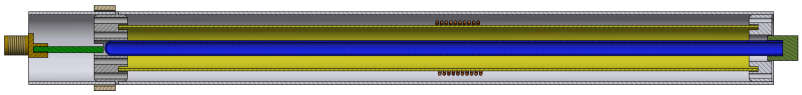

Figure 3.5: A diagram of the plunging probe deployed in E989. The probe is designed to work with motion system that allows the probe to be align with both the absolute and standard probes to determine the proper shifts needed to bring pNMR measurement from a frequency to an absolute magnetic field value. Figure reproduced from [13].

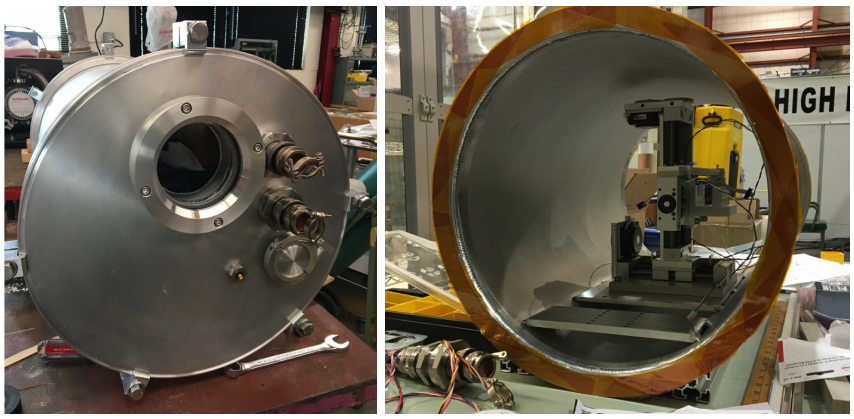

Figure 3.6: Images of some plunging probe system hardware. The apparatus pictured on the left holds the translation stage used to align the plunging probes with other pNMR probes, and the image on the right shows the translation stage itself. Figure reproduced from reference [13].

the standard probes (shown in figure 3.2.2). The plunging probe system deploys a motion controller stage inside a vacuum tight volume (see figure 3.2.2). The translation stage aligns the plunging probe with the standard pNMR probes in the trolley system and absolute calibration probes in situ to make frequency comparison measurements.

3.3 Magnetic Field Subsystems

3.3.1 Fixed Probe System (FPS)

The fixed probe system is a suite of stationary probes located above and below the muon storage volume. There are 378 probe in total, and they measure the field continuously while the muon beam is delivered to the storage ring. In this way, the magnetic field drift is

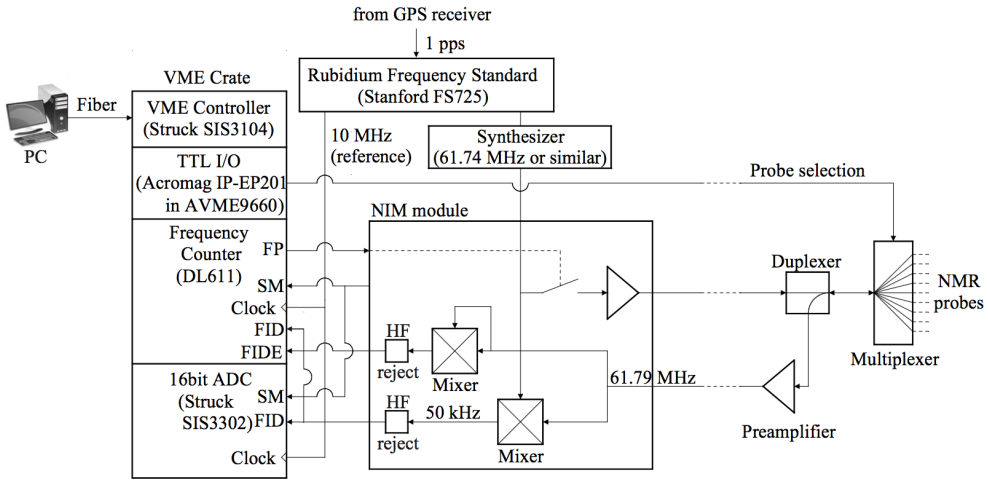

Figure 3.7: The pNMR probes signal requires a parallel system of controlled electronics equipment. The E989 reprises the NMR pulser design from the E821 experiments. The block diagram illustrates the measurement process and components of the system. In order to generate a pNMR signal, the first phase is generating a $\pi/2$ pulse to rotate the protonated volume. The $\pi/2$ pulse is generated by a TTL input trigger which is tunable from 4 µs to 7 µs. After the $\pi/2$ pulse, the free induction decay signal is mixed down against a very stable, 61.74 MHz rubidium frequency standard. The resulting signal feeds through a low-pass filter, and the filtered signal is fed into a waveform digitizer and recorded at a sampling rate of 10 MHz and sampling depth of 16 bits. The digitized waveform contains frequency information that represents the magnetic field. Figure reproduced from reference [14]

monitored at all times. The measurement logistics are detailed in figure 3.3.1.

The fixed probes are arranged in a very specific manner. They are secured inside grooves cut into the walls of the 12 vacuum chambers that reside inside the storage magnet gap. In E821, many of the pNMR probe grooves placed the active volume of the probe near pole and yoke boundaries where the magnetic field gradients were large. Due to the large gradients, the FID signal from the probe was unusable, and the E989 researchers decided to fix the problem. The grooves were extended 5 cm to shift the pNMR probes away from the magnetic field gradients and improve the FID signal.

Each probe resides at one of three radii: 7082 mm, 7112 mm, or 7142 The probe arrangement in azimuth repeats over each vacuum chamber effectively following the yoke sections, so yoke B with vacuum chamber 2 is the same azimuthal pattern as yoke C. The

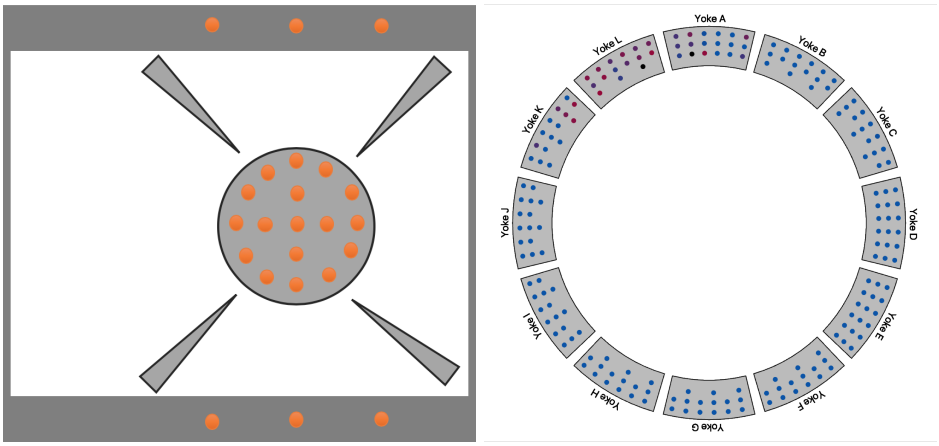

Figure 3.8: The left diagram shows the cross-sectional view of the storage ring with probe locations marked for the trolley and fixed probes. The three probes above and belows the storage volume illustrate the arrangement of fixed probes. The right diagram depicts the azimuthal arrangement of fixed probe in the storage ring. The probe arrangement shown is present both above and below the storage volume. Note that the FPS arrangement is not drawn to a realistic geometry.

specific angles over the chambers are: 3.72°, 6.64°, 12.53°, 17.53°, 22.53°, 26.89°. The one exception is yoke A/vacuum chamber 12 where the grooves for the FPS were not extended. In that case the probes reside at azimuthal angles of: 2.65°, 6.64°, 12.53°, 17.53°, 22.53°, 27.68°. The last detail in the probe arrangement is the number of probes. Most vacuum chambers hold 30 probes numbering 3-2-3-2-3-2 at each azimuthal position (top and bottom), but three chambers have 36 probes opting for a 3-3-3-3-3-3 scheme at each position. Figure 3.3.1 illustrates the arrangement of fixed probes. Note that the azimuthal spacing and radial spacing is drawn more symmetric in the figure and not to scale.

3.3.2 Trolley Probe System (TPS)

The trolley measures the magnetic field in the muon storage volume. It is a cylindrical shell with wheels that interface with rails protruding from the vacuum chambers which reside in the gap between the pole pieces as seen in figure 3.3.1. Internally the trolley contains an array of 17 of the standard pNMR probes as shown in figure 3.3.1 and the full set of electronics needed to produce and digitize FID signals. It cannot measure while muons are

Figure 3.9: An image of the trolley module itself. The trolley rides along a rail system inside the vacuum chambers. It travels through the muon storage volume.

injected though, so there is a naturally trade-off between measuring the magnetic storage field with the trolley and interpolating the field with the fixed probe system while the muons are delivered. In the trolley system the pNMR probes are arranged with one central device, a ring of 4 devices at a radius of 17.5 mm, and 12 more in an outer ring at 3.5 mm.

3.3.3 Fluxgate System (FS)

The Fluxgate System consists of four vector magnetometers. The four devices are Bartington Mag690-1000 three-axis magnetic field sensors. They are capable of measuring at readout rates of 8 kHz but limited in range to -10 G to 10 G. The limited range prevents the devices from making measurements too close to the storage field, but that is not their intended use anyway. The fluxgate sensors serve as a monitor to check for fast transient and periodic magnetic fields in the experimental hall. The FS picks up magnetic field signals that are too fast for the FPS, and enables the experiment to place a limit on the effect of such transient signals.

3.3.4 Absolute Probe System (APS)

Several unique probes are used as independent determinations of the free proton resonance frequency. E989 will make absolute field calibration measurements with the E821 water probe, a newly designed water probe, and a ^{3}He probe. The probes are engineered to best control for systematic effects that shift the free proton frequency.

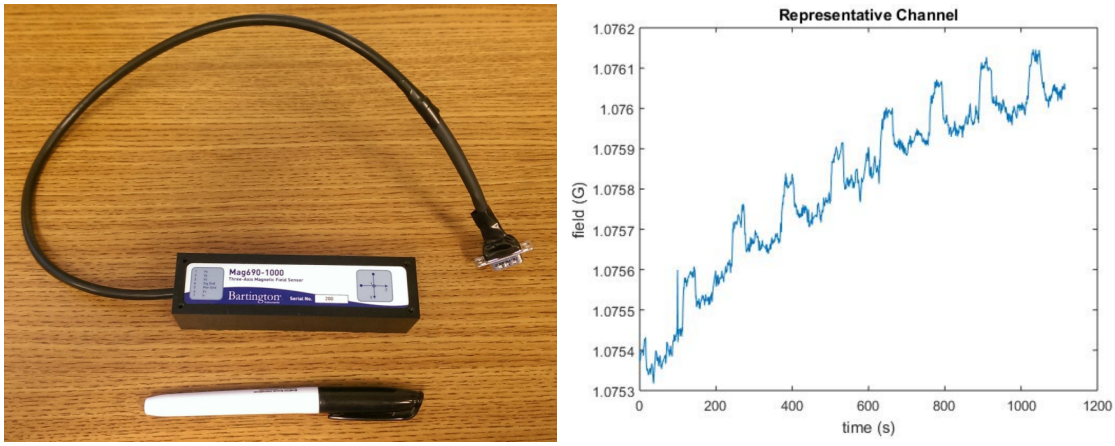

Figure 3.10: The left image is one of the fluxgate sensors. The right image is an average of measurements of the vertical magnetic field over about an hour. Similar structure was seen in both the fluxgate system and fixed probe system. Figures received from private communication with Alec Tewsley-Booth.

The original water probe was designed with intent to maximize precision. It used a highly spherical glass sample volume in order to minimize the bulk susceptibility correction. Sample impurities were also carefully monitored, and the probe materials were measured for magnetic perturbations. The measurement achieved an overall precision of 50 ppb [11].

The redesigned E989 water probe attempts to improve upon some of the systematics of the E821 water probe. The E989 probe uses more stringent testing of the magnetic susceptibility of probe materials. The new design also aims to improve machining tolerance for the spherical and cylindrical components of the probe. Stricter monitoring of impurities is also planned by characterizing the T_1 and T_2 relaxation times of high purity water. With the improvements, the measurements aim to achieve an uncertainty of 35 ppb. [14]

The principal motivation for the ^{3}He probes, a second type of absolute calibration probe, is a cross-check with very different systematics. For instance, the diamagnetic shielding factor, σ, on ^{3}He is larger compared to water probes, but it is known to higher precision.

$$\sigma_{H_2O} = 25.680(2.5) \quad \times 10^{-6} \tag{3.6}$$

$$\sigma_{^3He} = 59.967\ 43(10) \times 10^{-6} \tag{3.7}$$

The diamagnetic shielding also exhibits smaller temperature dependence. The bulk susceptibility correction, δ_b, is smaller for a ^{3}He probe due to the proton volume being gaseous. This also indicates that the measurement is less sensitive to the shape of the proton volume. The remainder of the design is similar to the water probe, but the materials are different providing another difference in systematics. Overall the ^{3}He absolute calibration probe has potential to reach higher precision than the water probe, and the goal for the measurement is again an uncertainty of 35 ppb. [14]

3.3.5 Plunging Probe System (PPS)

The relative relationship between the absolute calibration probe proton frequency and the probe frequency in the FPS and TPS is done with the plunging probe system. The plunging probes are brought into the same volume of magnetic field measured and corrected by the absolute probe to establish a calibration from the absolute probe, which does not enter the storage ring, to a system which has a ring entry point. The plunging probes are inserted into an azimuthal slice of the storage volume. The plunging probe is moved by a motor stage to be as near to each trolley probe as possible. In this way, the free proton resonance calibration from the absolute probe can propagate to the standard pNMR probes used to measure the g–2 magnetic field in situ.

3.4 Magnetic Field Reconstruction

The first important aspect of magnetic field reconstruction is understanding that the name itself, magnetic field, is not exactly precise. The hardware used to make magnetic field measurements mostly consists of pNMR magnetometers. Measurements using pNMR are proton precession frequencies, not magnetic field measurements. They act as a proxy for the magnetic field in a specific set of circumstances which are assumed in the $(g–2)_\mu$ magnetic storage field. The first assumption is that non-vertical components of the magnetic field are negligible, i.e., $B_r \ll B_z$ and $B_\phi \ll B_z$. That assumption is validated by radial and azimuthal field measurements, and the nullification of the radial field in the active shimming stage (see chapter 5). The second assumption is that the pNMR probes are aligned with the magnetic field in such a way that the spin precession frequency that the proton sample

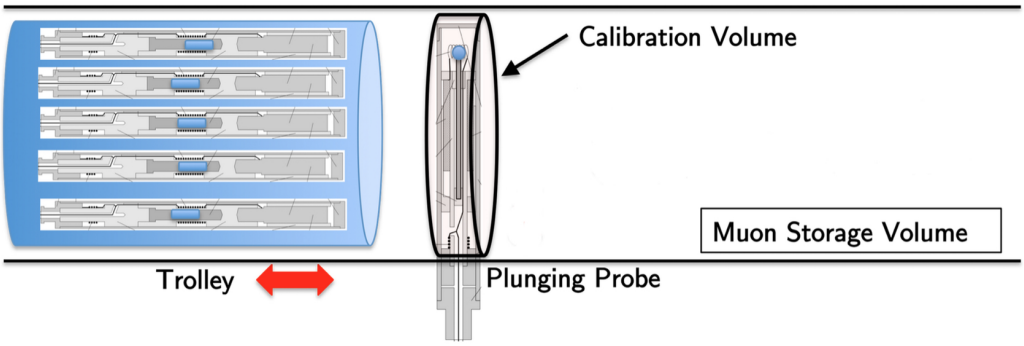

Figure 3.11: The diagram depicts how the plunging probe correlates pNMR frequencies with the trolley. The plunging probe is inserted into the calibration volume, and translated to the location of each trolley probe. The two probes measure the same field and compare frequencies to find the relative calibration correction. Figure reproduced from reference [21].

experiences is representative of the vertical magnetic field, B_z. This is validated by the same measurements of the radial magnetic field, and by alignment specifications on the probes within their respective device. The quantity which is truly reconstructed by the magnetic field reconstruction algorithms is the free proton spin precession distribution which acts as a proxy for the magnetic storage field or more precisely the magnetic flux density distribution.

With that caveat, the process of defining the g–2 magnetic storage field begins with the trolley system. The trolley system takes a set of measurements over the whole azimuth of the magnetic storage ring, and that data is used to define the magnetic flux density for any position in the muon storage volume at the specific time of the run, $\omega(\vec{r}, t_i)$. The magnetic field is expressed as a frequency to reflect the pNMR nature of the measurements. Between trolley runs, measurements from the fixed probes and the fluxgates are used to interpolate change in the flux density distribution. The changes extracted from the fixed probe data are a function of time only, $\delta\omega_{fps}(t)$. And, similarly the contributions to the interpolation from the fluxgates can be a function of time only, $\delta\omega_{fs}(t)$. Putting all the individual contributions together results in equation 3.8.

The algorithms for magnetic field construction are not manifest, and they must be designed and refined by the field team over the duration of the experiment. In light of the fact that the experiment is currently in the commissioning stage, the contributions cannot

be discussed in more than general terms. The magnetic flux density distribution is re-measured periodically, and only the most recent one contributes to the field definition. The most general form for the magnetic flux density distribution is then

$$\omega(\vec{r},t) = \sum_n \Theta(t-t_n)\left[\omega_n(\vec{r}) + \delta\omega(t)\right] \tag{3.8}$$

where $\omega_n(\vec{r})$ is the base magnetic field term over the 3D storage volume at time t_n, $\delta\omega(\vec{r},t)$ is an interpolation term to improve precision between runs, and Θ is the Heaviside function. The last step is calibrating the magnetic field determined by the trolley using the plunging probes and the absolute probes. Calibration shifts the flux density measured into terms of the free proton spin precession frequency.

$$\omega_0(\vec{r},t) = (1-\delta)\omega(\vec{r},t) \tag{3.9}$$

The measurement data and algorithms together define the calibrated magnetic flux density distribution over the muon storage volume.

3.4.1 Map of the Muon Storage Volume

The latest data from a trolley "full run", a set of measurements encompassing all azimuthal locations in the storage ring taken over a short period of time, serves as a fundamental definition of the field. The trolley continuously cycles through the 17 probe array while taking measurements at 34 Hz, two full cycles per second. The full run process takes over an hour which yields at least 7,200 measurements with the full array. Each field measurement is accompanied by multiple location measurements accurate to about 1 mm. Data from a full set of trolley measurements, $\vec{\omega}_{tps}$ and $\vec{\theta}$, then feeds into an algorithm, f, which returns the magnetic field at a location $\vec{r}$ as in equation 3.10. The algorithm represents the 3-dimensional magnetic storage field. The allotted uncertainty for constructing $\omega_n(\vec{r})$ is 30 ppb.

$$\omega_n(\vec{r}) = f(\vec{r},\vec{\omega}_{tps},\vec{\theta}) \tag{3.10}$$

3.4.2 Interpolation

Between trolley runs, the magnetic field still needs to be well known. The uncertainty budget on the interpolation of the 3D field is 30 ppb. There are two subsystems which contribute to the field interpolation.

$$\delta\omega(t) = \delta\omega_{fps}(\vec{f}, t_i) + \delta\omega_{fs}(\vec{x}, \vec{b}, t_i) \tag{3.11}$$

The FPS provides one data source used for interpolation. There are 378 fixed probes around the ring which measure the field at ∼1 Hz. A pNMR frequency is extracted from each probe, and the data is timestamped. Some algorithm, f, for field interpolation will be designed using all fixed probe measurements taken since the last trolley run, such as

$$\delta\omega_{fps}(t) = \sum_{t_i=t_n}^{t_i<t} f(\vec{\omega}_i, t_i) \tag{3.12}$$

The second arm of magnetic field interpolation is the fluxgate magnetometers. The FS measures the transient magnetic fields at high rates, ∼1 kHz. There are three fluxgate sensors in the array which can be moved around the ring to monitor local magnetic field transients at multiple locations. The transient magnetic fields measured in the fluxgates, $\vec{b}$, need to be well correlated and calibrated with the magnetic fields measured in the fixed probes, but they are a potentially powerful tool for discerning transients in the field at moderate, common frequencies, such as 15 Hz or 60 Hz. Each sensor needs a human defined probe position vector denoted by $\vec{x}_i$, and provides a vector magnetic field value, $\vec{b}_i$, and timestamp, t_i. If the transient fields are small, as they are expected to be, the fluxgate will be used to place a limit on the effect of transient fields.

$$\delta\omega_{fs}(t) = \sum_{t_j=t_n}^{t_j<t} f(\vec{x}_i, \vec{b}_i, t_j) \tag{3.13}$$

3.4.3 Calibration

Calibration is the final step in producing an absolute magnetic field map. The calibration flows down from the absolute probe in two stages. The absolute calibration probe and the

plunging probe both measure the field in a certain location of a monitored and understood field. The relationship is parameterized by a chemical shift model as given in equation 3.14.

$$\omega_{aps} = (1 - \delta_{pps \to aps})\omega_{pps} \tag{3.14}$$

The second stage of calibration is similar to the first. The plunging probe measures the field in approximately the same volume as the trolley probes. The calibration relationship can then be transferred in a similar way. Equation 3.15 is given as an example.

$$\omega_{pps} = (1 - \delta_{tps \to pps})\omega_{tps} \tag{3.15}$$

Bringing the magnetic field map and the calibration function together, one arrives at an expression for the magnetic field map over time.

$$\omega_0(\vec{r}, t) = (1 - \delta_{pps \to aps})(1 - \delta_{tps \to pps}) \sum_n \Theta(t - t_n) \left[\omega_{ts,n}(\vec{r}) + \delta\omega_{fps}(t) + \delta\omega_{fs}(t)\right] \tag{3.16}$$

This is the culmination of all magnetic field measurements and analysis. It is referred to as the magnetic field map, but more precisely it is the free proton spin precession distribution as a function of time inside the storage volume.

3.4.4 Beam Convolution

The final step in determining the quantity which goes into the expression for a_μ, equation 1.5, is convoluting the proton precession distribution with the muon beam distribution. The quantity really needs to represent the average magnetic field experienced by the stored muons. The muon distribution is measured by Fiber Harp Detectors and reconstructed using the Tracker Detectors (see [14] for more detail). The muon-weighted free proton precession frequency can be determined by convolving the calibrated magnetic flux density distribution, ω_0, with the muon density distribution, M.

$$\langle \omega_p \rangle = \iint M(\vec{r}, t)\, \omega_0(\vec{r}, t)\, dt\, dV \tag{3.17}$$

Chapter 4

FID ANALYSIS

The FID signals generated in the pNMR probes contain frequency information that represents the effective field seen by protons in the probe's active volume. During a measurement the net polarization vector of the protons rotates into the plane orthogonal to the magnetic field, then the protons undergo Larmor precession due to the field. The average precession frequency of the protons immediately after the $\pi/2$ pulse is the target of FID analysis.

4.1 Frequency Extraction Methods

An example FID is given in figure 4.1. The main characteristics of the signal are amplitude oscillations at the Larmor frequency of the proton sample, and a decaying envelope function returning the signal back to the baseline noise.

As in most data analyses, the first step is cleaning and characterizing the data. In the case of FIDs, the process begins with removing the baseline (and trend in some cases). The next step in data preprocessing is determination of the start time and stop time of the useful FID signal within the frame of the entire recorded waveform. The baseline value is calculated using an initial segment of the waveform which is free of signal. The signal duration is computed using the maximum amplitude around the signal baseline and applying an amplitude threshold to find the first and last sections of the waveform with an envelope above that threshold. After the FID has been characterized, it is ready for frequency extraction.

4.1.1 Zero Crossing Method (ZC)

The simplest method for determining the FID frequency involves counting zero crossings of the signal. The technique was used in the previous $(g–2)_\mu$ experiment in a hardware

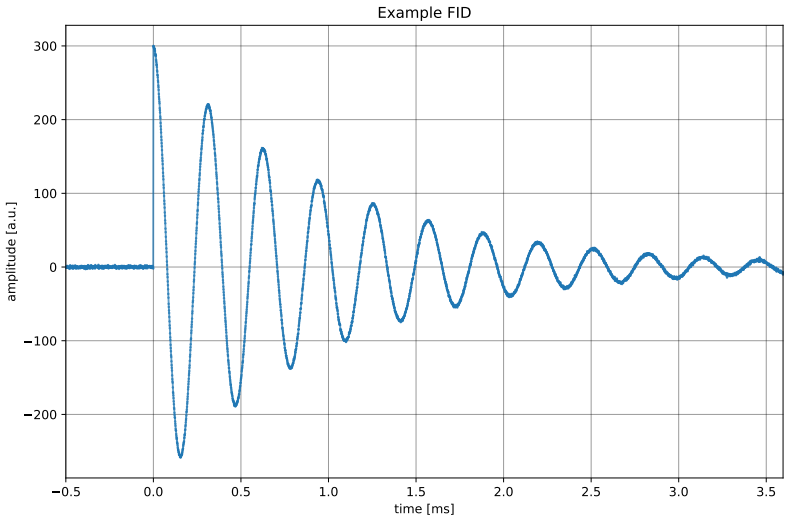

Figure 4.1: The example highlights the key characteristics of an FID waveform. The waveform starts abruptly after the $\pi/2$ pulse, oscillates at the Larmor frequency, and diminishes with some decaying envelope function which is often nearly exponential.

implementation, and has been re-implemented in software for E989. The zero crossing method is a time domain algorithm which counts all zero crossings in the signal (N), then determines the time of the first and last zero crossing using a polynomial fit (t_i and t_f). The FID frequency is then given by equation 4.1.

$$\omega_{zc} = \frac{1}{4\pi}\frac{N-1}{t_f - t_i} \tag{4.1}$$

The method prevents double counting of zeros due to noise fluctuations by requiring the signal cross an amplitude threshold before allowing another countable zero crossing. The threshold is referred to as the hysteresis threshold, and the concept is illustrated in figure 4.2.

The uncertainty associated with the zero crossing technique can be examined with a standard error analysis formulation given in equation 4.2 [34].

$$\frac{\delta f}{f} = \frac{1}{f}\sum_i \sqrt{\left|\frac{\partial f}{\partial x_i}\right|^2 (\delta x_i)^2} \tag{4.2}$$

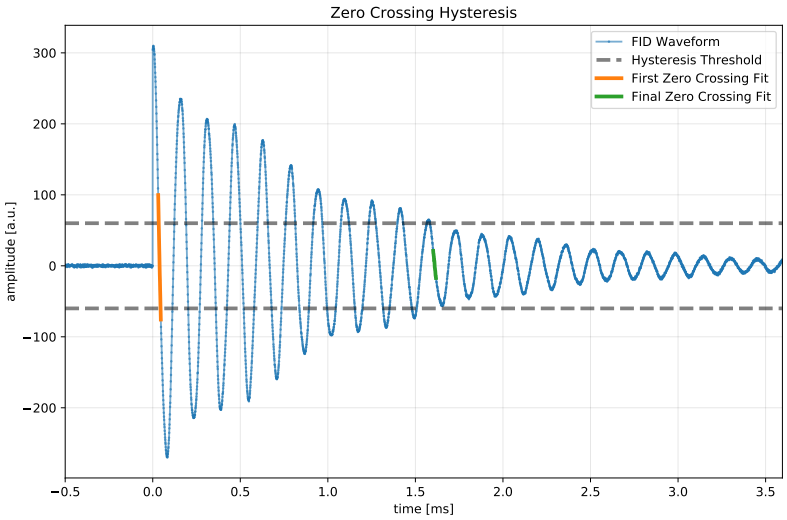

Figure 4.2: The figure illustrates aspects of the zero crossing frequency determination technique. The timing of first and last zero crossings are determined using a local polynomial fit which is represented by the orange and green lines. The number of zeros is incremented at a sign change, then frozen until the signal amplitude goes outside of the dotted gray lines. Using such a threshold is essential to the method as it suppresses erroneous zero counts.

In the case of the zero crossing algorithm, equation 4.2 yields:

$$\frac{\delta\omega}{\omega} = \sqrt{\left(\frac{\delta N}{N}\right)^2 + \left(\frac{\delta t_i}{t_f - t_i}\right)^2 + \left(\frac{\delta t_f}{t_f - t_i}\right)^2}. \tag{4.3}$$

For the first term in equation 4.3, larger values of N lower the uncertainty, and for the next two terms longer FID durations, $t_f - t_i = \Delta t$, lower the uncertainty. By using a proper hysteresis threshold, δN is suppressed to effectively zero. The start and stop times for the FID are determined by using a cubic fit around the first and last zero found in the algorithm. By assuming the waveform is an ideal FID (see equation 3.4) and expanding around the zero points, one finds that

$$\delta t = \frac{\delta y}{A\omega} = \frac{1}{S\omega} \tag{4.4}$$

where δy is the amplitude noise on the time domain signal and A is the signal amplitude. A substitution is then made to represent δt in terms of S the signal-to-noise in terms of

amplitudes. With the ideal FID assumption, δt_f can be expressed in terms of δt_i, $A_f = A_i e^{-(\Delta t)/\tau}$. Now combining equation 4.4, 4.3, and the preceding discussion, the uncertainty expression becomes

$$\frac{\delta\omega}{\omega} = \frac{1}{\omega}\frac{1}{S}\frac{\sqrt{1+e^{2\Delta t/\tau}}}{\Delta t} \tag{4.5}$$

It is informative to examine the uncertainty within typical limit values for the pNMR probe frequency. For a "bad" FID, the waveform might have a duration and decay constant of 100 µs, a frequency of 20 kHz, and a signal-to-noise of 100. In this case, the signal is expected to exhibit $\delta\omega/\omega \approx 3$ ppm. A "good" FID might have a duration and decay constant of 4 ms, a frequency of 50 kHz, and a signal-to-noise of 1000. In the good FID case, the frequency extraction is expected to exhibit $\delta\omega/\omega \approx 8$ ppb.

The zero counting technique is a robust and effective analysis method for FIDs. It was used in the previous $(g\text{–}2)_\mu$ experiment, and will be used as a frequency standard in E989.

4.1.2 Centroid Method (CN)

The centroid calculation is the simplest frequency domain analysis technique to extract the FID frequency. But before discussing the method, a few definitions are in order: the standard Fourier Transform and the discrete analogue of the Fourier Transform will be used to bring the signal into the frequency domain, and they are defined in equation 4.6 and 4.7 respectively.

$$\hat{f}(\omega) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-i\omega t}\, f(t)\, dt \tag{4.6}$$

$$X_k = \frac{1}{\sqrt{N}}\sum_{n=0}^{n=N} e^{-i2\pi kn/N} x_n \tag{4.7}$$

In equation 4.7, X_k is the component of the Discrete Fourier Transform representing the frequency $k/(N\Delta t)$ where N is the number of samples in the waveform, k is an integer from $-N/2$ to $N/2$, and Δt is the sampling period of the waveform.

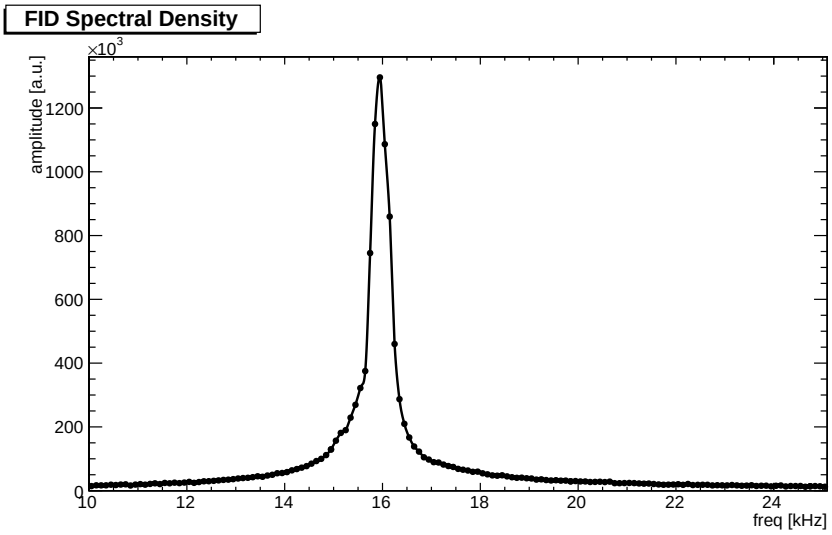

Figure 4.3: The plot depicts the spectral peak of a typical measured FID. The peak is asymmetric as the low frequency side exhibits a shallower curve. The peak is also fairly narrow as it only contains 20 or so points significantly above the noise and only 5 or 6 points above half the peak amplitude. Figure reproduced from reference [3].

Also important, the power spectral density (PSD) is the modulus squared of the Fourier Transform which is useful because it combines the real and imaginary frequency information into a single vector.

$$PSD(\omega) = |\hat{f}(\omega)|^2 \tag{4.8}$$

In the PSD, defined in equation 4.8, the FID frequency manifests as the peak amplitude. In the centroid method the frequency is computed as a frequency weighted sum of bins around the peak PSD value.

$$\omega_{cn} = \frac{\sum \omega_i |X_i|^2}{\sum |X_i|^2} \tag{4.9}$$

The peak in the PSD from FIDs is slightly asymmetric and fairly narrow as shown in figure 4.3, and the centroid method does not accommodate such deviations very well. It will not be a primary method in E989.

4.1.3 Analytical Peak Fit (AN)

For the idealized FID waveform defined in 3.4, an analytical solution for the spectral peak is calculable. The integral, equation 4.10, is fairly straightforward, since all the components can be converted into exponentials. Note that the Θ in equation 4.10 is the Heaviside step function.

$$\hat{f}(\omega) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-i\omega t}\left[\Theta(t)\left(A\sin(\omega_0 t - \phi_0)e^{-t/\tau}\right)\right]\, dt \tag{4.10}$$

$$\left|\hat{f}(\omega)\right|^2 = \left(\frac{A}{2\pi}\right)^2\left[\frac{(\gamma^2+\omega^2)\sin^2\phi_0 + 2\gamma\omega_0\sin\phi_0\cos\phi_0 + \omega_0^2\cos^2\phi_0}{(\gamma^2+\omega^2)^2 + 2(\gamma^2-\omega^2)\omega_0^2 + \omega_0^4}\right] \tag{4.11}$$

The analytical solution did not prove to be a robust basis for peak fitting with narrow width of the spectral peaks and the multi-variable nature of the expression. Still, it is a useful step into the next method. If one expands the solution in 4.11 around $\omega \to \omega_0 + \delta\omega$ and employs the assumptions: $\delta\omega \ll \omega_0$ and $\gamma^2 \ll \omega_0^2$, then the expression simplifies to equation 4.12.

$$\left|\hat{f}(\omega)\right|^2 \approx \left(\frac{A}{4\pi}\right)^2 \frac{1}{(\omega-\omega_0)^2+\gamma^2} \tag{4.12}$$

The simplified expression is an improperly normalized Lorentzian Distribution (or Cauchy Distribution as it is also called).

4.1.4 Lorentzian Peak Fit (LZ)

Peak fitting routines can improve in precision upon the centroid and more complex analytical techniques. Peak fitting works well using a Lorentzian Distribution, equation 4.13, in the domain around the maximum frequency bin. As shown in 4.1.3, the Lorentzian is an approximation of the analytical solution.

$$F(\omega) = \frac{1}{\pi}\frac{\Gamma/2}{(\omega-\omega_0)^2+(\Gamma/2)^2} \tag{4.13}$$

The Lorentzian peak fitting method is more stable and effective than other frequency domain methods, but not the most effective method. The expansion assumptions and ideal

waveform approximations are not always valid in FID signals, and those signals are handled worse by the Lorentzian method.

4.1.5 Polynomial Phase Fit (PH)

The most precise frequency determination technique involves calculating the phase of the FID waveform. Given the assumption that the original signal is harmonic, the Hilbert Transform (equation 4.14) gives the imaginary complement of a signal. Using the original signal and the Hilbert transform together, the phase of the FID is given by equation 4.15.

$$h(t) = IFFT[-i \cdot \mathrm{sgn}(\omega)\hat{f}(\omega)] \tag{4.14}$$

$$\phi(t) = \arctan(h(t)/f(t)) \tag{4.15}$$

The phase calculation using equation 4.15 does have the caveat that the domain is always $[-\pi, \pi]$, but unfolding the phase into the full domain is straightforward using thresholds on the allowing jump in phase between adjacent points. With a phase robustly determined for the signal, the FID frequency manifests as the linear slope of the phase at time zero, $\frac{\partial \phi}{\partial t}|_{t=0}$.

$$\omega_{ph} = a_1 \in \underset{a_0,a_1,a_2,a_3}{\mathrm{argmin}} \; |\phi - \left(a_0 + a_1 t + a_2 t^2 + a_3 t^3\right)| \tag{4.16}$$

A polynomial fit such as equation 4.16 is effective at extracting the frequency, because the frequency is just the linear coefficient. In fact, the phase information yields frequency as a function of time which is interesting for systematic checks. The phase fit method is the most precise method in simulated FIDs.

4.1.6 Sinusoid Fit

Normalizing the signal into a sinusoid is another approach for frequency extraction utilizing the Hilbert Transform. The envelope is calculable as the norm of the original signal and the Hilbert Transform

$$f_{env}(t) = \sqrt{|f(t)|^2 + |h(t)|^2}. \tag{4.17}$$

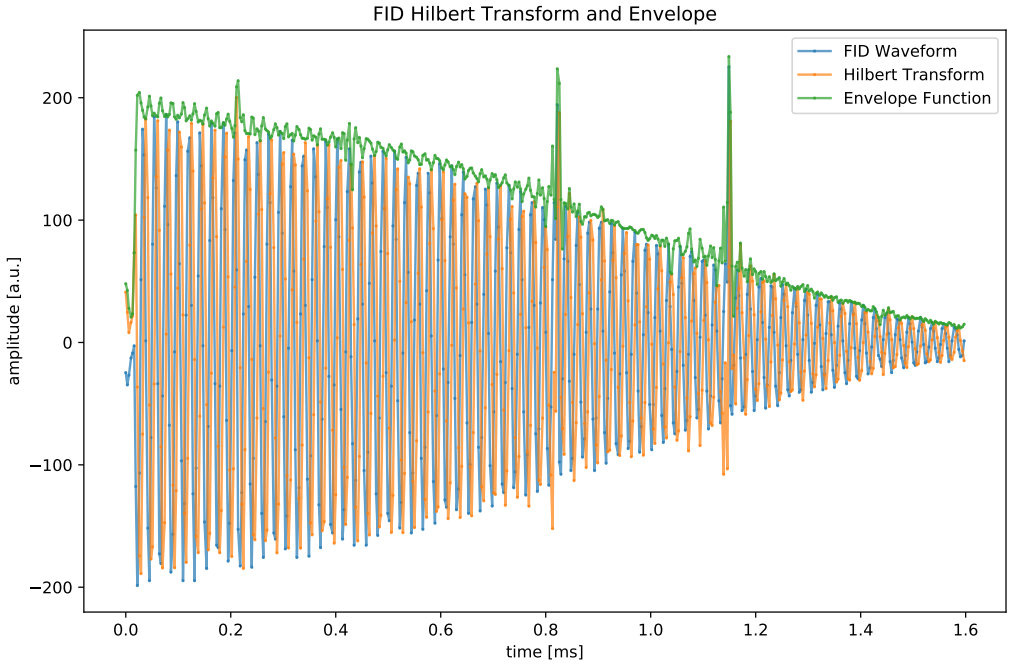

Figure 4.4: The plot shows the power of the Hilbert Transform. The original FID is depicted alongside its phase complement, the Hilbert Transform. And the magnitude of the two vectors together yield the third vector in the plot, the envelope function.

With the envelope in hand, the original signal can be normalized into an amplitude 1, sinusoidal signal. The frequency is then extracted by using a minimization routine, such as least squares.

$$\omega_{sn} = \underset{\omega}{\operatorname{argmin}} \left| f(t)/f_{env}(t) - \sin(\omega t - \phi_0)) \right| \tag{4.18}$$

4.2 Frequency Extraction Studies

Each FID frequency extraction technique must be tested and validated. The tests are performed with several different sources of data: ideal FIDs, simulated FIDs, and measured FIDs. Each dataset presents a test bed for determining the precision, accuracy and failure modes of each frequency extraction technique.

4.2.1 Ideal FIDs

The simplest case is idealized FIDs. The frequency is static in time and all techniques should perform well. If the technique does not perform with high precision on the simplest FIDs, then it is not viable for analyzing the E989 data. The effectiveness was tested on thousands of FIDs using a flat random distribution of frequencies from 46.8 kHz to 47.2 kHz. The test results are shown in 4.5.

4.2.2 Simulated FIDs

Integrating the Bloch Equations (equation 3.3) is a straightforward path which yields accurate, simulated FID waveforms. The waveforms are similar to the ideal FIDs with some possible deviations from the non-perfect, finite $\pi/2$ pulses. The simulation used a lower Larmor frequency of around 1 MHz as opposed to the real Larmor frequency of 61.79 MHz. The reason being that it is more practical to simulate at lower frequencies, because the integration step can be much larger. In the end, the entire system runs through a low-pass filter at 200 kHz, so all the higher frequency content is removed from the final FID signal. In that light simulations at 1 MHz are perfectly adequate. Noise is added to the simulated waveforms at 1/300 of the FID signal amplitude for frequency extraction tests. Example

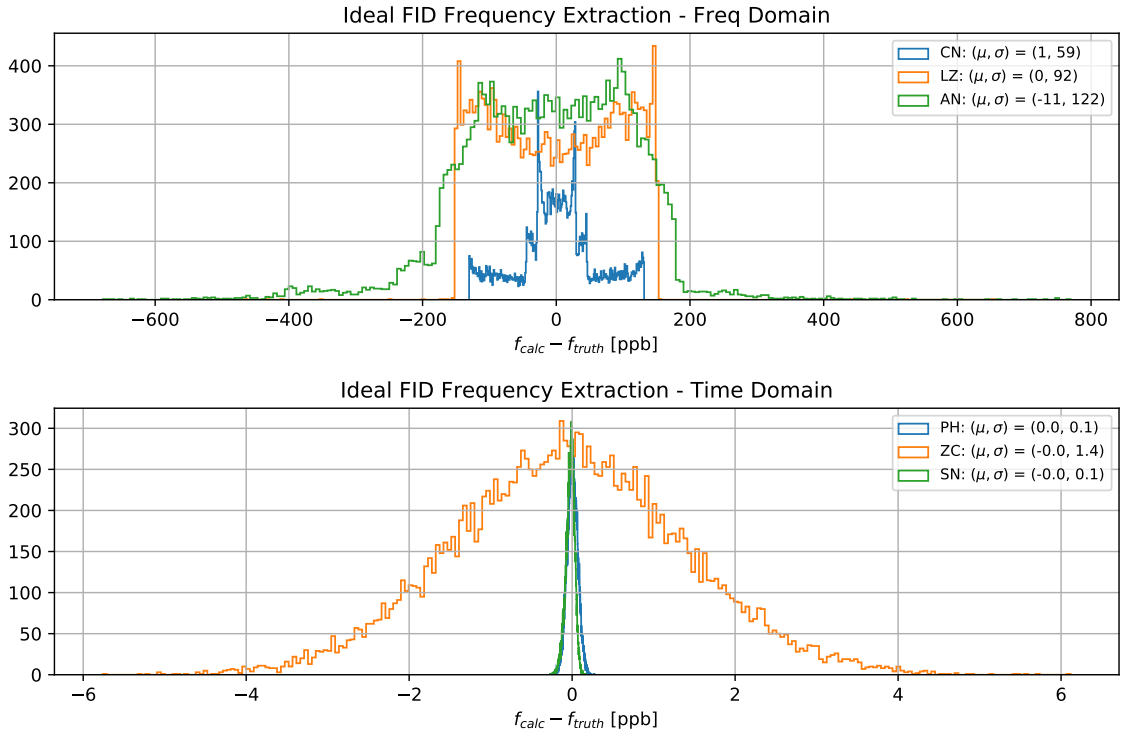

Figure 4.5: The plots depict precision tests of different FID frequency extraction on ideal FIDs. The methods are separated into frequency domain on the top, and time domain on the bottom. The time domain methods (zero crossing method (ZC), phase fit (PH), and sinusoid fit (SN)) all outperform the frequency domain methods (centroid method (CN), Lorentzian peak fit (LZ), and analytical fit (AN)) in this study, and in particular the sinusoidal fit and linear phase fit perform exceptionally.

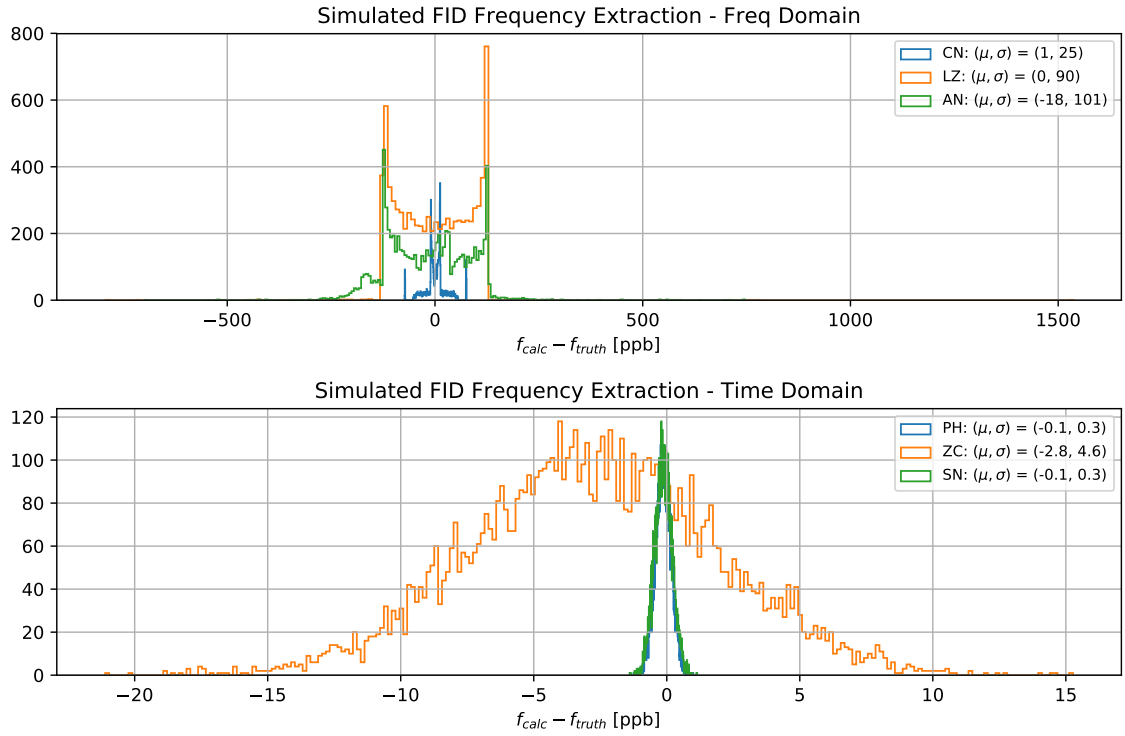

Figure 4.6: The plots exhibit tests of different FID frequency extraction methods on FIDs simulated by integrating the Bloch Equations. The methods are separated into frequency domain on the top, and time domain on the bottom. The time domain methods (zero counting method (ZC), phase fit (PH), and sinusoid fit (SN)) all outperform the frequency domain methods (centroid method (CN), Lorentzian peak fit (LZ), and analytical fit (AN)). As in the ideal FID study, the sinusoidal fit and linear phase fit perform very well in this study.

results are shown in figure 4.6.

4.2.3 Simulated Gradient FIDs

An anticipated complication with actual magnetic field measurements arises from the presence of gradients in the field. Any real magnetic field will have gradients over the volume of the probe which will distort the signal of the probe. These effects can be mimicked using a superposition of simulated FIDs over a small range. One notices node-like behavior in the envelope when gradients are present (and E821 waveforms exhibited this behavior). The

reason for node-like behavior is that individual FIDs all begin in phase but as different frequencies are evolving they lose phase coherence and cancel each other at one time. Then, as the phases continue to advance, the collection of spins can re-cohere at another time. The effect is similar to the amplitude fluctuations that accompany beat frequencies. For a purely linear gradient all components can cancel rendering a real zero in the FID amplitude. The simulations aim to test the effects of small gradients on the FID frequency extraction precision and accuracy.

The range of magnetic field gradient over the pNMR probes anticipated in the g–2 storage field are on the order of 10 ppb to 100 ppb. A collection of simulated FIDs was created using a Bloch Equation integrator with a range chosen to be 47 kHz ±1000 ppb. To create a gradient FID, around 20 waveforms from a collection of simulated FIDs were put into a weighted sum. The weights were chosen in a way that did not change the average frequency only the variation over the probe. For linear gradients this is simple, it is just the central value and evenly spaced set of frequencies on both sides. For instance, at 100 ppb and 5 waveforms, the gradient FID would be composed of the following frequencies: −100 ppb, −50 ppb, 0 ppb, 50 ppb, and 100 ppb (+ 47 kHz). For quadratic gradients, the weights are a bit less intuitive, but still simple enough. The frequencies were chosen to be even spaced around the central frequency again. As an example, let x be the linear distribution defined previously and the weights for the quadratic gradient will be defined in 4.19.

$$y = \frac{x^2 - \langle x^2 \rangle}{\max\left(x^2 - \langle x^2 \rangle\right)} \times 100 \text{ ppb} \tag{4.19}$$

The following frequencies would make up the quadratic gradient: 100 ppb, −50 ppb, −100 ppb, −50 ppb, and 100 ppb (+ 47 kHz).

The frequency extraction results from the gradient FIDs are presented in figures 4.8 & 4.9. In the linear case, the phase fit is essentially unperturbed while the zero crossing method gains uncertainty over the 1000 ppb gradient applied. The effect in the expected gradient range of ∼100 ppb may seem large, but bear in mind that 100 ppb is the largest expected gradient and the vast majority of pNMR measurements will be made in regions with much smaller gradients. In the quadratic case, much more dramatic deviations from

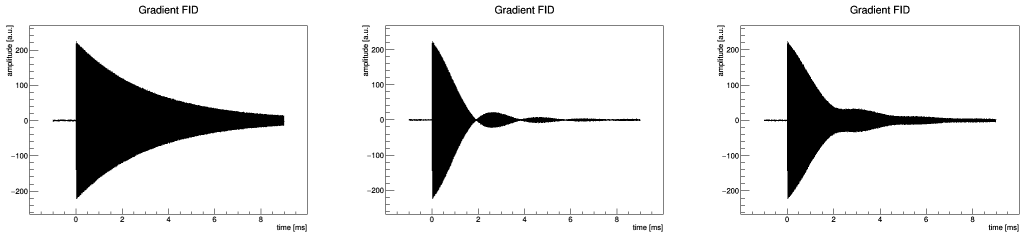

Figure 4.7: Example simulation waveforms made by superimposing a gradient of individual waveforms. The left waveform has no applied gradient. The middle has a linear gradient of 1000 ppb. Note that the beat frequencies put nearly complete nodes onto the waveform envelope. The right waveform has a quadratic gradient of 1000 ppb. Note the softer waveform distortions.

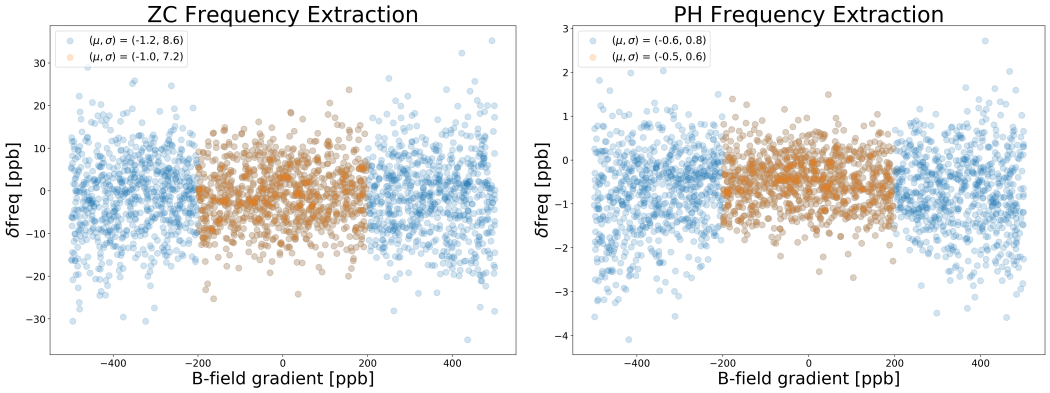

Figure 4.8: The plot depicts the deviation from the true input frequency in linear gradient FIDs. The deviations are negligible for the phase fit method, and small for the zero crossing method but not negligible at 1000 ppb.

the true frequency are present, but again the effects are small in the expected FID gradient domain. The effect at 200 ppb in all cases is less than 10 ppb.

4.2.4 FID Measurements

A stability dataset of waveforms from the field shimming stage was also available to test the FID algorithms. Testing against this dataset had a real difference in that the true precession frequency was not known. One can still build distributions to get a sense for the precision and stability of an analysis method.

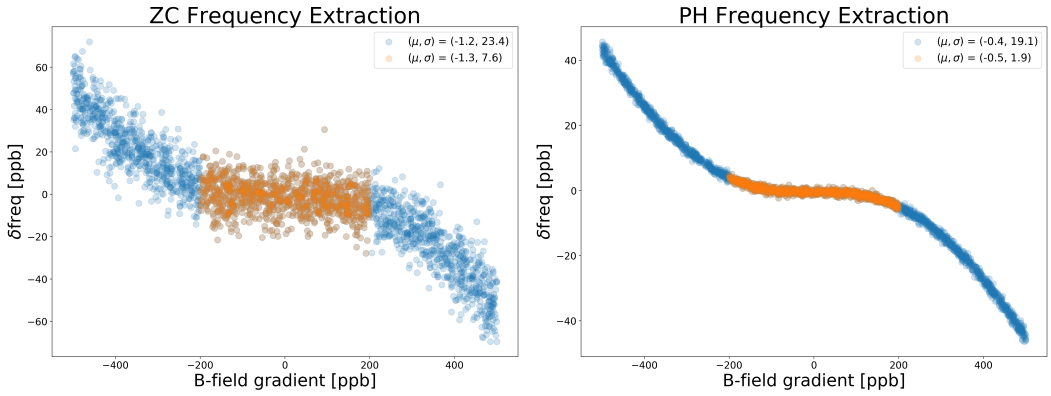

Figure 4.9: The plots show the deviation from the true input frequency for quadratic gradient FIDs. The deviations are much larger than the linear case, but still negligible at the expected gradients of a few hundred ppb.

Figures 4.10 & 4.11 show the analysis results for real FIDs. The data was first detrending by removing the average behavior over all 25 pNMR probes then removing polynomial drift up to third order. Zero crossing and phase appear similarly effective in the context of real data. Both methods have a precision of ∼10 ppb in the study. The phase fit central value has a shift of ∼15 ppb.

4.3 FID Frequency Results

Table 4.1: FID Analysis Summary

Data Type	Zero Crossing	Phase Fit
Ideal	1.4 ppb	0.1 ppb
Simulated	4.6 ppb	0.1 ppb
Linear Gradient	7.2 ppb	0.6 ppb
Quadratic Gradient	7.6 ppb	1.9 ppb
Measured	9.3 ppb	8.9 ppb

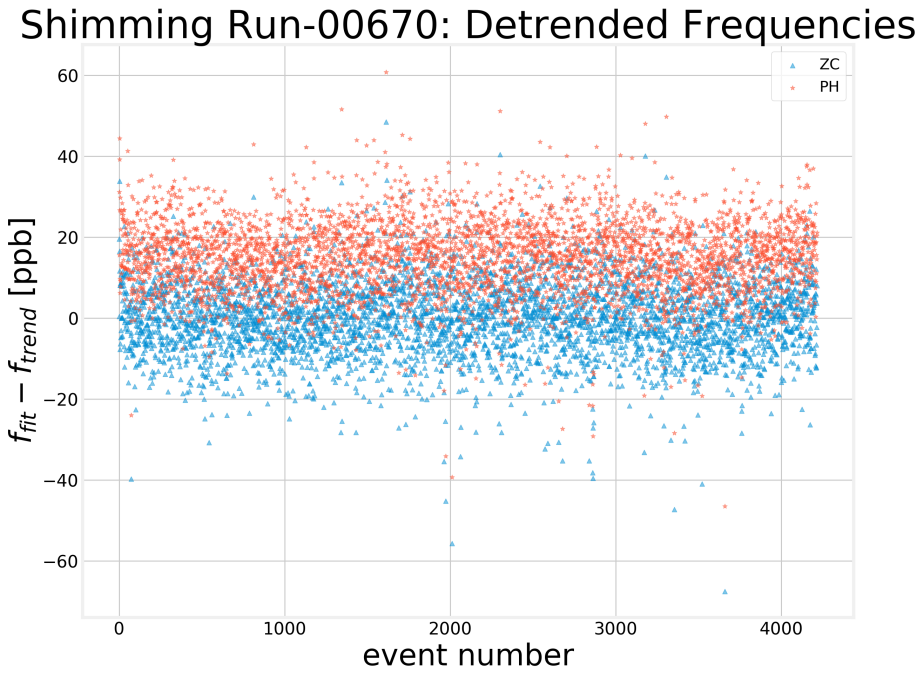

Figure 4.10: The detrended plot was produced by averaging all 25 probes in the shimming platform. Then, subtracting off the average from the central probe value. Lastly, fitting the difference to a cubic and subtracting off the cubic trend. A small trend remains in the data though, so detrending method was not completely successful. The result shows a small difference between the zero crossing and phase fit methods, but the distributions are within agreement.

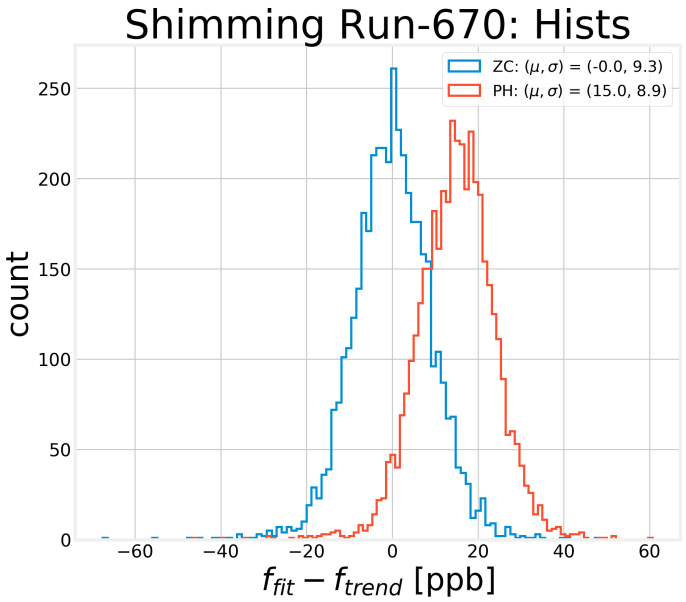

Figure 4.11: The resulting histogram of two different frequency extraction methods after detrending the signal. The two distributions nearly agree with a small shift in the central value. Interestingly, the phase fit method and zero crossing show similar precision on real data.

Chapter 5

SHIMMING THE MAGNETIC FIELD

5.1 Shimming Procedure

The g–2 magnetic field is subject to two types of improvements on the path to optimal uniformity. The first type is called "passive" or "rough" shimming. It occurs after the storage ring is reassembled and before the installation of the vacuum chambers in the magnet gap and detector subsystems has occurred. The rough shimming entails precise adjustments of physical, static shims. The second type is "active" or "fine" shimming. Active shimming is ongoing throughout the duration of the experiment. It uses dynamically controllable sources of magnetism such as electric currents for fine adjustments to the magnetic field uniformity.

5.1.1 Shimming Goals

The uniformity goals for field shimming are laid out in detail in the TDR [14] and summarized here. The overall uniformity goal on the azimuthally averaged storage magnetic field restricts deviations from the mean to less than 1 ppm. For each higher order multipole, the amplitude goal is 100 ppb when averaged. For reference, the field from the previous experiment is shown in figure 5.1. The E989 magnetic field goals are essentially a factor of two more uniform than E821.

The uniformity goals for rough shimming are less strict than the overall shimming goals. The focuses for rough shimming are twofold. One is on minimizing gradients and deviations around the ring in all multipoles but especially the dipole. The general limit for the first goal is a peak-to-peak variation of 50 ppm for the dipole field and 10 ppm for the higher order multipoles. The other goal is to bring the higher order moments within range of the active shimming hardware. This goal aims for the azimuthally averaged value of each higher order multipole be less than 5 ppm.

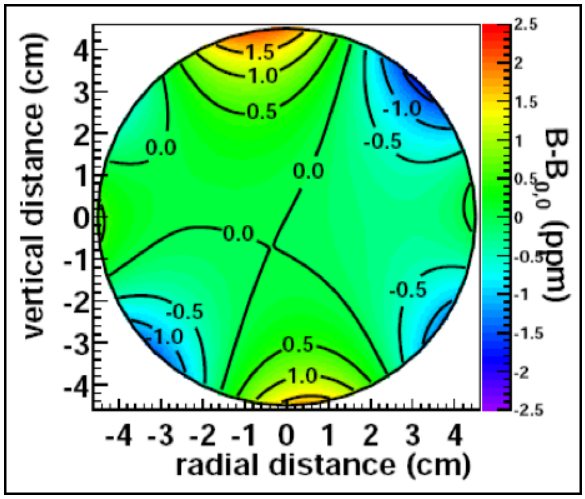

Figure 5.1: The plot depicts the final magnetic field for the E821 g–2 experiment averaged over azimuth. The 4.5 cm circle shown is the cross-section of the muon storage volume. The uniformity goals for E989 push the entire field to have deviations smaller than 1 ppm in the azimuthal average.

5.1.2 Passive Shimming

The passive shimming process makes large, static changes to the magnetic storage field. It uses the intrinsic shim kit of the storage ring and additional shims as needed. Each different component of the shim kit must be well understood, so the shims undergo a calibration procedure. The operating procedure is to measure the magnetic field as a reference, adjust the shim of interest, re-measure the magnetic field, subtract the adjusted magnetic field data from the reference magnetic field data, and use the difference in magnetic field to build a model of the shim effects. The calibration process is repeated for each shim type until a feasible model can be constructed. With a working model for each type of shim, an optimization plan for the entire storage field can be calculated.

The implementation of optimization routine occurs in multiple stages as well. The first stage is physical adjustments of the largest pieces of the storage ring, the yokes and the pole pieces. Next, the shims designed to affect higher order multipoles need to be adjusted properly. And, finally the lowest order multipoles including the dipole and normal quadrupole receive shim compensation. That was the planned procedure based on E821

reports and the magnet design, but the field team found that one more stage was needed to reach the uniformity goals. The final stage added laminations of small steel foils all around the magnet.

5.1.3 Active Shimming

The active shimming stage makes dynamic changes to the azimuthally averaged field. This phase of the experiment will continue as long as the experiment is running. The active shimming has two essential tools. The first component is a set of 100 wires spaced across the surfaces of the top and bottom pole. Called the surface current coils, these wires drive controllable currents which adjust the azimuthally averaged multipole content of the magnetic field within several ppm. The second component is active feedback to the magnetic field power supply to minimize field drift. The power supply feedback influences the set point of the magnet's main current supply which allows it to effectively set the overall strength of the magnetic field. The active shimming will not be discussed further in this document as it is not yet completed work like the rough shimming.

5.2 Shimming Hardware

With the rough field shimming plan outlined, the shimming discussion continues with details on the magnetic design shims and tools available. The storage ring has already been described in some detail in chapter 2. The discussion here will begin with hardware in the storage magnet's built-in shimming kit, and follow with an enumeration of the measurement devices used in the shimming process.

5.2.1 The Shims

The storage ring was engineered with several built-in knobs to tune the field locally in azimuth. Each of the shims was tailored to have leverage on different complementary aspects of the field. Figure 5.2 gives a visual overview of the shims.

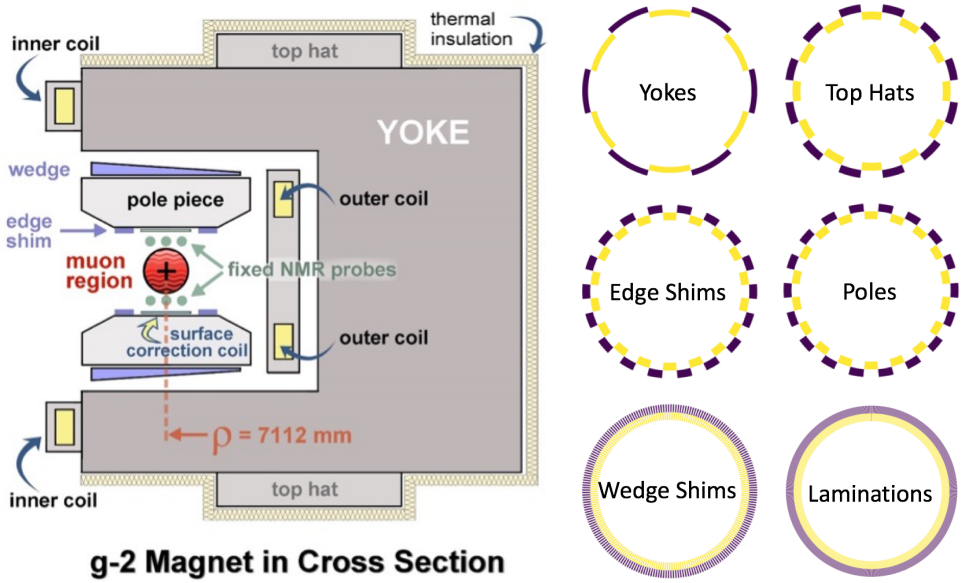

Figure 5.2: The cross-sectional view on the left exposes the location of all the shimming hardware. The top hat, edge shims, and wedge shims are shown in context with the pole and yoke hardware. The ring diagrams on the right exhibit the azimuthal frequency of each piece of hardware: 30° yoke sections, 15° degree top hats, 10° degree pole pieces, 10° edge shims, 0.833° wedge shims, and 0.244° spacing on the iron laminations. Note that the laminations are not part of the built-in shimming kit.

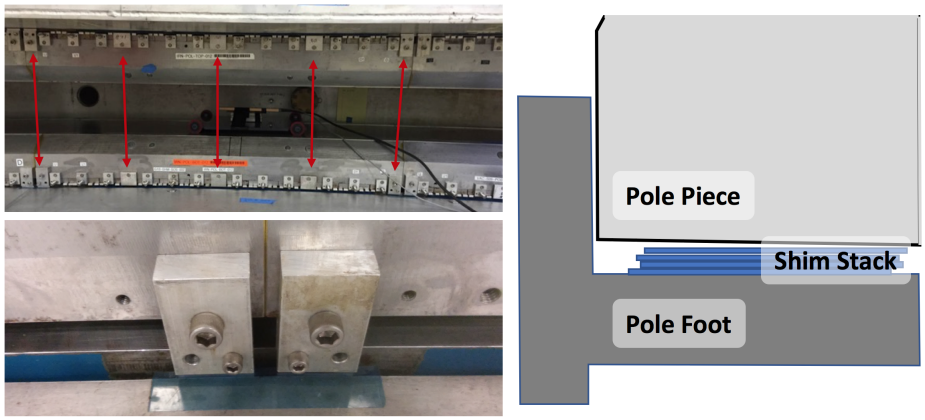

Figure 5.3: Images and diagrams highlighting the pole feet. The top image shows full poles which have pole feet at the interface between the pole and the yoke. The locations of the pole feet are denoted by red arrows. The bottom image zooms in on two adjacent pole feet. The diagram depicts the shimming method used to adjust the poles. The feet were removed and a stack of thin metal shim stock pieces was inserted to elevate that region of the pole surface.

Pole Feet

The uniformity of the gap between the upper and lower pole pieces determines the magnetic field strength between two poles to first order. The pole pieces interface with the yoke on adjustable mounts dubbed "pole feet". Adjustments to the pole feet can alter the plane of the pole pieces by up ∼200 µm with precision of ∼5 µm. The pole feet also offer a small amount of leverage to reshape the pole surface itself. Figure 5.3 depicts the pole feet and the method for using the pole feet to adjust the pole pieces.

Top Hat Shims

The simplest shims are steel plates called the top hats. See figure 5.2 for a diagrammatic depiction. The top hats run in 15 degree segments, two per yoke, top and bottom. Adjusting the top hats can induce a large effect on the average field in its respective region without affecting the higher order multipoles. Adjustments are made by adding plastic spacer shims between the yoke and the top hats. Adding material allows more magnetic flux to leak and decreases the magnetic field. The size of the spacer stack for top hats was limited to 2 mm.

Edge Shims

Each pole piece has two steel runners on the outside of its uniform, flat surface. With top and bottom sets of so-called edge shims, the design imparted a complex handle on many different higher order multipoles. Adjustments to the edge shims in E821 were done by grinding the initially over-sized edge shims down to a custom size. Two other avenues exist for edge shim adjustments. The effect of the edge shim could be augmented by addition of thin (25 µm) steel shim stock cut to the same shape as the edge shim, or the effect could be diminished by similar addition of a plastic shim layer. All options were evaluated for E989. Figure 5.4 shows the edge shim model from an OPERA2D magnetic field simulation.

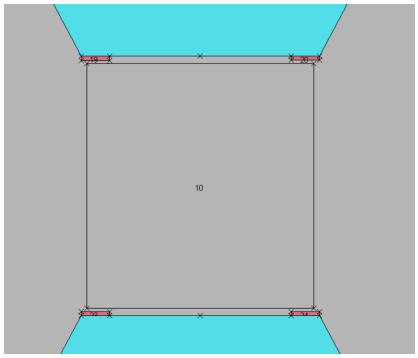

Figure 5.4: The image depicts four edge shims at the edges of an upper and lower pole piece. It was rendered using an OPERA2D model of the storage ring.

Wedge Shims

The storage ring contains 864 wedge shims, i.e., 12 per pole piece, which occupy the space between the pole pieces and the main yoke. They have an angled design in order to effect both the quadrupole and dipole fields. The wedge shims were designed to be 50 mm thick at their inner radius, 165 mm thick at their outer radius, and 530 mm in length giving an inclination of 1.24° [9]. Figure 5.5 shows an image of a wedge shim in context between the yoke and the pole piece.

Figure 5.5: The image depicts a real wedge shim between the yoke in the pole piece. The protruding bolt on the left is used to adjust the edge shim radial position. The white spacer is used to ensure the wedge shim position is secure through power cycles of the storage magnet.

Laminations

The final set of shims utilized in rough shimming were a step beyond the magnet design. To achieve the rough shimming uniformity goals, a set of small, thin steel strips were affixed to a thin sheet of plastic which was then attached the pole surface. The laminations were massed, and placed at 0.244° intervals at three different radii all around the storage ring. In regions where fixed probes were known to be, the central lamination strips were aligned in the azimuthal direction rather than the radial direction to lessen gradients detrimental to the fixed probe FID signal. A mock up of a single pole's worth of laminations is depicted in figure 5.6.

5.2.2 Measurement Devices

Shimming Cart

The primary tool used in the rough shimming process was a mobile platform engineered from low-magnetism materials to make measurements throughout the muon storage region. The so-called "shimming cart" used the lower pole surface as a track. A rendered model of the shimming cart is shown in figure 5.7. The cart was equipped with magnetometers, thermometers, distance sensors, and laser tracker beacons to measure the magnetic field,

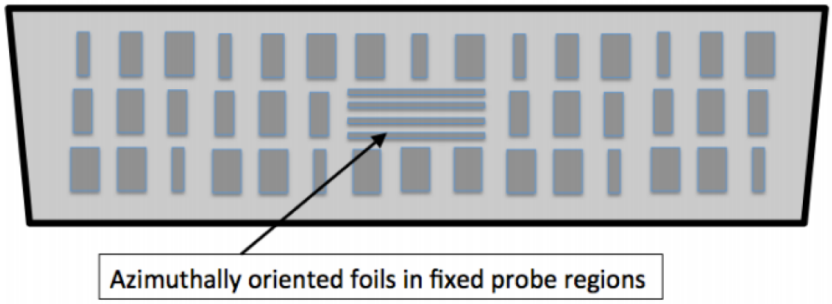

Figure 5.6: A simplified diagram of the lamination patterns laid down on pole pieces. The real laminations had 41 azimuthal regions and two picket fences per pole piece.

the local temperature, the pole gap and relay its relative position in the storage ring.

All materials were chosen to make as small of a magnetic perturbation as possible. The platform and wheels of the cart were constructed from PEEK, the side walls were made from quartz, and fixed probe matrix was made of aluminum. The cart traversed the ring with the help of a stepper motor and flexible cabling to make steps forward. A typical step size was 3 mm, and the probes and Capacitec performed measurements at each location before the cart moved on. A full scan of the ring with the standard step size took about 3 hours.

pNMR Probes

The magnetic field was measured through a matrix of 25 pNMR probes secured in a frame which is aligned in the azimuthal plane of the storage volume. The probes were arranged with a single probe in the center of the storage region, an 8 probe ring at a radius of 22.5 mm from the center probe at angular increments of $\pi/4$ starting on the negative y-axis, and a 16 probe ring at a radius of 45 mm from the center probe at angular increments of $\pi/8$ The active volume of the pNMR probes extends 10 mm in the azimuthal direction and 2 mm in the radial direction. The 25 probes are aligned within 1 mm in azimuth and much better within the azimuthal plane. The probe pattern can be viewed in figure 5.7, and it facilitated measuring the full cross-section of the muon storage volume with sensitivity to the first 17

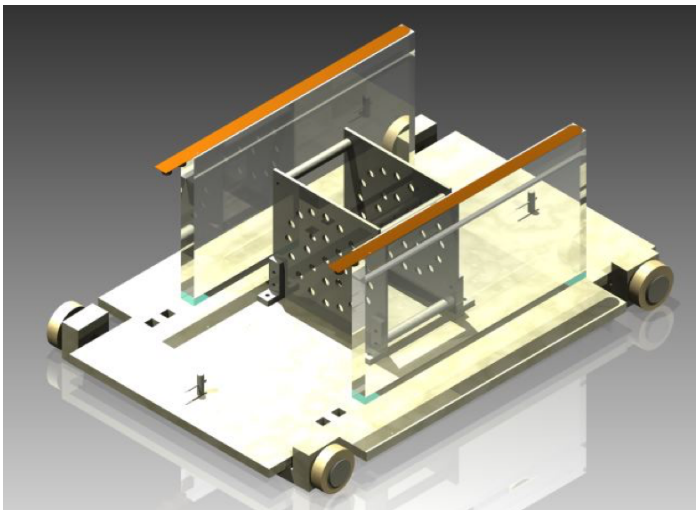

Figure 5.7: The image is a rendered model of the shimming cart. In the center of the image, the pNMR matrix which holds 25 probes at three different radii is shown. On each side of the probe matrix, there are two quartz plates with a copper colored strip on the top. That strip is the Capacitec sensor and sensors are also present on the underside of the quartz. The wheels and the base plate were fabricated out of PEEK to minimize magnetic perturbations. On each corner of the inside quartz plate, a corner cube reflector was fastened to the real shimming cart to use as lock points with the laser tracker. A temperature sensor was mounted onto the base plate behind the pNMR probes.

2D magnetic field multipoles (dipole plus 8 normal and 8 skew multipoles).

Capacitec Sensors

The cart also possessed four capacitive distance sensors which measured the spacing from the sensor to the pole surface. The devices were model 4100 SL Capacitec Sensors, and they were sensitive to distances as small as 5 µm. The sensors were on a ribbon cable which was fixed to the top and bottom of each quartz plate giving readings for relative distance to the pole surface at an inner radius of 7002.5 mm and an outer radius of 7221.5 mm. By summing the readout from each sensor and calibrating the size of the quartz, the Capacitec system produced data to define the pole gap, the distance between each pole surface.

Temperature Probe

An Analog Devices 590 temperature probe was placed on the cart near the pNMR probes. The device has potential to correlate signal distortions of the pNMR frequencies, and drift in the magnetic field due to temperature changes. The probe could also be compared against several other temperature probes operating around the experimental hall. The temperature data was used minimally though as it had minimal utility.

Laser Tracker

A commercial, API laser tracker was used to follow the location of four reflective corner cube mirrors mounted at the corners of the inner quartz plate (see figures 5.7 and 5.8 for context). Installed in the center of the ring it was used to determine the cart position for each measurement taken with the shimming cart. The most important value from the laser tracker was the azimuthal angle, ϕ, but the device also reported height, z, and radius, r. At the center of the storage ring, the laser tracker measured the position of corner cubes on the shimming cart with the precision, $\delta\phi \approx 0.001°$, $\delta r \approx 3.5\,\mu m$, and $\delta z \approx 15\,\mu m$. With imperfect alignment though, the tilt of the laser tracker made these uncertainties into $\delta\phi \approx 0.1°$, $\delta r \approx 25.0\,\mu m$, and $\delta z \approx 25\,\mu m$

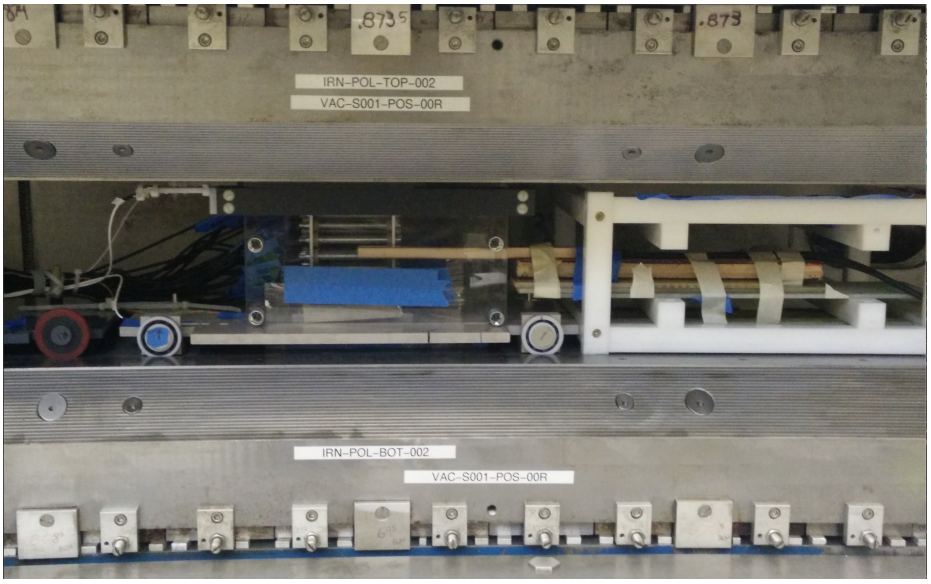

Figure 5.8: The image shows the shimming cart and the Metrolab NMR probe (with carrier structure) during a calibration cross-check of the shimming cart magnetic field measurement values. The magnetic field reported by each system differed by only a few ppm.

Metrolab

The Metrolab PT 2025 is a wide-range commercial NMR system which complements the shimming cart probe measurements. The "Metrolab" as it was referred to had an absolute accuracy of 5 ppm and a precision of 0.1 ppm. A probe carrier structure was fabricated to stably and reliably align the probe inside the magnet gap as shown in figure 5.8. The commercial NMR device was used to verify that the magnetic field was near the target value of 1.4513 T. And, it was also used to monitor field drift.

Tilt Sensor

The tilt sensor was a heavily utilized piece of equipment. The device contained two Frederick's 0719 Tilt Sensors in a custom carrier structure. Researchers from the University of Washington designed the setup and compiled the two mounted commercial, electrolytic bubble levels onto a 29 cm by 12 cm aluminum base plate. The aluminum plate had three spherical feet both top and bottom to define a consistent plane. The tilt sensors were able

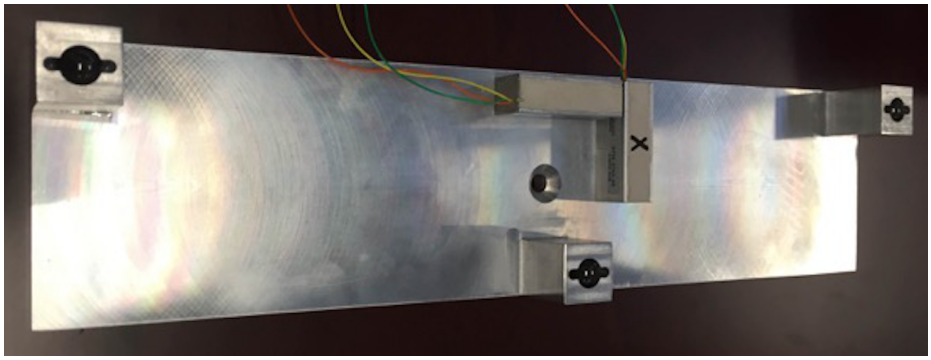

Figure 5.9: The image shows the tilt sensor apparatus used in rough shimming. The device had two tilt sensors on orthogonal axes. The tilt sensors are the small, long boxes in the center of the aluminum base plate. The three spherical feet define a consistent plane for the apparatus to rest on any surface. They are located with one at each long end and one in the center.

to read out at a precision of about 2 µrad, but the sensors took more than 30 minutes to fully settle into the highest accuracy reading. The majority of measurements were made with a lower precision of 15 µrad as a compromise to reduce measurement time down to a more manageable 3 to 5 minutes.

5.3 Adjustments

From the start, each step in the plan for rough shimming was well defined, though the necessity of each step was not as clear. After calibrating the measurement devices and subsequently the state of the storage ring and storage field, it was decided which stages of adjustments were necessary. The first stage involved leveling the pole surfaces with respect to gravity. The second was a full iteration of adjusting the feet on each pole piece. The third was a survey of the necessity of high order multipole adjustments via edge shims. The fourth was full optimization of the field using the built-in shimming kit. And, the final round of adjustments implemented a new system of shims, over 8,000 specifically massed iron foils, which really pushed the field uniformity to a new level.

5.3.1 Ring Leveling

At an intermediate point in the shimming process, it became clear that the plane of the storage ring and the plane defined by gravity were not aligned. Furthermore measurements indicated that plane of the ring was different from the plane that had been measured upon re-assembly of the ring. It would seem that the concrete floor of the experimental hall had shifted by a millimeter or two. In addition to improving the ring alignment, the adjustment procedure also afforded an opportunity to see the effect that future floor settling might have on the magnetic field.

The Plan

The height of the shimming cart as determined by the laser tracker clearly shows a modulation around the magnetic storage ring (see figure 5.10). The deviations of the ring plane were about 1 mm from the average value. The tilt plane was also apparent in the radial tilts of the pole pieces around the ring. An adjustment plan was crafted to correct the height and radial tilts of each yoke. The plan to level the ring used only adjustments on the support legs of each yoke. With yoke leg adjustments the interface between each yoke needs to remain fixed to prevent introducing additional stresses on the system. Therefore, the adjustment plan locked yokes together on the boundaries. The plan used linear fits of the laser in small angular windows on each side of the yoke boundaries to determine the target height for the inside of the yoke. The radial tilt measurements were used to calculate differential adjustments between the inner and outer legs of a yoke while of course minding the fixed interface.

Implementation

A team of mechanical engineers led by Eric Voirin performed the yoke adjustments. Adjustments were made incrementally in steps of approximately 100 µm to minimize differential stress between the yokes. During the process, the shimming laser tracker was replaced with the Fermilab alignment group's Hamar laser. Progress was measured periodically using a floating corner cube placed at three specific pole positions around the ring. The measure-

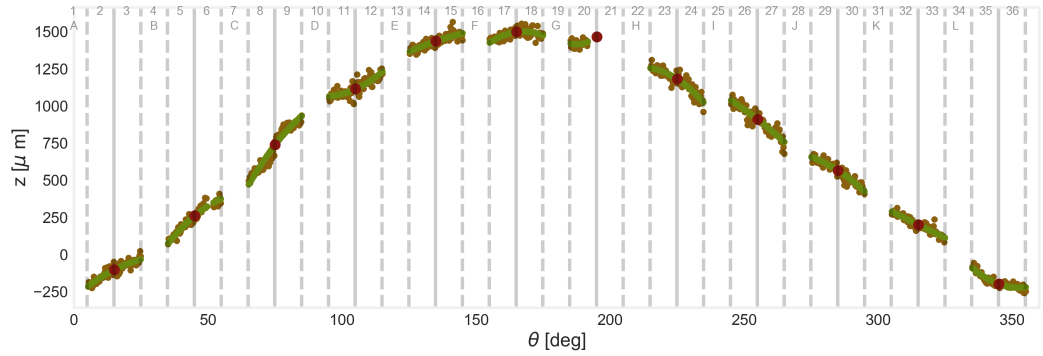

Figure 5.10: The plot depicts the height measurements of the shimming cart with small red dots. The large red dots represented the height of the pole surfaces at the interface between two yoke section. The edge around each interface was characterized with a linear fit shown in green. The storage ring plane defined by the pole surfaces had a rather large tilt plane into it. The tilt plane was removed using only adjustments on the yoke legs. The adjustments were devised to both bring the plane of the pole surfaces close to the plane defined by gravity and adjust the radial tilt of yokes large deviations.

ments with a sine fit overlay are depicted in figure 5.11. The progress plot highlights the well-controlled adjustment procedure. In total the whole ordeal took three days, two days for the initial plan and one more day of minor adjustments.

Results

The ring leveling was a success. Yoke leg adjustments managed to remove the tilt plane from the ring and help flatten the radial tilt of the pole (see figure 5.12). Afterwards the ring plane was flat to $\pm 150\,\mu\text{m}$. While flattening the ring simplified general $(g\text{–}2)_\mu$ detector alignment and made future adjustments simpler for the field team, one of the most important results was the effect on the field. In theory, the field should not depend on the ring plane or the relative yoke orientations, but the ring leveling provided an opportunity to symbiotically measure the effect. Figure 5.13 contains a plot of the difference between the azimuthally averaged field before and after leveling. There is a small change in the average dipole field which cannot be ascribed to a single specific cause with certainty (field drift, adjusted hysteresis from moving yokes, etc.). The higher order symmetries remain essentially unchanged though, and that is the more important result.

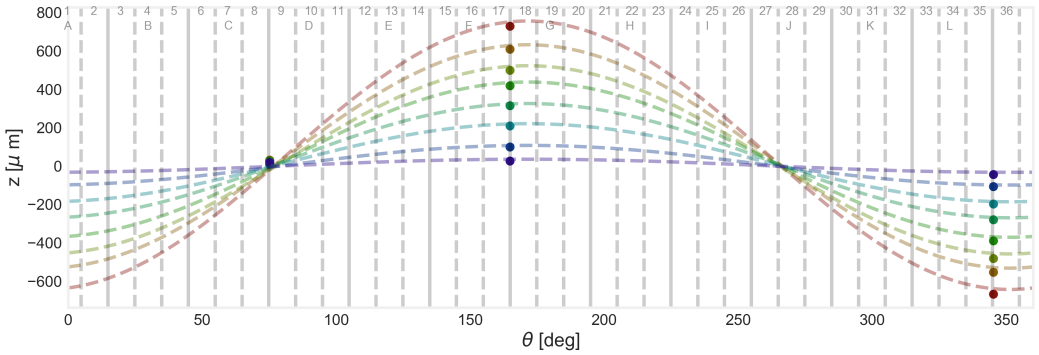

Figure 5.11: The image depicts iterations of yoke leg adjustments made to level the storage ring. Three locations around the ring were measured using the Hamar laser system. Those locations are represented by the points in the plot. The points were fit to sine wait to illustrate how the leveling effects in the unmeasured sections of the ring.

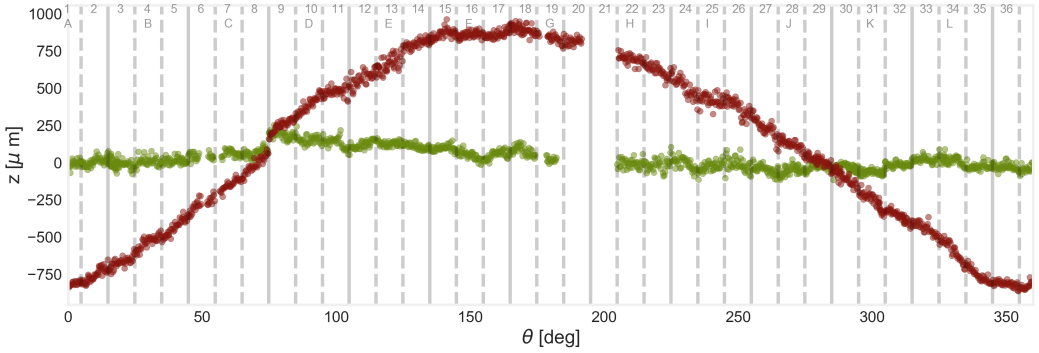

Figure 5.12: The image shows the original ring plane in red and the final ring plane in green. The plot shows the success of the procedure. The deviations in the ring plane were reduced from $\pm 1000\,\mu m$ to $\pm 150\,\mu m$.

Average Field Change From Ring Leveling

Orientation	Dipole	Quadrupole	Sextupole	Octupole
Normal	-9.97	0.01	0.03	-0.01
Skew	N/A	-0.09	0.08	-0.04

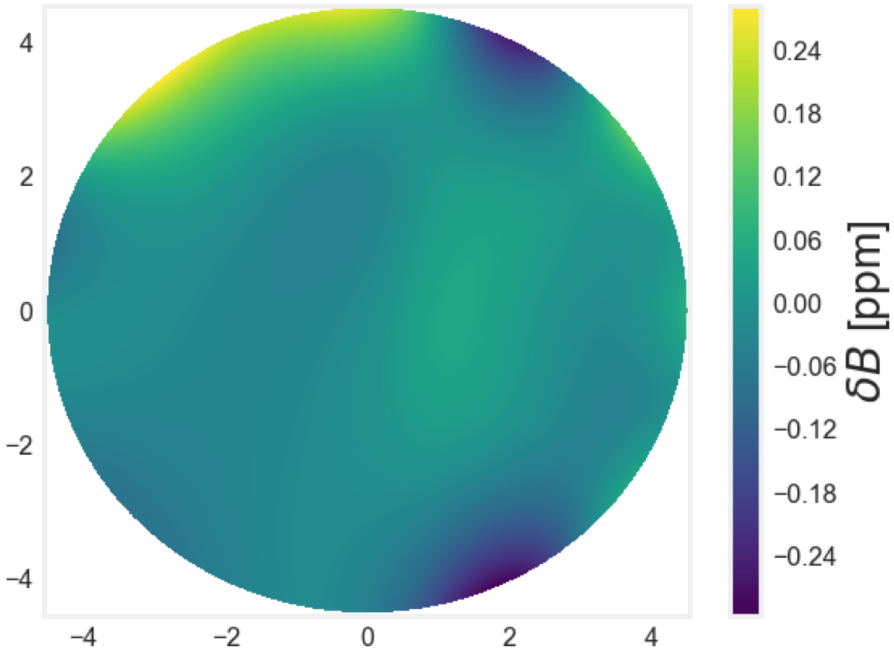

Figure 5.13: The figure highlights the results of an important test from the ring leveling process. The plots show the magnetic field averaged over the azimuth before and after ring leveling. The takeaway lesson is that there might be a small possible effect on the average field (though this could also be drift and other effects), and there is virtually no effect on higher order multiples. If necessary, the ring could be leveled again without significant changes to the field uniformity.

5.3.2 Pole Moves

Adjusting the orientation and shape of the pole pieces was the first substantial stage in optimizing the uniformity of g–2 storage field. Having previously established calibrations for the field effects from pole adjustments, the g–2 field team tested the model on a few of the most aberrant poles with some success. After a few tests, it became clear that adjustments could not be made on a pole-by-pole basis. The problem with a local approach was that it did not guarantee closure at the completion of a round of pole adjustments. The cost of adjusting a single pole was fairly large taking hours to crane out, replace shims in the pole feet, and crane back in, so the plan needed to achieve as much as possible with a single pass of pole adjustments. What was needed was a global model of the pole geometry.

The Bottom Pole Model

Building an accurate global model of the pole surfaces was not a simple task with the given instrumentation. Four measurements were available to use: the laser tracker, the tilt sensor, the Capacitec sensors, and the magnetic field measurements. The model ignored magnetic field data as it was difficult to decouple the effect of current shim positions, impurities, and perturbations from hardware that broke the azimuthal symmetry. In the right combination, the sensor data is just enough to build a complete model of the pole surfaces.

The bottom poles were characterized using a combination of tilt sensor data and laser measurements. The shimming cart rides directly on the bottom poles along the edge shims at the inner radius of the pole (see figure 5.7). The height value of the shimming cart then directly represents the inner radius of the bottom pole surface. From the inner radius of the bottom pole, the outer radius can be extrapolated with tilt sensor measurements. Each pole was broadly characterized with a radial and azimuthal tilt measurement taken from the center defining a rough plane for the pole. The pole interfaces were characterized with a set of three tilt measurements near the gap: one on the upstream pole, one straddling both poles, and one on the downstream pole. The pole step measurement is very sensitive to average elevation changes and rotation mismatches between poles. The pole model folds the laser height data in with a radial tilt model that smoothly varies across the pole, and

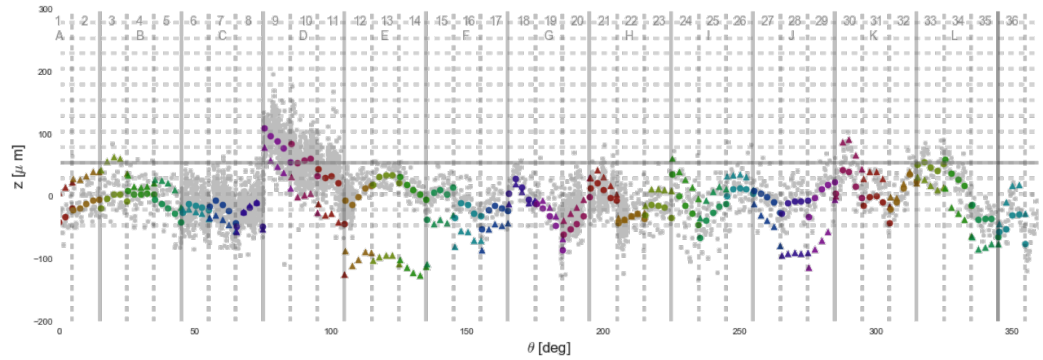

Figure 5.14: The plot visually represents the plan for adjusting pole feet using circles for the inner feet, triangles for the outer feet against their azimuth location. The laser data which anchored the pole surface model is also plotted in scatter form. Notice that the inner feet changes are minor adjustments on the laser data, and the outer feet changes fold in the radial tilt measurements on the poles.

then, restricts the pole interfaces with the pole step tilt measurements. An entire global plan for the bottom pole moves was made from this model, see figure 5.14

The Top Pole Model

The model for the top pole surfaces had to be constructed on the foundation of the bottom pole model. This situation was not ideal, since the top model compounds on errors in the bottom model. Starting from the model for the bottom poles, the sum of the two inner Capacitec measurements acted as a proxy for the gap between the upper and lower poles, so the sum of the bottom pole model inner radius and the two inner Capacitec values represented the inner band of the top pole model. With an established value for the inner band of the top poles, the same extrapolation to the outer band and pole interfaces was employed as with the bottom poles. The top and bottom outer bands could then be used to validate the pole model overall by comparing the predicted gap in the outer bands to the gap measured by the outer Capacitec sensors. While the model was not perfect, it was complete. The adjustment plan was ready for a full sweep of pole moves, top and bottom.

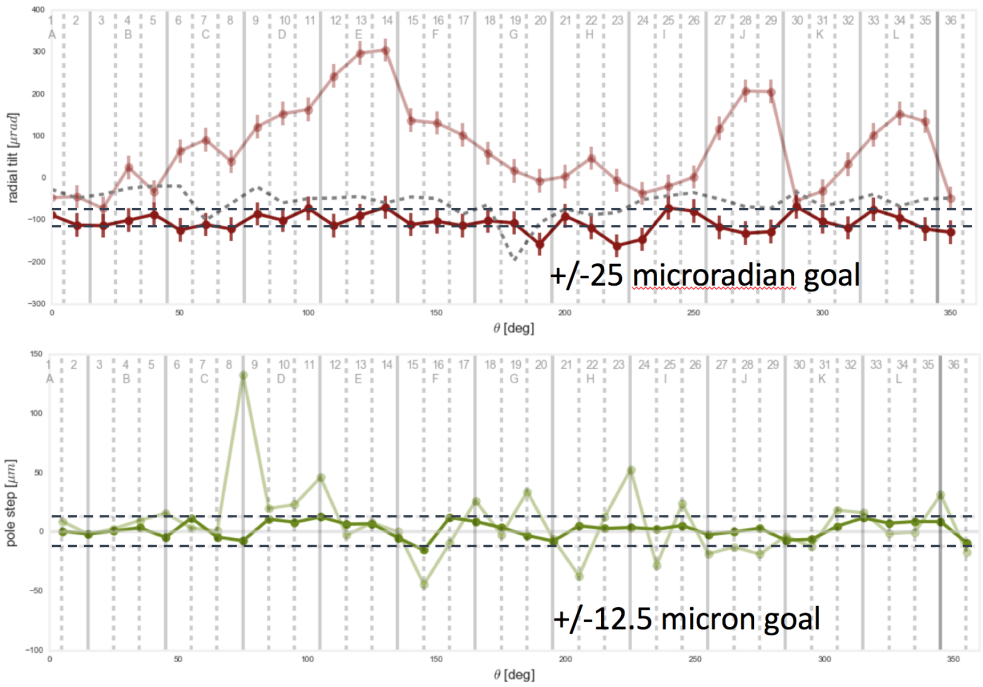

Figure 5.15: The radial tilts in red and the azimuthal tilts in green of the bottom poles are shown before pole moves in lighter color and after in darkened color. The dotted horizontal bands shown represent the target range for the tilt uniformity. The bottom pole movements brought most tilts into the target range.

Pole Moves Round One

Implementing the pole adjustments proved to be an involved procedure. A typical day went as follows: John Najdzion would come in early around 6 am and carefully transport 3 or 4 poles onto work blocks, the field team would come with micrometers and assorted shim stock to implement a prescription of changes to 1/4 mil (6 µm), and then John Najdzion would re-seat the poles into the ring. The first full round of pole movements took over a month to implement. The adjustments were able to successfully change the pole geometry though as shown in figure 5.15.

Pole Moves Round Two

The model was not perfect nor was the predictability of pole feet adjustments. The results from the first pass at pole adjustments were good, but a second round of smaller tweaks was needed to reach the target geometry. The primary goal of the second pass at pole adjustments was to eliminate large remaining height and tilt deviations and close remaining pole steps in the center of each pole interface to within 0.5 mil (12.5 µm). Many of the adjustments were able to be implemented without full removal of the poles, since the front pole feet were accessible after untorquing the pole's stabilizing bolts and propping the weight of the bottom poles up on a "jack stand". The "jack stand" was an adjustable pole foot prototype developed during the E821 era. The top poles did not require a jack stand, since gravity assisted in spacing the poles from the yoke in that case. In each specific case, the field team weighed the value of another attempted full pole adjustment against small tweaks on the inner feet.

Pole Move Results

While still making pole adjustments, the field team began to implement shim adjustments, so the entirety of the field improvements shown in figure 5.16 is not due to the pole moves. The main improvements from the pole moves are the elimination of sharp gradients in the average field that were caused by large elevation changes in the pole surface due to pole interface mismatches, along with the elimination of the average normal quadrupole moment by properly setting the radial tilts top and bottom.

5.3.3 Edge Shim Study

The edge shims are the most powerful tool available for adjusting multipoles higher than quadrupoles. However, the edge shims are not as trivial to adjust as other shims. The edge shims are originally ordered at a larger thickness than the expected necessary thickness. Then the effects of changing the thickness are calibrated. Magnetic field measurements can then be used to determine the optimal thicknesses of the edge shims. The edge shims are shipped off to be re-finished, and repeating as many times are needed. The process is costly

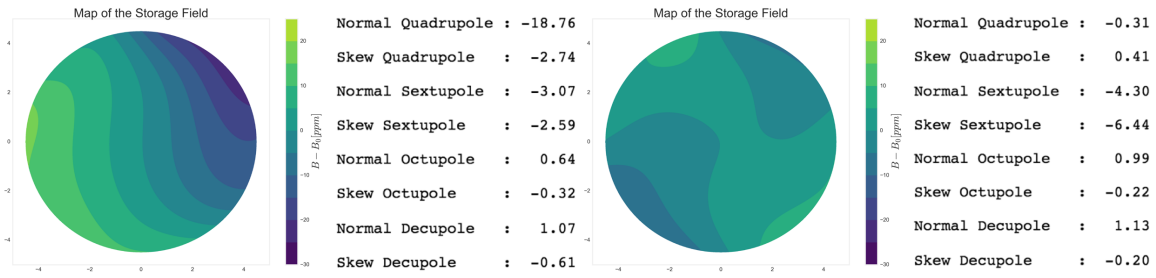

Figure 5.16: The azimuthally averaged field plots are shown with field prior to pole moves on the left and the field after pole moves on the right. A little care is necessary in interpreting attributions for the improvements though, because some shim adjustments were made before the pole movements had completed. The major improvement shown in the plot is the elimination of the -18 ppm normal quadrupole moment. A smaller improvement in the average skew quadrupole is also evident.

and time consuming, and therefore would only be performed if truly necessary.

Edge Shim Calibration

Several prototype steel shims were ordered for calibration and proof of principle testing of edge shim adjustments. One prototype edge shim of each type (inner lower, outer lower, inner upper, and outer upper) was installed in the place of a current edge shim at azimuthally isolated locations around the ring. The difference in thickness was measured and used to calibrate the multipole effects of the edge shims. The results are shown figure 5.17.

In addition to E821 style legwork, the E989 team set up tests using capton as a spacer under edge shims in lieu of grinding down shims and some tests using steel shim stock in lieu of implementing thicker edge shims. The capton spacer succeeded in reducing the influence of the edge shims, and the additional shim stock did indeed enhance the influence of the edge shims. The diminishing effects of the capton spacer were roughly one-tenth the size of the augmenting effects of the shim stock for the same size of spacer.

Edgelets

A new idea that arose during edge shim calibration was to implement edge shims in a more localized form. The angular size was chosen to be one degree segments and steel could only

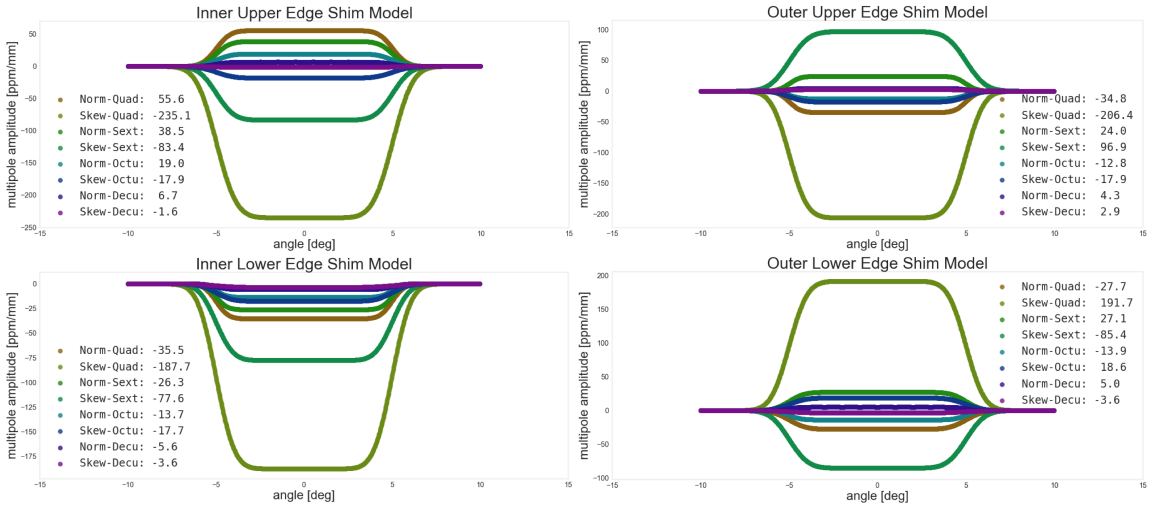

Figure 5.17: The models for edge shims derived from edge shim calibration data and edgelet shim tests. Each type of edge shim induced different higher order multipole effects. The effect shown is for a 1 mm change.

be added not removed, but the approach seemed promising. The model was based on the full edge shim model with edge effects being compartmentalized into one degree segments. The edgelet shim model went forward for a validation test. The test attempted to fix one of the most aberrant, localized regions of multipoles. The implementation did not achieve the success predicted by the model which signifies a need for more model complexity than originally expected. In the end the edgelet patching model was abandoned.

Implementation

The final footprint of the edge shims remained nearly identical to the initial footprint of the edge shims. Additionally a few steel shims of 25 µm were added in a few sections, but for the most part no adjustments to the edge shims were deemed necessary. One of the major drivers for acting with a light touch on edge shims was the promise of the optimization model using steel laminations.

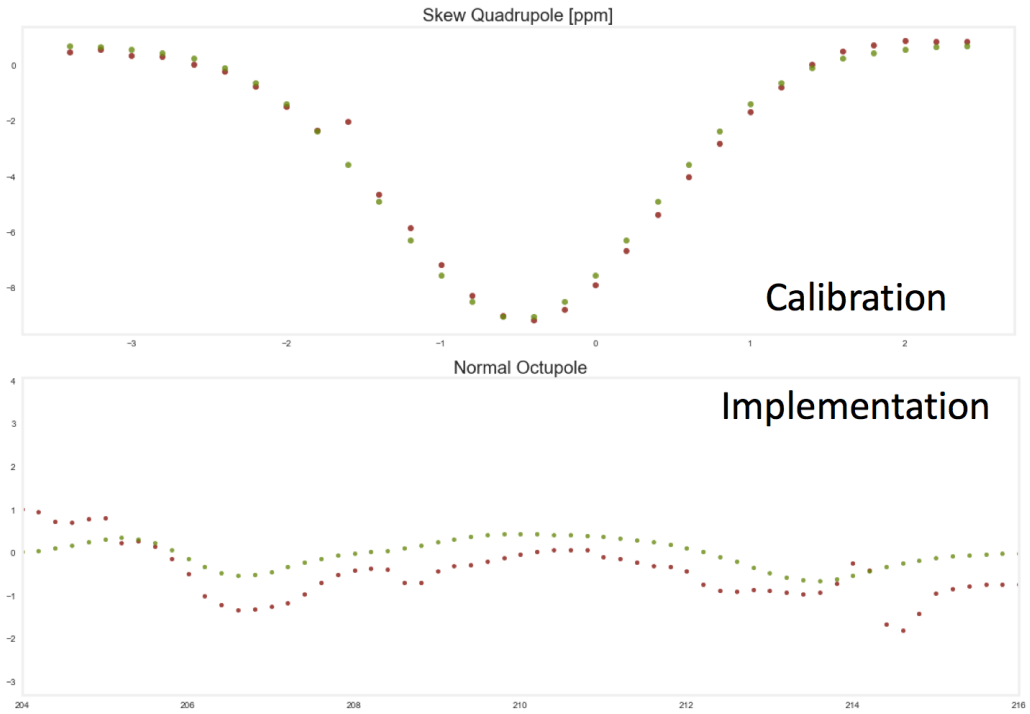

Figure 5.18: The top plot shows a calibration test for a one degree segment of edgelet shims. The measurement is in red, and the model is in green. The edgelet model agreed similarly well with other multipoles, a bit worse for some. The bottom plot shows predicted effects of the edgelet model in green and measured results of implementation across a full pole piece in red. The model was qualitatively correct, but was in need of more accurate calibration or increased complexity to be precise on the ppm level.

5.3.4 Shim Optimization

For the purpose of this section, shim optimization refers to optimizing the easily tunable shimming knobs, the top hats and the wedge shims. The top hats affect the average field and the wedge shims have an effect on the dipole and the quadrupole fields making these knobs a natural last stage in shimming the storage magnetic by design.

Calibration

The top hats were well understood fairly early. The data taken in figure 5.19 was actually taken using the Metrolab prior to full utilization of the shimming cart. The wedge shim calibration proved more difficult to pin down exactly (see figure 5.20). One of the confounding factors came from the constraint that the field must be powered down, the wedge shim adjusted and then the field re-powered to measure the effect. Field drift on the order of a few ppm was normal, but this made the calibrations looking for changes on the order of tens of ppm difficult, especially when the range was narrow such as single wedge calibrations. With these difficulties, it was necessary to calibrate, model, implement and validate the wedge shims for a few iterations before the model reached acceptable precision.

Field Model

The field optimization routine took the problem and cast it into a standard linear framework. If the problem is to be linear, then each knob needs to have a single dependent variable. Fortunately, the field shimming problem lends itself to the linear framework nicely if it is cast in terms of multipoles

$$\mathbf{y} = \mathbf{M}\mathbf{x} + \mathbf{b}$$

where $\mathbf{y}$ is the stack of residual field multipoles,

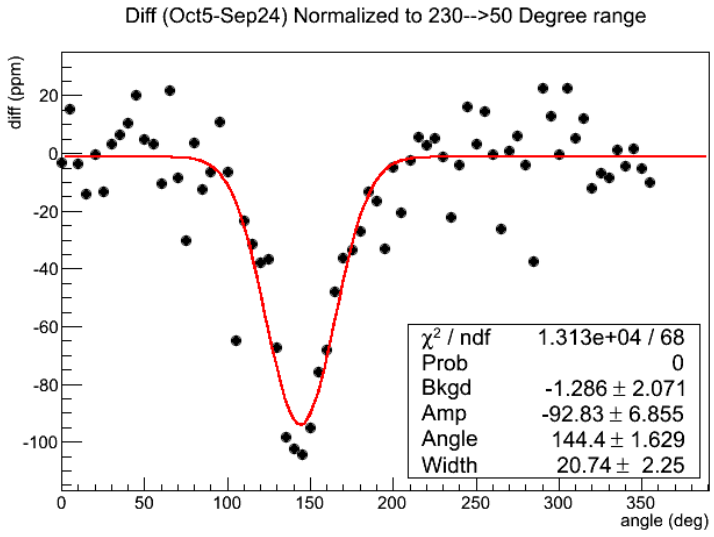

Figure 5.19: The calibration for the top hat shims. Between the two days that data was collected a 0.5 mm spacer was placed under one of the top hats. Normalized by size and adding a factor two for changes that would be done on both top and bottom hats, the calibration comes out to 338 ppm/cm with an RMS width of 20.74°.

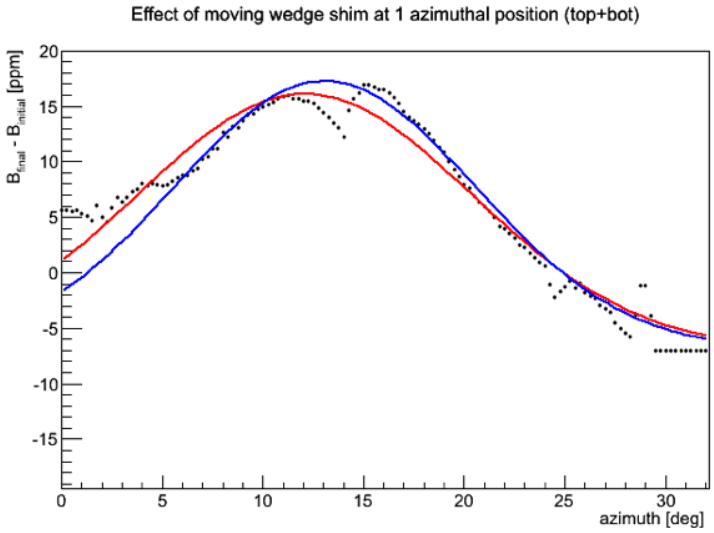

Figure 5.20: The plot shows an example wedge shim calibration. The size of the effect from the wedge shims was the same order as drift effects in the magnetic field. As a result, the wedge shims required multiple calibrations and validation steps to created a feasible model.

$$\mathbf{y} = \begin{bmatrix} \delta\vec{B}_{dipole} \\ \delta\vec{B}_{nquad} \\ \delta\vec{B}_{squad} \\ \delta\vec{B}_{nsext} \\ \delta\vec{B}_{ssext} \\ \vdots \end{bmatrix}$$

M contains the shim models,

$$\mathbf{M} = \begin{bmatrix} \vec{B}_{top-hat-01,dipole} & \cdots & \vec{B}_{wedge-01,dipole} & \cdots \\ \vec{B}_{top-hat-01,nquad} & \cdots & \vec{B}_{wedge-01,nquad} & \cdots \\ \vec{B}_{top-hat-01,squad} & \cdots & \vec{B}_{wedge-01,squad} & \cdots \\ \vec{B}_{top-hat-01,nsext} & \cdots & \vec{B}_{wedge-01,nsext} & \cdots \\ \vec{B}_{top-hat-01,ssext} & \cdots & \vec{B}_{wedge-01,ssext} & \ddots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix}$$

x is the dependent variable, shim position,

$$\mathbf{x} = \begin{bmatrix} x_{top-hat-01} \\ \vdots \\ x_{wedge-01} \\ \vdots \\ x_{edge-01} \\ \vdots \end{bmatrix}$$

and b is most recently measured field values

$$\mathbf{b} = \begin{bmatrix} \vec{B}_{dipole} \\ \vec{B}_{nquad} \\ \vec{B}_{squad} \\ \vec{B}_{nsext} \\ \vec{B}_{ssext} \\ \vdots \end{bmatrix}.$$

Each magnetic field mapping was interpolated into a grid of 1800 points and the first 9 multipoles were optimized. The dimension of y is then $16,200 \times 1$, and similarly b has a dimension of $1 \times 16,200$. The dimension of M is $16,200 \times 601$ from 432 wedge shims, 144 edge shims, 24 top hats, and 1 magnet current. The wedge shims and top hats were restricted to move symmetrically on top and bottom, so only half went into the optimization. The edge shim model had a separate entry for each type.

Solving the linear optimization problem is quite standard. The harder part is defining calibrated shim models and plugging into the linear format. Each shimming knob underwent calibration measurements, simulations, and had a precedent from E821.

Table 5.1: Shim Model Parameters

Multipole	Gain [ppm/mm]		Width [degrees]			
	Top Hat		Upper Wedge		Lower Wedge	
Dipole	-372.0	20.74	3.40	15.0	3.4	15.0
N-Quad	0.0	0.0	-0.18	4.0	-0.18	4.0
S-Quad	0.0	0.0	0.28	3.5	-0.28	3.5

The model was run through a linear least squares optimization scheme. Each knob was restricted to be within a finite range. Each multipole received a weight to normalize to the uniformity goals, for instance the usual weighting scheme placed a value of 1.0 on the dipole and 5.0 on the higher multipoles to reflect the relative goals of 50 ppm peak-to-peak on the dipole and 10 ppm peak-to-peak on the multipoles. The initial setting of each shim had to be known beforehand also, and these were measured by hand. The resulting model was useful in deciding the necessity of adjusting edge shims and iterating on the easily tunable shims to optimize the storage field (figure 5.21).

Results

The average field uniformity had gone from a standard deviation of 280 ppm to 40 ppm. Similarly most of the higher order multipoles had been shrunk to values below 1 ppm. At

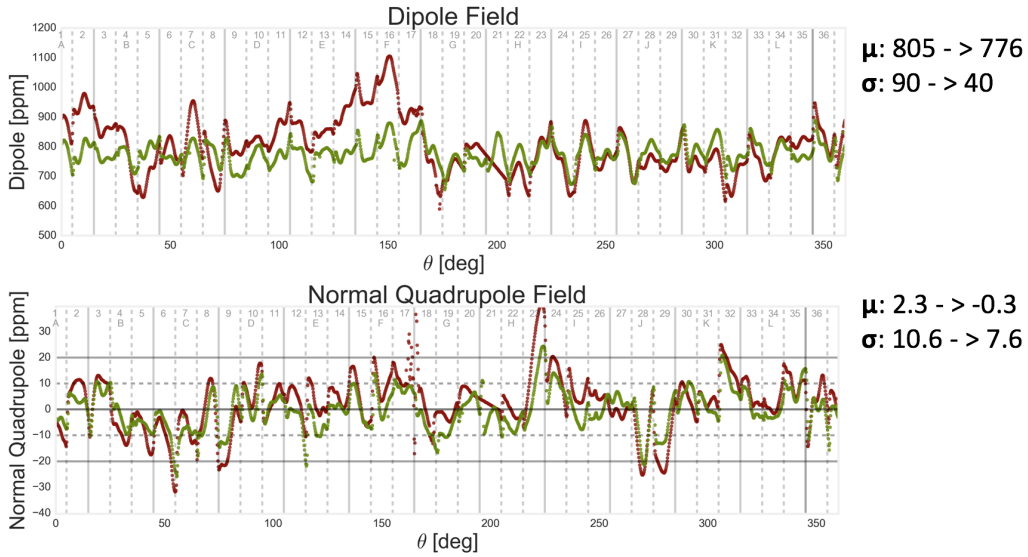

Figure 5.21: An example of one of the later rounds of shim adjustments. The wedge calibration had been tuned in properly, and as one can see, the optimization predicted strong improvements in the dipole normal quadrupole uniformity. Similar results were achieved with the standard deviation going from 90 ppm to 40 ppm and the normal quadrupole central value and uniformity improving to -0.3 ppm and 7.6 ppm respectively.

the end of optimizing the field using the built-in shimming knobs, the state of the magnetic field was similar to that of E821.

5.3.5 Laminations

The design flexibility of the storage magnet fell short of the uniformity goals for E989 due to limited angular resolution of the shims, so the field team had to expand the shimming infrastructure. The next stage of shimming placed thin, 25 µm to 50 µm, strips of iron onto the pole surfaces. In all over 8,000 strips of iron were placed to adjust the field uniformity to an unprecedented level.

Model

David Kawall implemented the model for optimization using iron laminations. The model treated the metal strips as thin, fully magnetized iron dipoles with three orders of magnetic

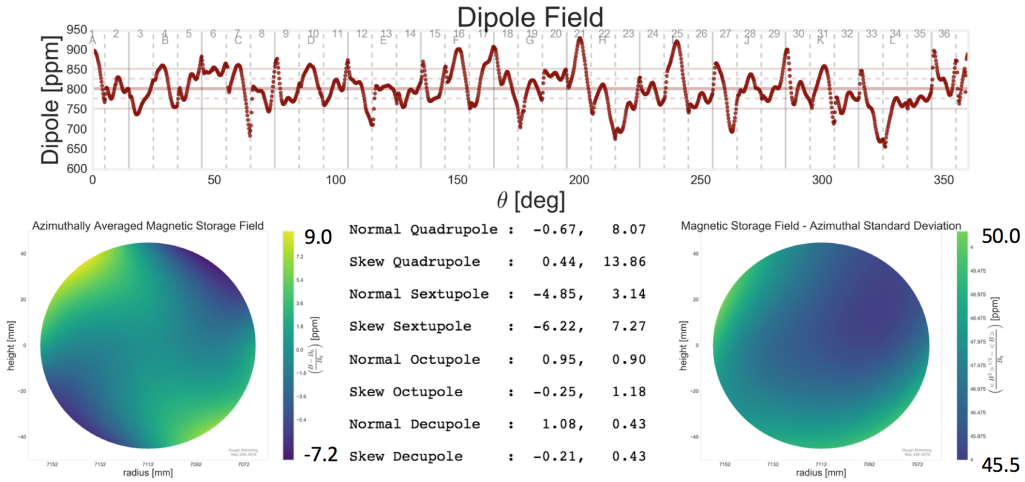

Figure 5.22: The final rough shimming dipole field result for E989 in red, compare with the PRD field standard deviation of 30 ppm. The horizontal bands indicated ±25 ppm around the central value which was target for E989. And, below the azimuthally average multipole content of the field. The results are on par with the running field from E821.

images manifesting inside the pole piece. The model split each pole into 41 azimuthal sections and 3 radial sections to give a lever on both higher field asymmetries and variations of the dipole field localized in azimuth.

Another concern arose about strong gradients from the lamination edges, since strong gradients could be a detriment to the fixed probe signals. The solution was to design an azimuthal grated pattern known as a "picket fence" region. Using azimuthal strips cut down on the size of gradients in the azimuthal direction which shorted the fixed probe FIDs. The final design is visible in figure 5.23. The picket fences successfully mitigated the detrimental effects of the iron foils as illustrated in figure 5.24.

Calibration

The lamination model included a few parameters that needed to be calibrated with the actual materials. The material purchased for laminations was Precision Brand 1010 shim stock with all material from the same lot to ensure similar chemical composition. Cutting a strip from the iron shim stock was the first step in calibrating the material. The strips were

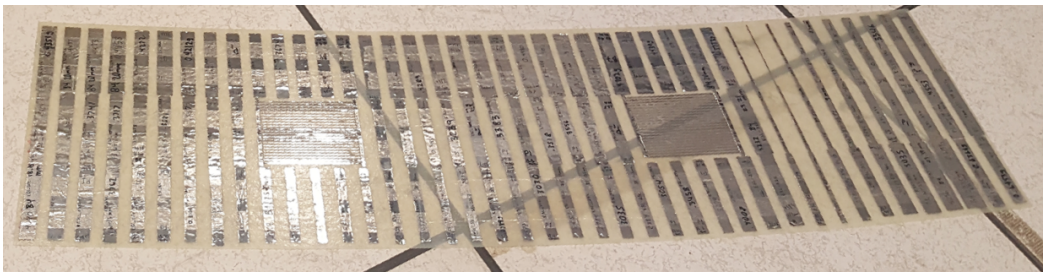

Figure 5.23: The image is an example of the lamination shims implemented as the final stage of rough shimming. Each G10 lamination shell has iron foils at 3 radii and 41 azimuthal positions. The exception is at the "picket fence" regions where the radial strips are replaced by longitudinal strips with a linear gradient. The picket fence was design to mitigate detrimental magnetic field gradients in regions where the fixed pNMR probes reside.

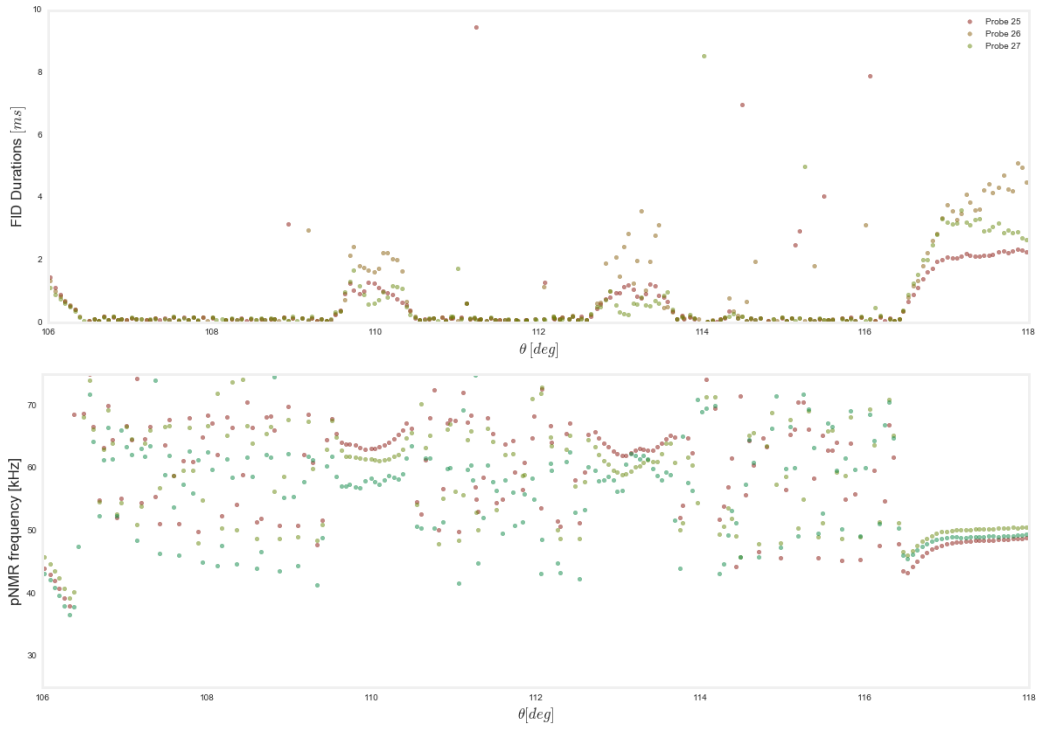

Figure 5.24: The setup for both plots shown above includes picket fence laminations deployed in two locations, 110° and 113°. The three pNMR probes depicted rode on the top of the shimming cart in a location similar to fixed probe locations within the vacuum chamber. The first plot shows how the FID duration increases from essentially no signal to a signal of $\sim$1 µs. The second plot illustrates how the frequency extraction from the fixed probe stabilizes over the picket fence regions.

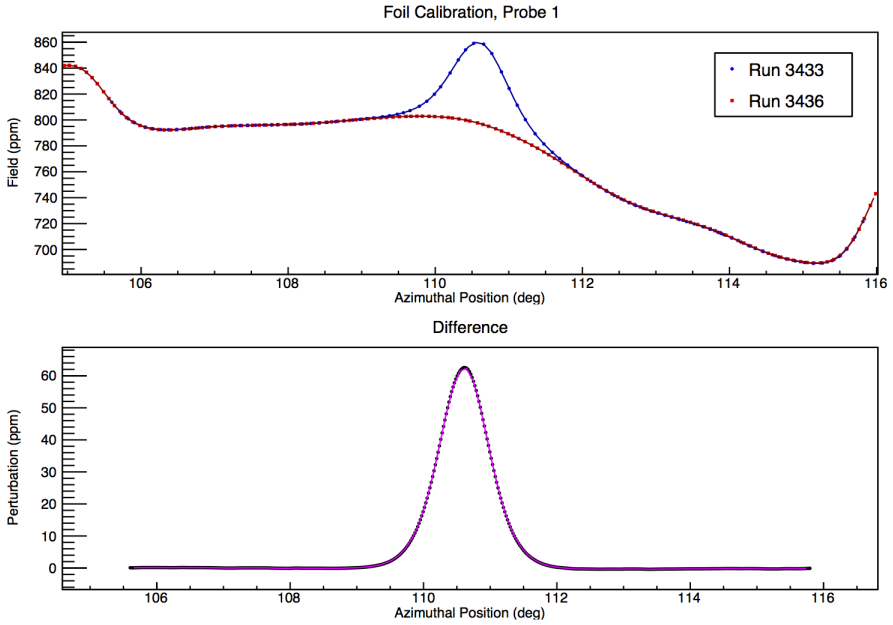

Figure 5.25: The calibration results with the raw measurements before adding the calibration strip in red and after in blue. The difference is compared to the model prediction in the lower plot which shows excellent agreement. All magnetic field values are field perturbation at the location of the central pNMR probe in the shimming cart.

typically the full width of the sheet, 12", and a few cm wide. The dimensions were all well measured and the sample was massed. Once the sample had been physically characterized, it was taped to a G10 sheet and affixed to a pole surface. After the sample was stabilized the shimming cart ran over the sample while measuring the field perturbations with its pNMR probes as shown in figure 5.25. The result was compared to predictions made by the foil model with adjustments allowed for the magnetic saturation and order of image dipoles to improve the model. The calibration process was performed for many batches of shim stock to ensure uniformity.

Implementation

The procedure between modeling the laminations, and fastening them onto the pole surface was rather involved. First, some of the high spikes in the field were lowered with the wedge shims by 50 ppm to 100 ppm, because the laminations could only add to the current field

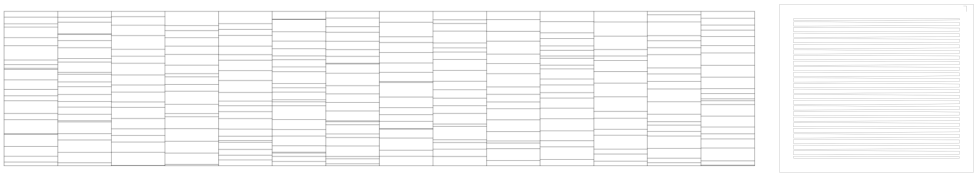

Figure 5.26: The left image depicts a typical distribution of foils cut by the laser cut to be sorted into mass bins. The right image shows a typical picket fence pattern sent to be laser cut and fit into the laminations.

value. The foils then allowed for another shot at improving the uniformity of those regions. Then, outputs from the lamination simulation were examined to create a target distribution of foils to fabricate with a laser cutter. The individual foils and picket fences needed to be cleaned, massed, sorted and drafted into the lamination plan. Each lamination board, made from G10 plastic, was cut, outfitted with the proper laminations and picket fences, and affixed to the pole surfaces using epoxy.

Foil production was a multi-stage process. The initial test batch of foils were cut by hand using a paper trimmer at FNAL. The process was slow and inaccurate in producing the desired foil masses. At the University of Washington, a laser cutting machine accelerated process by cutting hundreds of strips per hour. The patterns cut by the laser cutter were designed in Mathematica by Martin Fertl to produce a distribution that largely overlapped with the foil distribution requested by the lamination optimization. The laser cutter was used once again to cut the picket fence designs.

The laser cut foils were then shipped onward to the next stage of lamination production. At FNAL, the foils were cleaned in ethanol to remove plastic residue from the laser cutting process. During laser cutting the shim stock was placed on a substrate of plastic and thick paper of which the laser melted and sputtered a bit onto the metal. The cleaning was important, because the most important aspect of the strips was the total mass of the strip and residue would disturb the mass measurement. After cleaning, each strip was massed to the nearest 0.1 mg and sorted into a manual histogram made with a set of cups.

The next stage involved actually placing the foils onto G10 laminations shaped like pole pieces. The foils were inventoried using a shared spreadsheet, and another team of

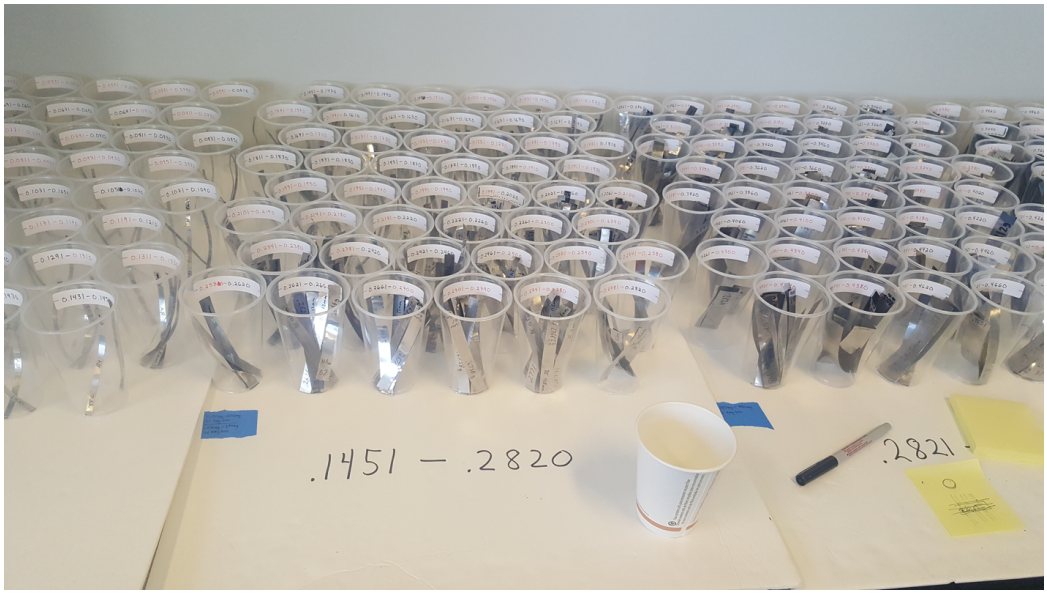

Figure 5.27: The image shows the manual histogram resulting from over 10,000 steel foils sorted by mass. The array of sorted foils helped to match the set of foil prescribed by the lamination optimization quickly.

workers tracked the inventory matching the prescriptions for each individual pole lamination. When a full prescription was available, the strips were affixed to the lamination using very strong and thin double sided tape (3M-9485PC). The finished pieces were then stored until completed tops and bottoms were paired and the relevant area of the ring was ready for installation.

The last stage of the lamination process was installation onto the pole surface. Lamination installation was done with the magnetic field powered down. Testing found that double sided tape was insufficient to affix the laminations stably onto the pole surfaces, so Stycast 1266 epoxy was employed instead. The epoxy was mixed in a controlled area, applied generously to the laminations. Then, the laminations were azimuthally aligned with the ring clocking. And finally a makeshift rig was assembled to keep both top and bottom laminations under even pressure while the epoxy cured. Figure 5.28 depicts the final stage in progress and post completion. The magnetic field was then ramped to verify the effects of the newly installed laminations.

Figure 5.28: The left image shows the full rig used to apply pressure and stabilize the laminations while the epoxy cured over 4 h to 6 h. The rig used a combination of plywood pole covers, curved plastic rails, small metal shims, and wooden wedges to distribute the pressure across all sections of the lamination. The right image shows the finalized laminations affixed to the pole surfaces.

The whole process was arduous and involved assistance from a large group of people throughout. 8 high school students, 2 undergrads, 4 graduate students, 2 post-docs, 1 high school teacher, 1 professor, 2 research scientists, 2 technical engineers, and one project manager lent their efforts to make the lamination procedure successful. The field team has lots of gratitude for all the honorary members that helped implement the laminations.

Effects

The results from the laminations were impressive. In figure 5.29, the improvement in the Yoke E region is clear compared to the other regions in the current iteration, and the magnetic field presented in the E821 PRD [3]. A similar improvement is seen in figure 5.30 which stacks the multipoles weighted by their E821 beam presence.

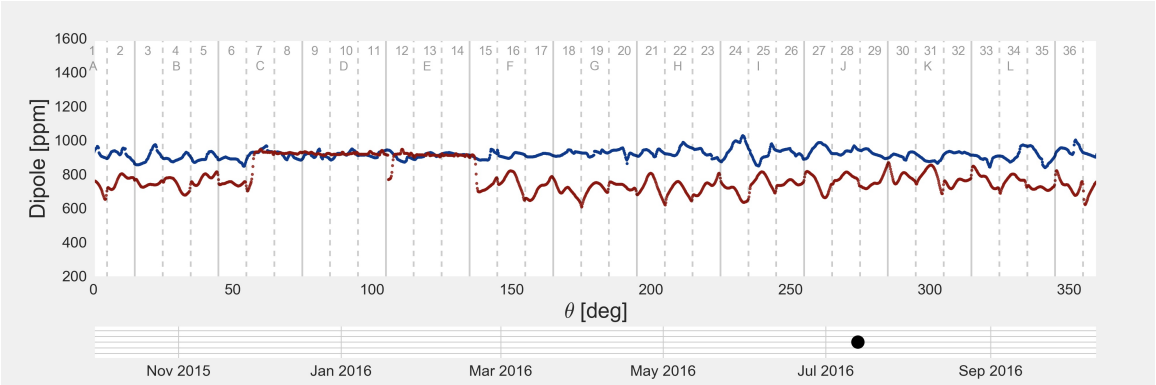

Figure 5.29: The E989 dipole field after installing laminations over part of the ring is shown in red, and the E821 PRD dipole field is shown in blue for reference. The laminations pushed the uniformity down to a ± 15 ppm band in the region of installation, well within the dipole magnetic field specification.

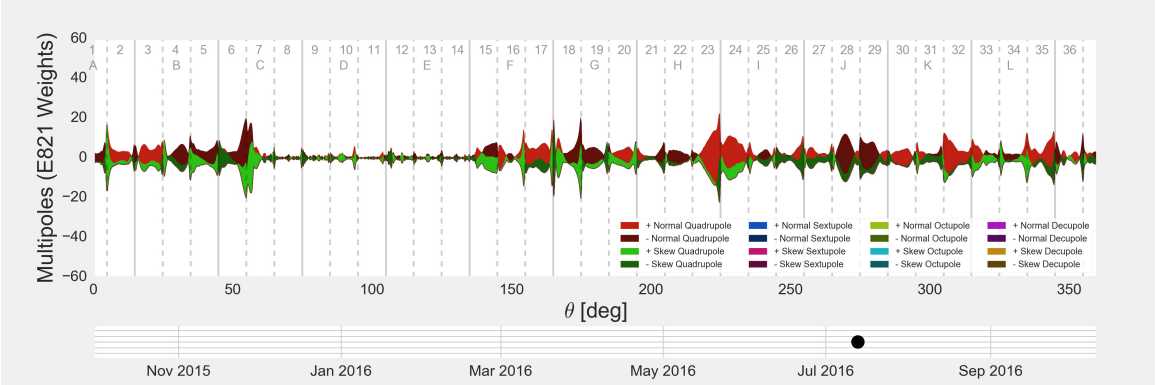

Figure 5.30: The E989 field multipoles after installing laminations over part of the ring are weighted by their E821 beam contributions and stacked as a proxy for the total uncertainty contribution from higher order multipoles in each azimuthal section of the ring.

5.4 Results

The field team started the shimming journey with the goal of a two-fold improvement over the E821. In the dipole field, that target is a peak-to-peak variation of 50 ppm. And in the field multipoles, that target is an azimuthal average under 1 ppm. Recall that these goals include active shimming with the surface coils, so the rough shimming has slightly more lax constraints on the azimuthal average.

5.4.1 Design Shimming

The first stage of progress comes from the set of built-in shimming knobs. The average field is shown in figure 5.31 as a comparison with the E821 PRD dipole field. The field at this point was not a large improvement over the efforts of the E821 team. The peak-to-peak variation after this stage is around 100 ppm.

5.4.2 Further Shimming

The implementation of the laminations pushed the peak-to-peak value of the dipole field to 30 ppm, well below the target range. The results are depicted in figure 5.33. The results for higher order field multipoles were similarly impressive. The azimuthal variation is largely suppressed. The plot in figure 5.34 shows each multipole weighted by the E821 beam composition for that particular multipole. A few of the multipoles have means that remain relatively large (a few ppm), but these do not appear in the figure. The reason for this is that those multipoles, mostly skew sextupole, were suppressed in the distribution of muons in E821 and are expected to be similarly suppressed in E989. The average value can still be adjusted by the active shimming mechanisms, the surface current coils. The surface current coil system was not commissioned at the time rough shimming ended.

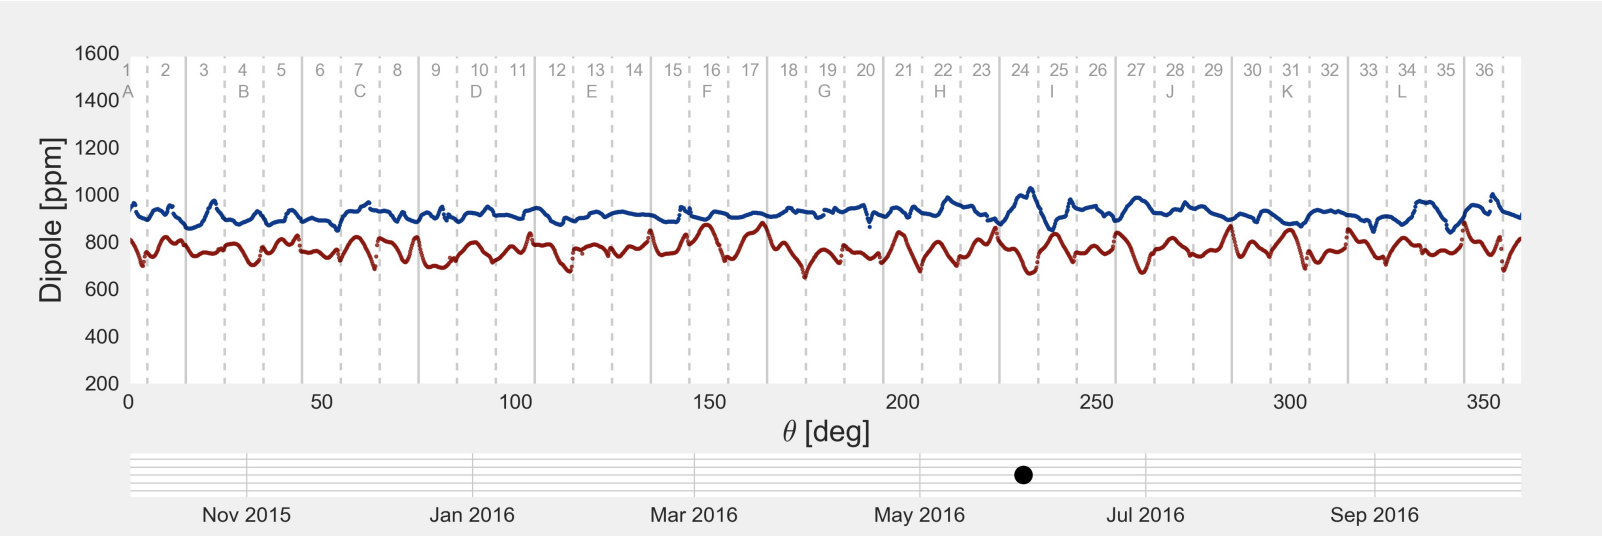

Figure 5.31: The plot depicts the rough shimming dipole field compared to E821 after full optimization using the built-in shimming kit. The E821 field has a better standard deviation, 30 ppm versus 40 ppm, and a similar peak-to-peak value for the average field.

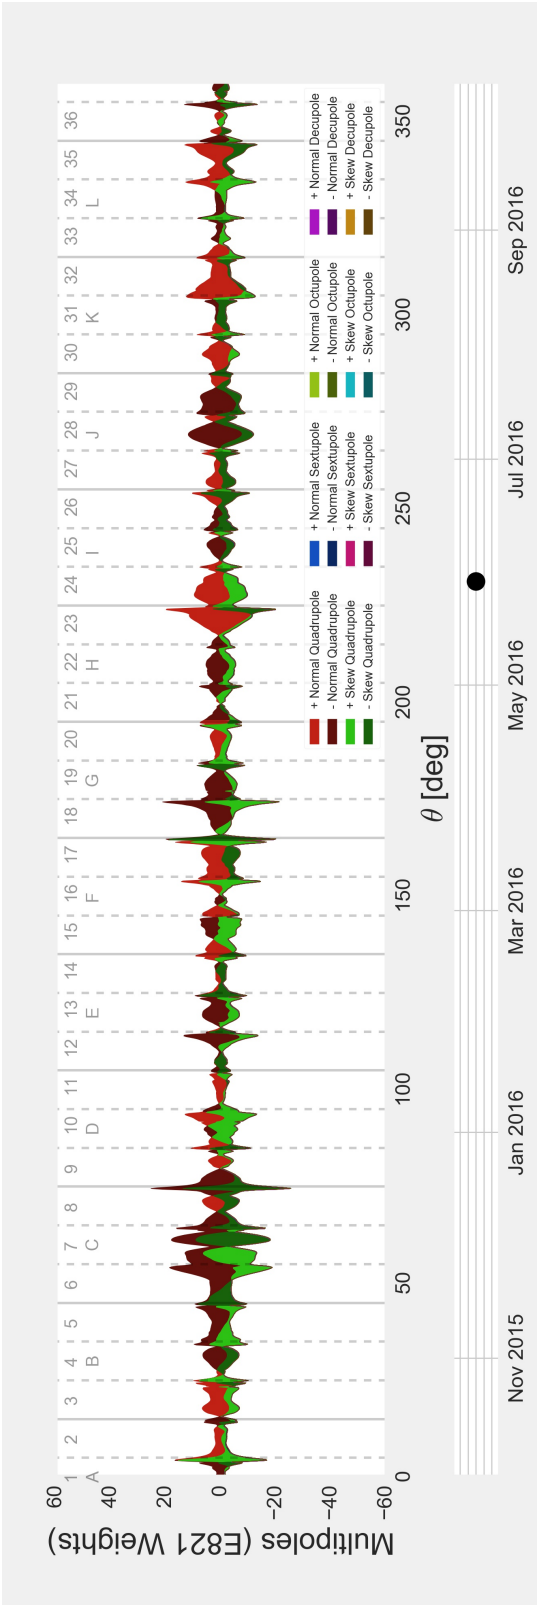

Figure 5.32: The plot shows the rough shimming multipoles after full optimization with the built-in shim kit. The dominant sources of uncertainty appear as the colors with the most area in the plot, i.e., the normal and skew quadrupole terms.

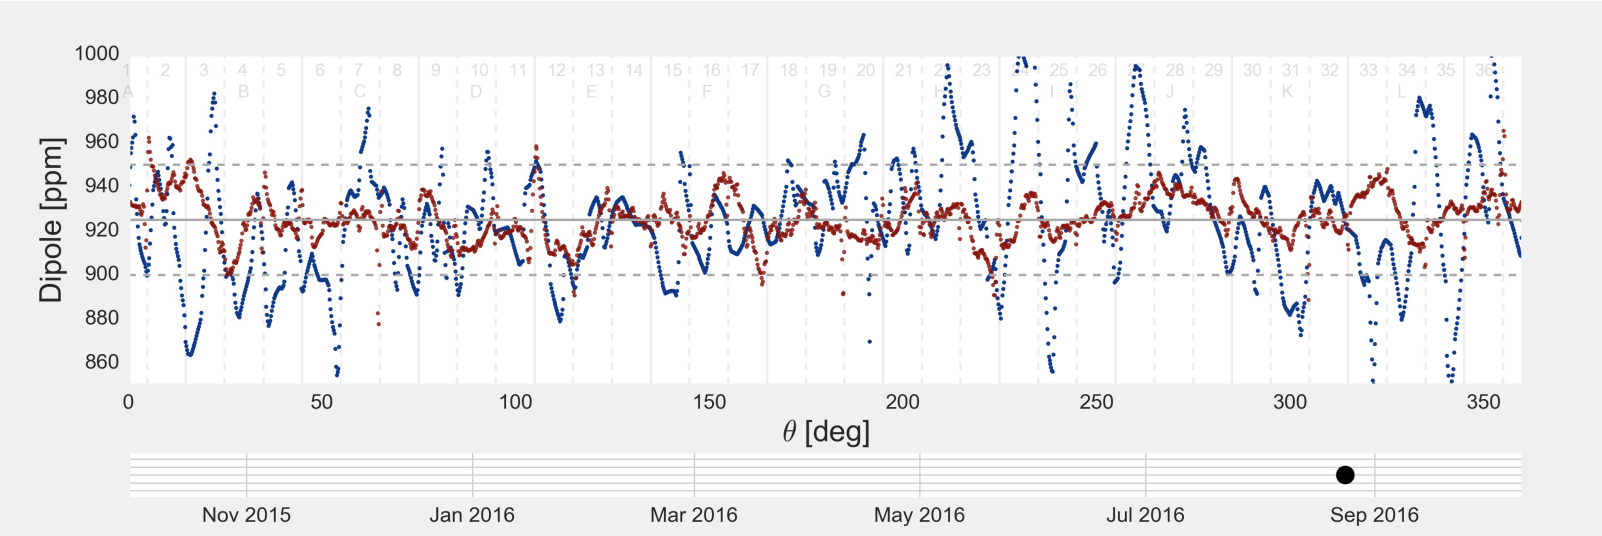

Figure 5.33: The final rough shimming dipole field for E989 in red compared with the PRD field plot for E821. The horizontal bands indicate ±25 ppm around the central value which was target for E989. The resulting dipole field beat the peak-to-peak target by almost a factor of two and E821 by a factor of three.

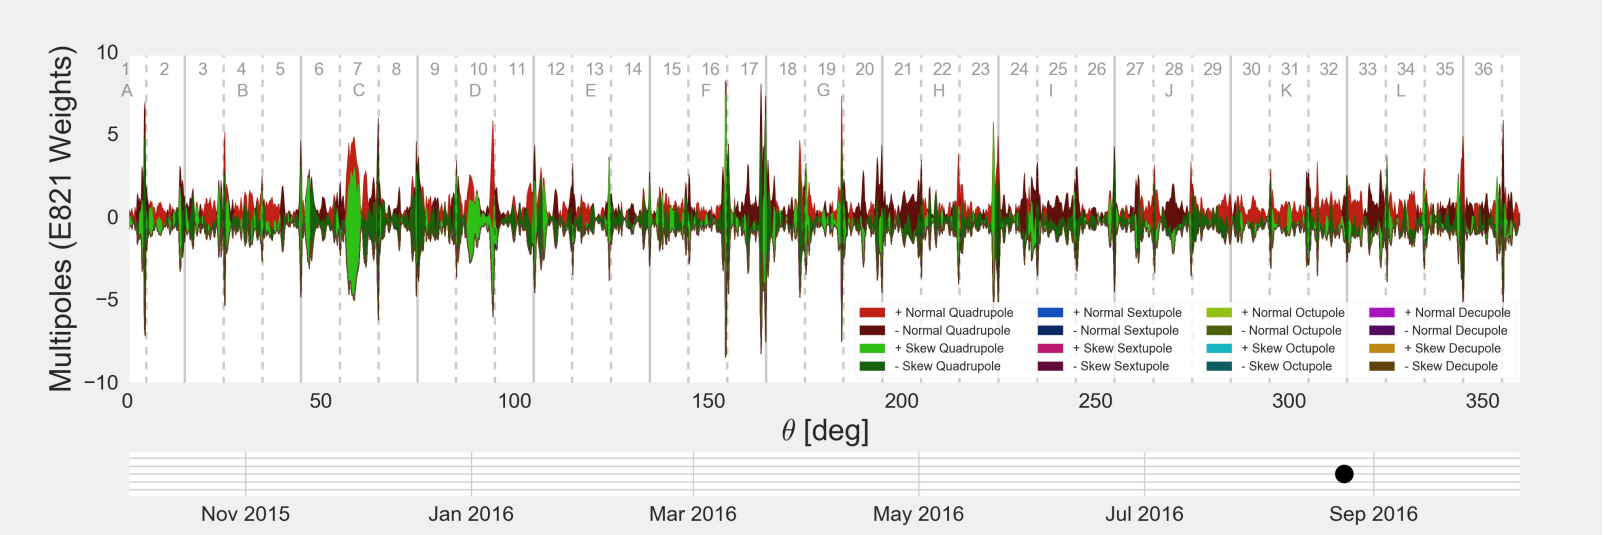

Figure 5.34: The plot shows each multipole weighted by the E821 beam composition for that particular multipole. A few of the multipoles have mean values that remain relatively large (a few ppm), but these do not appear in the figure. The reason for this is that those multipoles, mostly skew sextupole, were suppressed in the distribution of muons in E821 and are expected to be similarly suppressed in E989.

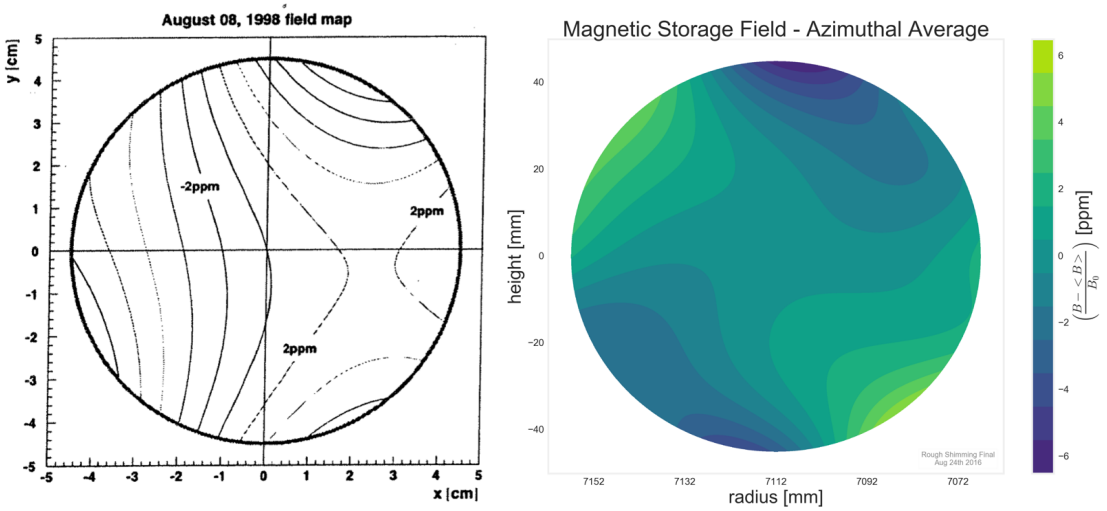

Figure 5.35: The left plot shows the azimuthally averaged magnetic field values after rough shimming was completed for E821 and the right plot shows the same for E989. The magnitude of the E989 field ranges over ± 6 ppm, but that is predominantly the skew sextupole term which the surface coils will nullify.

Chapter 6

CONCLUSIONS

The FNAL muon g–2 experiment continues to progress toward the 140 ppb measurement of the anomalous magnetic moment of the muon. Progress on the magnetic field measurement sub-domain in particular is detailed in this document. The analysis techniques for magnetic field measurements were extensively studied for precision levels and systematic effects. The magnetic storage field has undergone the passive shimming stage to achieve a new level of field uniformity for the g–2 storage ring.

6.1 FID Analysis

Magnetic field measurements in the $(g\text{–}2)_\mu$ experiment are subject to strong uncertainty goals. The technique employed to measure the magnetic field is a set of custom pulsed nuclear magnetic resonance (pNMR) magnetometers. Each measurement pulsed the pNMR probe to initiate a free induction decay (FID) signal where the frequency content represents the magnetic field. The soft uncertainty goal for each individual pNMR measurement is stated as 10 ppb. Within this domain, the measurement uncertainty is small enough that the statistics are negligible in the overall magnetic field uncertainty.

Several FID analysis methods were tested in detail. Three frequency domain and three time domain techniques were studied on FID datasets. Each technique was tested against idealized functional FIDs, simulated FIDs based on the Bloch equations, gradient FIDs based on superpositions of simulated FIDs, and real FID signals from manufactured pNMR probes in real magnetic fields.

The tests highlighted one method which showed the highest overall precision level. The most precise technique in the studies uses the complex signal information given by a Hilbert Transform in conjunction with the original signal. The phase information is extracted from the combinations of the signals, and the frequency can be found as the slope of a polynomial

fit. The conclusion of the study is to use the phase fit method for an optimal value and the zero crossing method as a standard value for comparison. The summary table is given first in table 4.3 of chapter 4, and included again here.

Table 6.1: FID Analysis Summary

Data Type	Zero Crossing	Phase Fit
Ideal	1.4 ppb	0.1 ppb
Simulated	4.6 ppb	0.1 ppb
Linear Gradient	7.2 ppb	0.6 ppb
Quadratic Gradient	7.6 ppb	1.9 ppb
Measured	9.3 ppb	8.9 ppb

6.2 Magnetic Field Shimming

Uniformity of the magnetic storage field is an undeniably important aspect of the FNAL $(g–2)_\mu$ experiment. The magnetic field uniformity contributes to the 30 ppb uncertainty allotted to magnetic field maps over the muon storage volume, the 30 ppb uncertainty placed on interpolating the magnetic field between full field mappings, and the 10 ppb uncertainty allowed for convolution of the muon distribution with the magnetic storage field. It is not the only contributor to these uncertainties, but maximizing uniformity lessens pressure for achieving the aforementioned goals.

The initial stage of magnetic field uniformity optimizations has been completed. The passive shimming stage took place over 11 months from September 2015 to August 2016. The process involved the careful adjustment of physics shims. Shim material is sometimes ferric to pull field with the shim and sometimes minimally magnetic to lessen flux capture. In all around 10,000 shims made of iron, steel, aluminum and G10 were placed with precision to adjust the physical geometry of the magnetic storage ring and the quality of the magnetic field.

The first uniformity goal was to reduce the peak-to-peak variation of the average mag-

netic field to ±25 ppm and a standard deviation of 25 ppm. The passive shimming stage achieved a peak-to-peak variation of ±25 ppm and a standard deviation of 15 ppm. The resulting plot of the magnetic dipole field is given first in chapter 5, and again here in figure 6.1.

The uniformity goal for the higher order multipoles (symmetries of the field) was an average value less than 5 ppm with a peak-to-peak value of less than 10 ppm. The multipoles couple to multipoles in the muon distribution which are not measured as of yet, but are most likely to manifest as normal multipoles. The passive shimming adjustments pushed the peak-to-peak of all normal multipoles within the target range, and the skew multipoles close to the target range. The average value for all multipoles is below the target of 5 ppm which puts them within range of the active shimming hardware. The plot of the azimuthal average was shown in Chapter 5 and reprises here in figure 6.2. The active shimming apparatus can control currents on the surface of the pole pieces to cancel small remaining multipole in the magnetic field.

The field uniformity still has a few possible routes for improvement. The active shimming stage has yet to make field adjustments. Additionally one set of shims are still accessible for minor adjustments even with $(g–2)_\mu$ detector systems hardware populating the inside of the storage ring. In total the passive shimming stage met and exceeded all uniformity goals.

6.3 Outlook

The FNAL $(g–2)_\mu$ experiment continues installation and commissioning of all subsystems. In the summer of 2017, muons reached the storage ring for the first time during a hardware commissioning run. The commission run will end midsummer, and muons will not return for several months. All subsystems will have a respite to fix unforeseen problems and improve robustness. The next time muons return to the $(g–2)_\mu$ ring, the data marks the start of making an unprecedented precision measurement of the anomalous magnetic moment of the muon.

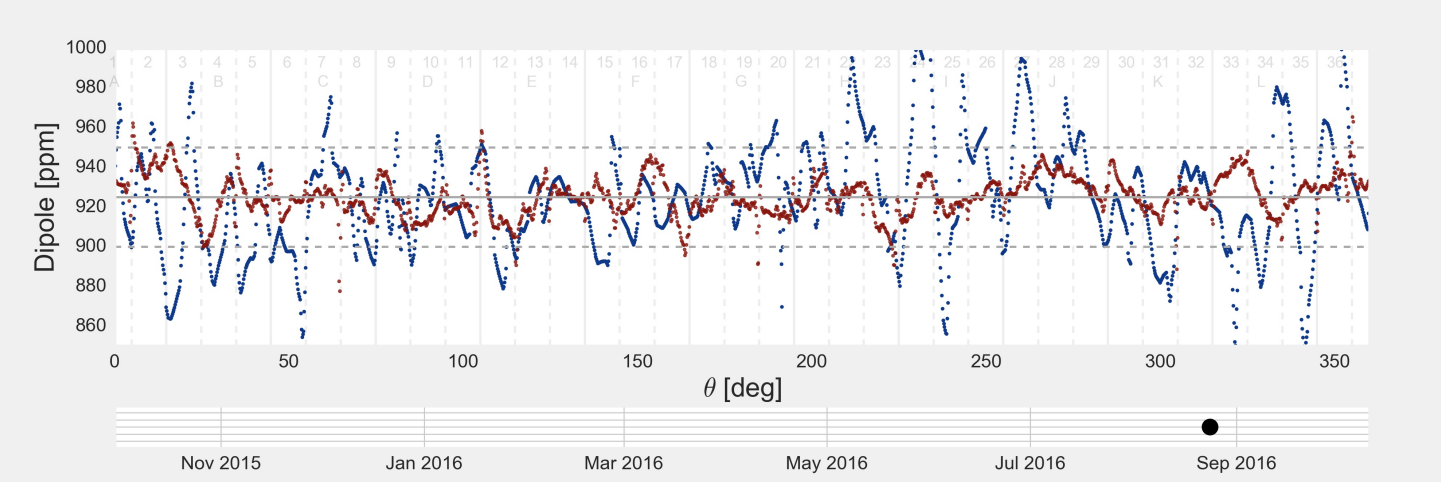

Figure 6.1: The final rough shimming result for E989 in red compared with the PRD field plot for E821. The horizontal bands indicate ±25 ppm around the central value which was target for E989. The result field beat the target by a factor of two and E821 field results by a factor of three.

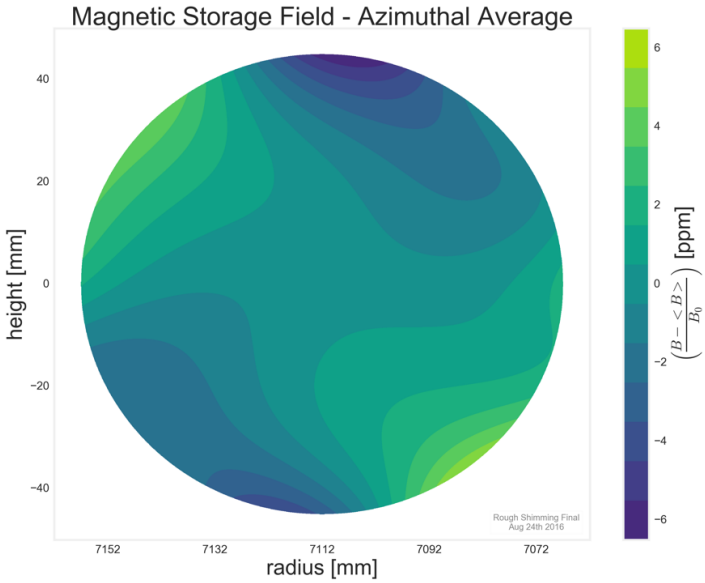

Figure 6.2: The left plot plot shows the azimuthally averaged magnetic field values after rough shimming was completed for E821 and the right plot shows the same for E989. The magnitude of the E989 field ranges over ± 6 ppm, but that is predominantly the skew sextupole term which the surface coils can eliminate.

BIBLIOGRAPHY

[1] *The Anomalous Magnetic Moment of the Muon.* Springer Berlin Heidelberg, 2007.

[2] Tatsumi Aoyama, Masashi Hayakawa, Toichiro Kinoshita, and Makiko Nio. Complete Tenth-Order QED Contribution to the Muon g-2. *Phys. Rev. Lett.*, 109:111808, 2012.

[3] G. W. Bennett et al. Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL. *Phys. Rev.*, D73:072003, 2006.

[4] Thomas Blum, Norman Christ, Masashi Hayakawa, Taku Izubuchi, Luchang Jin, Chulwoo Jung, and Christoph Lehner. Connected and Leading Disconnected Hadronic Light-by-Light Contribution to the Muon Anomalous Magnetic Moment with a Physical Pion Mass. *Phys. Rev. Lett.*, 118(2):022005, 2017.

[5] Thomas Blum, Norman Christ, Masashi Hayakawa, Taku Izubuchi, Luchang Jin, Chulwoo Jung, and Christoph Lehner. Using infinite volume, continuum QED and lattice QCD for the hadronic light-by-light contribution to the muon anomalous magnetic moment. 2017.

[6] Thomas Blum, Achim Denig, Ivan Logashenko, Eduardo de Rafael, B. Lee Roberts, Thomas Teubner, and Graziano Venanzoni. The Muon (g-2) Theory Value: Present and Future. 2013.

[7] G. Charpak, F. J. M. Farley, R. L. Garwin, T. Muller, J. C. Sens, V. L. Telegdi, and A. Zichichi. Measurement of the anomalous magnetic moment of the muon. *Phys. Rev. Lett.*, 6:128–132, Feb 1961.

[8] Andrzej Czarnecki and William J. Marciano. Muon anomalous magnetic moment: A harbinger for "new physics". *Phys. Rev. D*, 64:013014, Jun 2001.

[9] G. T. Danby et al. The Brookhaven muon storage ring magnet. *Nucl. Instrum. Meth.*, A457:151–174, 2001.

[10] F. J. M. Farley and Y. K. Semertzidis. The 47 years of muon g-2. *Prog. Part. Nucl. Phys.*, 52:1–83, 2004.

[11] Xiang Fei, V.W. Hughes, and Ralf Prigl. Precision measurement of the magnetic field in terms of the free-proton nmr frequency. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 394(3):349 – 356, 1997.

[12] Martin Fertl. private communication.

[13] David Flay. private communication.

[14] J. Grange et al. Muon (g-2) Technical Design Report. 2015.

[15] D. Hanneke, S. Fogwell Hoogerheide, and G. Gabrielse. Cavity control of a single-electron quantum cyclotron: Measuring the electron magnetic moment. *Phys. Rev. A*, 83:052122, May 2011.

[16] David Hertzog. private communication.

[17] John David Jackson. *Classical electrodynamics*. Wiley, New York, NY, 3rd ed. edition, 1999.

[18] Fred Jegerlehner. Muon g-2 Theory: the Hadronic Part. 2017.

[19] Fred Jegerlehner and Andreas Nyffeler. The Muon g-2. *Phys. Rept.*, 477:1–110, 2009.

[20] Klaus Jungmann. internal document.

[21] Dave Kawall. private communication.

[22] P. Kusch and H. M. Foley. The magnetic moment of the electron. *Phys. Rev.*, 74:250–263, Aug 1948.

[23] L.D. Landau and E.M. Lifshitz. *Statistical Physics*. Number v. 5. Elsevier Science, 2013.

[24] T. D. Lee and C. N. Yang. Question of parity conservation in weak interactions. *Phys. Rev.*, 104:254–258, Oct 1956.

[25] Weiwen Liu et al. High precision measurements of the ground state hyperfine structure interval of muonium and of the muon magnetic moment. *Phys. Rev. Lett.*, 82:711–714, 1999.

[26] W. H. Louisell, R. W. Pidd, and H. R. Crane. An experimental measurement of the gyromagnetic ratio of the free electron. *Phys. Rev.*, 94:7–16, Apr 1954.

[27] James P. Miller, Eduardo de Rafael, and B. Lee Roberts. Muon (g-2): Experiment and theory. *Rept. Prog. Phys.*, 70:795, 2007.

[28] Peter J. Mohr, David B. Newell, and Barry N. Taylor. Codata recommended values of the fundamental physical constants: 2014. *Rev. Mod. Phys.*, 88:035009, Sep 2016.

[29] Bill Morse. private communication.

[30] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.

[31] John F. Schenck. The role of magnetic susceptibility in magnetic resonance imaging: Mri magnetic compatibility of the first and second kinds. *Medical Physics*, 23(6):815–850, 1996.

[32] Julian Schwinger. On quantum-electrodynamics and the magnetic moment of the electron. *Phys. Rev.*, 73:416–417, Feb 1948.

[33] O. W. Srensen. James keeler. understanding nmr spectroscopy. *Magnetic Resonance in Chemistry*, 44(8):820–820, 2006.

[34] J.R. Taylor. *Introduction To Error Analysis: The Study of Uncertainties in Physical Measurements*. A series of books in physics. University Science Books, 1997.

[35] Alec Tewsley-Booth. private communication.

[36] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1415, Feb 1957.