Measuring the anomalous precession frequency ω_a for the Muon $g-2$ experiment

Jason Hempstead

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

Doctor of Philosophy

University of Washington

2021

Reading Committee:

David Hertzog, Chair

Alejandro Garcia

Shih-Chieh Hsu

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Measuring the anomalous precession frequency ω_a
for the Muon $g-2$ experiment

Jason Hempstead

Chair of the Supervisory Committee:
Professor David Hertzog
Physics

Announced in April 2021, the Run-1 result of the Fermilab Muon $g-2$ experiment increased the discrepancy between the measured value of the magnetic anomaly of the muon, a_μ, and the value predicted by Standard Model theory to over 4σ. A key component of the a_μ value is the measurement of the anomalous precession frequency ω_a, made possible by the parity-violating nature of the weak decay of muons. The energies of decay positrons carry information about the spin direction of the parent muons, so a suite of 24 calorimeters records the hit times and energies of decay positrons to fit for the oscillation ω_a. To perform a precise ω_a measurement, the nearly 1300 channels present in the calorimeter system must be aligned both in time and energy scale. In this thesis, the physical hardware of the calorimeter system and the data-processing software will be described; a blinded, preliminary analysis of Run-2 ω_a data will be shown, highlighting a possible improvement from the Run $-$ 1 analysis; and some systematic studies related to the detector hardware will be assessed in the context of Run $-$ 1 data.

TABLE OF CONTENTS

Page

LIST OF FIGURES

LIST OF TABLES

GLOSSARY

TDR: the Technical Design Report for the E989 experiment at Fermilab.

E989: the current generation of the Muon $g - 2$ experiment, located at Fermi National Accelerator Laboratory.

FERMILAB: Fermi National Accelerator Laboratory.

SIPM: Silicon PhotoMultiplier. A type of photodetector that is compact and robust in magnetic environments.

FFT: Fast Fourier Transform. A method of finding frequencies present in a collection of data.

CBO: Coherent Betatron Oscillation. The radial oscillation of the muon beam aliased at the cyclotron frequency.

SM: Standard Model. The laws governing particle physics.

ACKNOWLEDGMENTS

First I would like to thank my advisor David Hertzog, for helping me successfully finish this thesis, in a confusing time for both the world and myself. Our "physics chats," about $g-2$ and otherwise, were incredibly important for how I approach the field of precision physics.

My graduate school career would have quite different without the influence of Jarek Kaspar. For the experience and confidence with hardware systems he's helped me develop (and several rides to the hospital), I owe him many beers. I owe the same amount of gratitude to Aaron Fienberg for the help he has given me with software, data analysis, and anything else related to being a phenomenal physicist. His ability to make otherwise complex topics seem understandable will never fail to amaze me. To the rest of the UW $g-2$ group, our Fermilab complaints, grid struggles, silly whiteboard quotes, and quite essential pub lunches have made graduate school not so miserable.

Similarly, I need to thank the Muon $g-2$ collaboration for allowing me opportunities to grow and providing support along the way. Mark Lancaster and Brendan Casey made me feel like I had their respect as a peer despite their considerable seniority. The entire tracker/UK group felt like an adoptive family whenever I was at Fermilab.

Throughout my entire career at the University of Washington, Jayadev Athreya has been an incredible mentor. After our time working together at the University of Illinois, our serendipitously simultaneous move to Seattle could not have been more welcome. Thank you for continuing to push me, constantly reminding me of my sometimes lost love of the sciences.

There are far too many names to list every person who has helped me grow into the person and physicist I am, but I would like to thank my entire network of friends who supported me during this time. Some sent me physics papers to read, some went on bike rides with me, and some refused to laugh at my objectively funny jokes, but all of them kept me sane throughout the last six years.

CONTRIBUTIONS TO THE MUON $g-2$ COLLABORATION

The author began working with the Muon $g-2$ experiment during the summer of 2016. He took part in the final test beam run at SLAC National Accelerator Laboratory, where a near-production-level prototype of the Muon $g-2$ calorimeter (to be described in the text) was being tested. After the test beam, the author spent that summer working on quality control of components for the calorimeters, namely the SiPMs and PbF_2 crystals.

Following that, the author moved to Fermilab (a recurring theme) in order to help build, install, and commission the 24-station calorimeter system. See Figure 1 for a picture of the author building one of the calorimeters. One specifically daunting task was to de-magnetize the HDMI cables used to provide communication and voltage to the SiPMs. As a result, the author has stripped the shielding off of well over 1000 HDMI cables. The author was present in the counting room as the first muon beam was injected and stored in the Muon $g-2$ storage ring in the Spring of 2017. After the commissioning run was completed, the author led an effort to make the calorimeters more robust to electronics noise; this involved changing the signal cable connection for the 1296 SiPM channels. The other performed a number of other tune-ups to the calorimeters in this time, including remaking the cables that provide power for the SiPM amplifiers and communication system.

For Run $-$ 1, the author was the "Calorimeter coordinator" within the Muon $g-2$ collaboration, responsible for the continued health and performance of the calorimeter system hardware. Prior to the designated Run $-$ 1 data-taking period, it was discovered that the calorimeters had an unacceptably poor performance following beam injection. The author designed a test stand at the University of Washington to recreate the effect and test hardware improvements designed by the electrical engineers at the Center for Experimental Nuclear Physics and Astrophysics. The effort was massively successful, as will be shown in the text

Figure 1: The author assembling one of the Muon $g-2$ calorimeters. The active volume is at the right where a piece of black Tedlar rests on top of the array of PbF_2 crystals. The black HDMI cables provide communication and bias voltage to the SiPMs. The blue signal cables to the left have yet to be plugged into the SiPMs.

of the thesis. The proposed solution was adopted by the collaboration, and the author was responsible for installation in Winter 2018, ahead of $\text{Run}-1$.

The author once again moved to Fermilab full-time for $\text{Run}-2$, at which point he was promoted to the "Detector coordinator." In addition to his responsibilities with the calorimeter system, this involved ensuring maximal up-time of the straw-tracking detectors and various beam monitors. The first challenge came in the form of rebuilding a full calorimeter *in situ* after a water leak had rendered about half of the crystals unusable.

The calorimeter channels' gains had drifted rather far during the course of $\text{Run}-1$, due mostly to the temperature instability in the experimental hall. Before $\text{Run}-2$ began, the author was tasked with determining the voltage and programmable gain set-points for the 1296 SiPM channels. Instead of using the same values as used for $\text{Run}-1$, the author chose to make the calorimeters more robust to temperature variations. Following this, the author developed robust methods for aligning the energy scales and timing of the 1296 SiPM channels.

During $\text{Run}-3$, the author was selected as one of the "Run coordinators" — a sign of the trust garnered from the collaboration, as this position is typically reserved for postdoctoral researchers (and above). This position involves setting priorities for the day-to-day activities of the experiment and interfacing with the Fermilab accelerator division.

Almost continuously throughout this time, the author was intimately involved in upkeep and improvement of one of the main calorimeter reconstruction methods. The author helped with one of the main $\text{Run}-1$ analyses to determine ω_a, primarily by designing and carrying out systematic studies of the reconstruction method. The author helped push production of $\text{Run}-2$ data while testing alternative methods of handling the data and performing the ω_a analysis.

Chapter 1

MEASURING THE MAGNETIC ANOMALY a_μ

The Standard Model (SM) has provided an astounding number of verifiable (and verified) predictions about the behavior of subatomic particles. While highly successful, the SM does fail to explain some observations, such as the existence of dark matter. Highly precise tests of theory-based predictions are necessary to understand how the SM may be extended to fully explain the natural world.

One such precision measurement involves the magnetic moment of the muon. This introductory chapter will recount the recent announcement by the Muon $g-2$ collaboration showing how the predicted SM value holds up to the experimentally measured value. Following that, the techniques for both the SM calculation and the experimental measurement will be summarized.

1.1 Fermilab a_μ result

The Muon $g-2$ experiment at Fermilab (E989) released Run-1 results in April of 2021, much to the excitement of the physics community. This is the most precise single measurement of the anomalous magnetic moment of the muon a_μ to date. The result strengthened the long-standing discrepancy between the predicted and measured values of a_μ, as reported by the Muon $g-2$ experiment at Brookhaven National Lab (E821). Taken in isolation, the 0.46 parts-per-million (ppm) Fermilab result lies $3.3\,\sigma$ from the SM prediction [1]. When combined with the 0.54 ppm Brookhaven result [2], the experimental value is $4.2\,\sigma$ from the predicted value [1]. The comparison of these individual results and the average value can be seen in Figure 1.1.

Measurements of the muon's magnetic moment have been a test of the theoretical frame-

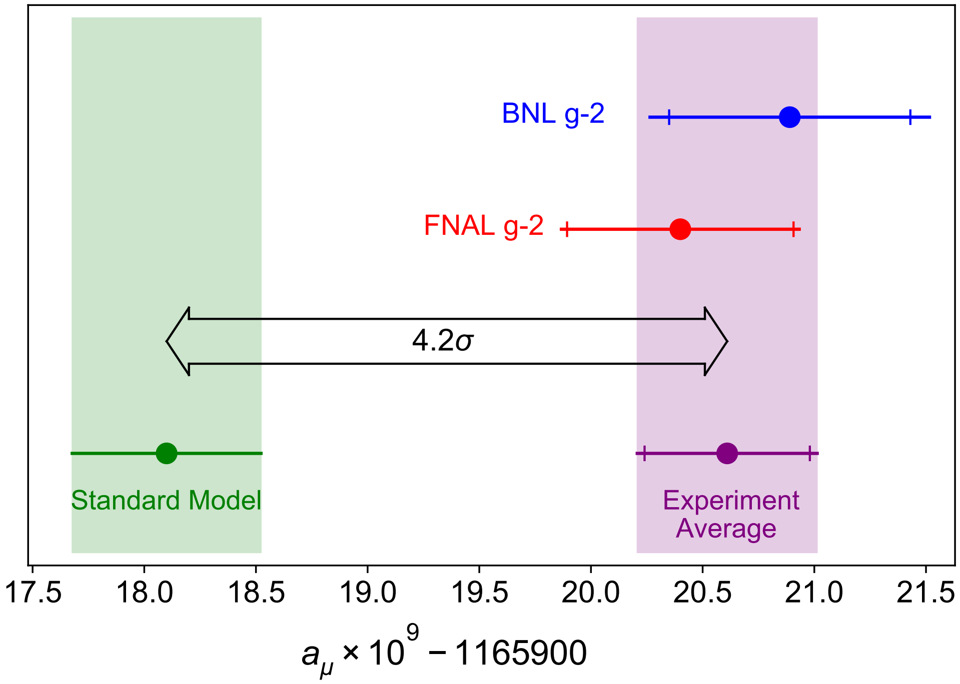

Figure 1.1: The two most recent measurements of a_μ, performed at Brookhaven and Fermilab. When the weighted average is taken, the experimentally measured value is 4.2σ higher than the value predicted by the Standard Model. Image reproduced from [1].

work of particle physics since the late 1950s, when experiments at Nevis Labs demonstrated that the weak decay of the muon violates parity [3]. Subsequent experiments at CERN and Brookhaven improved the accuracy of the measured value as efforts reduced uncertainty on the theoretical value at similar rates. The most up-to-date Standard Model prediction of a_μ^{SM} is the culmination of a years-long effort by the Muon $g-2$ Theory Initiative to reach a global consensus [4].

1.2 The magnetic moment of the muon

Classically, an electric charge with angular momentum has a magnetic moment $\vec{m}$ given by

$$\vec{m} = \frac{q}{2m}\vec{L}, \tag{1.1}$$

where q is the charge, m is the mass, and $\vec{L}$ is the total angular momentum. When in an external magnetic field, the (leading order) torque on the charge is

$$\vec{N} = \vec{m} \times \vec{B}_0, \tag{1.2}$$

where $\vec{B}_0$ is the first order term in an expansion of the magnetic field. The potential energy of the charge is given by

$$U = -\vec{m} \cdot \vec{B}. \tag{1.3}$$

Equation 1.1 does not adequately describe how particles with spin (intrinsic angular momentum) behave in magnetic fields. Rather, the magnetic moment $\vec{\mu}$ of such particles, like the muon and electron, is given by

$$\vec{\mu} = \frac{ge}{2mc}\vec{s}, \tag{1.4}$$

where e is the elementary charge, m is the particle mass, c is the speed of light, and $\vec{s}$ is the spin vector. For spin-1/2 particles, $\vec{s}$ can assume quantized values of $\pm\frac{\hbar}{2}$ on any axis. The parameter g is a dimensionless quantity (called the g-factor) with a value $g \approx 2$ for muons and electrons [5].

When moving at velocity $\vec{v}$ relative to external magnetic and electric fields $\vec{B}$ and $\vec{E}$, the spin will rotate at a rate given by[1]

$$\left(\frac{d\vec{s}}{dt}\right)_{\text{rest frame}} = \vec{\mu} \times \left(\vec{B} - \frac{\vec{v}}{c} \times \vec{E}\right) \tag{1.5}$$

---

[1] This form of the spin rotation ignores a term arising from a non-zero muon electric dipole moment (EDM). The muon EDM has been shown to be small [6], so this is a safe assumption.

in the muon's rest frame, where $\vec{v}$ is the velocity of the muon relative to the external fields and c is the speed of light. Thomas introduced the translation to the non-rotating lab frame through

$$\left(\frac{d\vec{s}}{dt}\right)_{\text{lab frame}} = \left(\frac{d\vec{s}}{dt}\right)_{\text{rest frame}} + \vec{\omega}_T \times \vec{s}. \tag{1.6}$$

The Thomas precession term $\vec{\omega}_T$ is given by

$$\vec{\omega}_T = \frac{1}{c^2}\left(\frac{\gamma^2}{1+\gamma}\right)\frac{d\vec{v}}{dt} \times \vec{v}, \tag{1.7}$$

where γ is the Lorentz boost factor

$$\gamma = \sqrt{\frac{1}{1-\frac{v^2}{c^2}}}.$$

For relativistic motion in the external fields, with $\vec{\beta} = \frac{\vec{v}}{c}$:

$$\frac{d\vec{v}}{dt} = c\frac{d\vec{\beta}}{dt} = \frac{e}{\gamma m}\left(\vec{E} + \vec{\beta} \times \vec{B} - \left(\vec{\beta} \cdot \vec{E}\right)\vec{\beta}\right) \tag{1.8}$$

Substituting this into Equation 1.6 leads to the familiar form of the expression for the rate of change of the spin direction [7]:

$$\frac{d\vec{s}}{dt} = \frac{e}{mc}\vec{s} \times \left[\left(\frac{g-2}{2} + \frac{1}{\gamma}\right)\vec{B} - \frac{g-2}{2}\frac{\gamma}{\gamma+1}\left(\vec{\beta} \cdot \vec{B}\right)\vec{\beta} - \left(\frac{g-2}{2} + \frac{1}{\gamma+1}\right)\vec{\beta} \times \vec{E}\right]. \tag{1.9}$$

This has been indicatively written in terms of $\frac{g-2}{2}$, a quantity that will be defined in the next section.

1.3 Calculating a_μ in the Standard Model

While the formalism above directs how an experiment would be performed, it does not guide how one would predict a value of a_μ in the Standard Model. Following the description in Reference [8], the Dirac equation for a spinor field ψ in an external field A reads

$$\left(i\not{\partial} - e\not{A} - m\right)\psi = 0, \tag{1.10}$$

leading to

$$\left[(i\partial_\mu - eA_\mu)^2 - \frac{e}{2}F_{\mu\nu}\sigma^{\mu\nu} - m^2\right]\psi = 0. \tag{1.11}$$

The $F_{\mu\nu}\sigma^{\mu\nu}$ term corresponds to the magnetic dipole interaction with the external field; $F_{\mu\nu}$ is the electromagnetic field strength tensor, and the $\sigma^{\mu\nu}$ are the relativistic spin matrices. In the non-relativistic limit, this produces the Hamiltonian

$$H = \frac{\vec{p}^2}{2m} + V(r) + \frac{e}{2m}\vec{B}\cdot\left(\vec{L} + 2\vec{S}\right). \tag{1.12}$$

Note that there is an inherent assumption that $g = 2$ in the equation above.

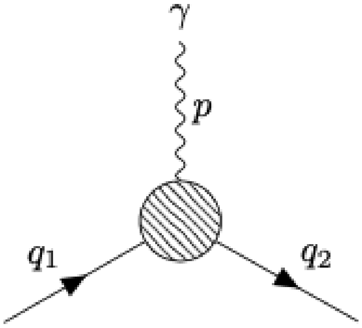

Figure 1.2: A generic interaction between a fermion and an external electromagnetic field. Within the blob could be any number of particles and interactions. Equation 1.13 allows parametrization of whatever processes take place. Image reproduced from [8].

The matrix element for a generic interaction between a muon and an external field, as seen in Figure 1.2, can be parametrized as:

$$i\mathcal{M} = (-ie)\bar{u}\left(q_2\right)\left[F_1\left(\frac{p^2}{m^2}\right)\gamma^\mu + \frac{i\sigma^{\mu\nu}}{2m}p_\nu F_2\left(\frac{p^2}{m^2}\right)\right]u\left(q_1\right), \tag{1.13}$$

where F_1 and F_2 denote form factors that describe the interaction, $\bar{u}(q_1)$ and $u(q_2)$ are Dirac spinors, $p^\mu = q_2^\mu - q_1^\mu$ is the momentum exchange with the external field, m is the mass of the lepton, and γ^μ are the Dirac gamma matrices. The leading order diagram seen in Figure 1.3a has the matrix element

$$i\mathcal{M}_0 = (-ie)\bar{u}\left(q_2\right)\gamma^\mu u\left(q_1\right), \tag{1.14}$$

meaning $F_1 = 1$ and $F_2 = 0$ at tree level. Identifying that F_2 modifies the magnetic moment and measurements will be made in the energy region $p << m_\mu$, the g-factor of the muon is [8]

$$g = 2 + 2F_2(0). \tag{1.15}$$

Using the definition of the magnetic anomaly a_μ:

$$a_\mu \equiv \frac{g-2}{2} = F_2(0). \tag{1.16}$$

Thus, by calculating the F_2 form factor in any generic interaction, the contribution to the magnetic moment can be found.

The first correction to $g \approx 2$ arises at second-order from a photon loop in the muon vertex, as seen in Figure 1.3. This is the famous Schwinger term, resulting in a contribution of $\alpha/2\pi$. The calculation of the Schwinger term is rather straightforward; an accessible walkthrough can be found in [9].

(a) First order diagram

(b) Schwinger term

Figure 1.3: The lowest-order contributions to the muon's magnetic moment. On the left is the tree-level interaction with a photon. The right shows a photon exchange within the vertex, contributing $\alpha/2\pi$ to a_μ.

Additional terms come from loops such as those seen in Figure 1.4. The contributions are split into four categories: quantum electrodynamics (QED: photon and lepton loops), electroweak (W/Z-boson and Higgs loops), hadronic vacuum polarization (HVP), and hadronic light-by-light (HLbL). The contributions to a_μ^{SM} and associated errors for each of the terms are listed in Table 1.1. Note that while it is the largest contribution, the error on the QED contribution is miniscule. The upcoming sections will walk through each contribution and give a brief overview of how each is calculated. For a more comprehensive review of the current calculation of a_μ^{SM} and the individual contributions, the reader is directed to [4].

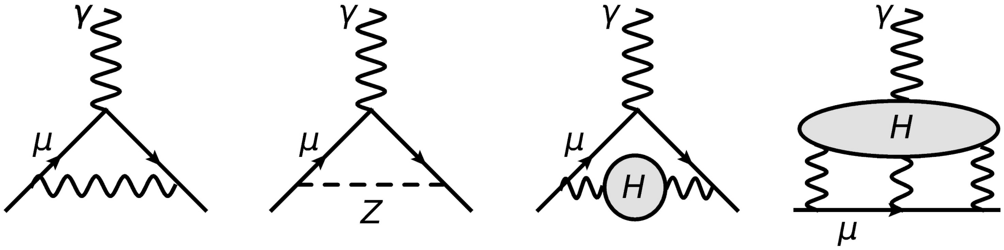

Figure 1.4: Examples of loops that enter into the calculation of a_μ^{SM} . From left to right are a QED term, an electroweak term, a hadronic vacuum polarization (HVP) term, and a hadronic light-by-light (HLbL) term. Image reproduced from [1].

1.3.1 Quantum electrodynamics (QED) and electroweak loops

The most well-understood contributions to a_μ^{SM} arise from photon and lepton loops in the muon-photon vertex. The calculation is split into lepton mass-dependent and -independent terms. For the lepton mass-independent terms, each additional photon loop carries an additional order of the fine structure constant $\frac{\alpha}{\pi} \approx \frac{1}{430}$, meaning a reduction in several orders of magnitude. Up to four loops, the mass-independent term has been calculated near-analytically, involving the evaluation of 891 Feynman diagrams [36]. This calculation agrees with fully numerical calculations of the same term. The lepton mass-independent contributions have been calculated numerically up to five loops (tenth-order). This involved evaluations of 12672 Feynman diagrams [10].

Table 1.1: Contributions to a_μ^{SM} and associated errors. Note that while it it is the largest contribution, the QED term has the smallest error. The hadronic terms dominate the total uncertainty. Also shown are the measured $\mathrm{Run}-1$ value from Fermilab and the new average a_μ^{exp} value. The average experimental value is 4.2σ higher than the calculated Standard Model value.

Term	Value $\times$ 10^{11}	Error $\times$ 10^{11}	
QED	116 584 718.931	0.104	[10, 11]
Electroweak	153.6	1.0	[12, 13]
HVP	6845	40	[14–20]
HLbL	92	18	[21–35]
Total SM value a_μ^{SM}	116 591 810	43	[4]
a_μ^{FNAL}	116 592 040	54	[1]
Average measured a_μ^{exp}	116 592 061	41	[1]
$\Delta a_\mu = a_\mu^{exp} - a_\mu^{SM}$	251	59	[1]

The mass-dependent QED terms are expanded into terms involving $\frac{m_\mu}{m_e} > 1$ and $\frac{m_\mu}{m_\tau} < 1$. Analytic expressions for these terms exist up to sixth-order, while eighth- and tenth-order contributions are calculated numerically. The twelfth-order contribution has been shown to be smaller than is important in this context, but an estimation of this term is included in the error quoted on the QED term. Remarkably, the total error on the QED loop calculations is less than one part in a trillion [10].[2]

Electroweak loops, such as the second from the left in Figure 1.4, are generally suppressed

---

[2] The calculated contributions from QED loops depend heavily on the experimental value of the fine structure constant α. Two experiments measuring α using the recoil of ^{133}Cs [37] and ^{87}Rb [38] atoms find values more than 5σ apart. While unimportant on the scale of a_μ^{SM}, this difference would make a noticeable shift in the value of a_e^{SM}. The Rb-experiment predicts a value of a_e^{SM} $+1.6\sigma$ from the experimental value [39] whereas the less precise Cs-recoil data gives a value of a_e^{SM} -2.4σ from a_e^{exp} [37].

by the very small mass ratios $\left(\frac{m_\mu}{m_{Z/W/\text{Higgs}}}\right)^2$. The one-loop terms involving W- and Z-boson exchange are known analytically, whereas the Higgs exchange diagram is negligible. Two-loop terms are calculated numerically, and the resultant value for the electroweak contribution to a_μ^{SM} is known to sub-percent level precision [13].

1.3.2 Quantum chromodynamics (QCD) loops

Unlike the QED contributions, the perturbation theory expansions at low energy cannot be used to calculate the contributions from QCD terms entering as loops in the muon-photon vertex. Instead, techniques relating experimental e^+e^- data to hadronic cross-sections and lattice QCD are used.

Hadronic vacuum polarization (HVP)

The HVP term encapsulates diagrams in which the virtual photon in the muon-photon vertex decays into hadrons, as seen in the third diagram of Figure 1.4. Integrals of this form can be calculated through the use of dispersion relations. In the case of the magnetic anomaly of the muon [15]:

$$a_\mu^{\text{HVP, LO}} = \frac{\alpha^2}{3\pi^2} \int_{M_\pi^2}^{\infty} \frac{K(s)}{s} R(s) ds, \tag{1.17}$$

where s is an energy scale, $K(s)$ is a kernel function that behaves as $K(s)\ m_\mu^2/3s$ at low energies, and $R(s)$ is the hadronic R-ratio given by

$$R(s) = \frac{\sigma^0(e^+e^- \rightarrow \text{hadrons} + \gamma)}{4\pi\alpha^2/3s}. \tag{1.18}$$

Here, $\sigma^0(e^+e^- \rightarrow \text{hadrons} + \gamma)$ is the experimentally measured bare cross-section. The $1/s$ term in the integral of Equation 1.17 leads to the low energy region dominating this term. This is the leading order term; higher order terms are significant but small.

Experimental e^+e^- data is taken from numerous experiments, and combining results requires some ingenuity. Herein lies the reason for the large error bar on the HVP calculation. In particular, the data from KLOE and BABAR do not agree in the $\pi^+\pi^-$ channel, which accounts for approximately 3/4 of the hadronic contribution to a_μ^{SM} [4]. To account for the tension, a conservative systematic error was assessed on the combined $\pi^+\pi^-$ cross section in the evaluation of $R(s)$. The other experiments with data in this channel (CMD-2/SND/BESIII/CLEO) do not favor the results from either KLOE or BABAR [4]. As more e^+e^- becomes available, reduction in the size of the error on the HVP calculation will be possible. The various channels and the total value for $R(s)$ can be seen in Figure 1.5.

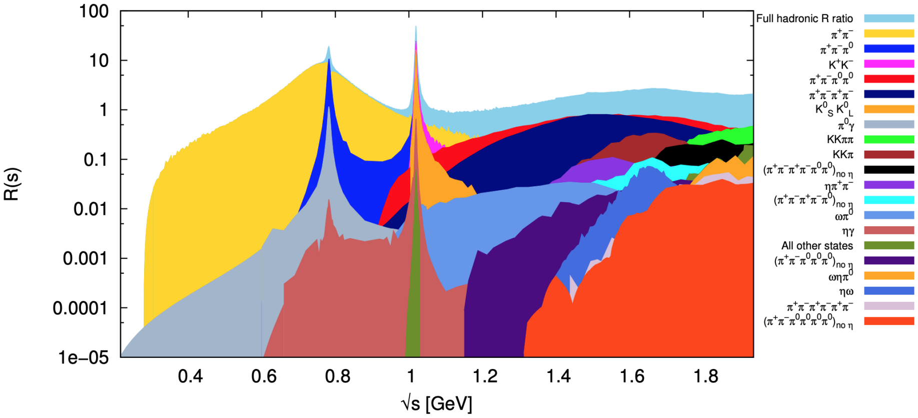

Figure 1.5: Hadronic channels contributing to the $R(s)$ term in Equation 1.17. The $\pi^+\pi^-$channel shown in yellow contributes approximately 75% of the total a_μ^{HVP} value, in part because of the heavy weighting at low energies. Image reproduced from [15].

Hadronic light-by-light (HLbL)

The hadronic light-by-light term, seen furthest right in Figure 1.4, involves hadronic blobs that exchange three photons with the muon and interact with the external field. While the smallest contribution to a_μ^{SM}, the HLbL term contributes a not insignificant error on the Standard Model value, as it is difficult to calculate. Hadronic model calculations were typically used to calculate this value [40, 41], but more recently dispersion relations were crafted to perform the calculation in the same style for the HVP term [42]. The difficulty lies in crafting dispersion relations for each intermediate state; at present, the pion pole and pion box contributions are known very well with the two-pion cut contribution known to 10% accuracy. The slightly less precise lattice QCD results support the value determined from the dispersion relations [34].

1.4 Measuring a_μ with a magnetic storage ring

One key element allowing the Muon $g-2$ experiment in its current form is the parity-violating nature of the weak decay. Since neutrinos emitted in the decay of a muon will have a well-defined chirality, the resultant daughter electron (or positron for antimuons) will carry information about the parent muon's spin. The same parity-violating weak interaction governs the decay of pions into muons. Since pion beams are easy to create, a muon beam with a high degree of polarization is possible. The polarized muons can then be put into a magnetic field, and the decay products will give information about the spin direction at the time of decay.

Returning to Equation 1.9, an approximation will made in the case of $\vec{E} = \vec{\beta} \cdot \vec{B} = 0$. This leads to a rotation in the spin direction at a rate of

$$\frac{d\hat{s}}{dt} = \omega_s = \frac{eB}{mc}\left(\frac{g-2}{2} + \frac{1}{\gamma}\right), \tag{1.19}$$

The motion in the magnetic field leads to cyclotron motion, where the direction of travel changes at a rate of

$$\frac{d\hat{\beta}}{dt} = \omega_c = \frac{eB}{\gamma mc}, \tag{1.20}$$

under the same assumptions about the $\vec{B}$ and $\vec{E}$ fields given above. The difference between the two frequencies is given by:

$$\omega_a = \omega_s - \omega_c = \frac{eB}{mc}\frac{g-2}{2} = \frac{eB}{mc}a_\mu, \tag{1.21}$$

by the definition of the magnetic anomaly $a_\mu \equiv \frac{g-2}{2}$ given before. Thus, the spin precesses about the momentum as the muon travels within a magnetic field if the g-factor of the muon is non-zero. Corrections to this equation for the presence of an electric field (used for containment) and non-transverse motion in the magnetic field will be required.

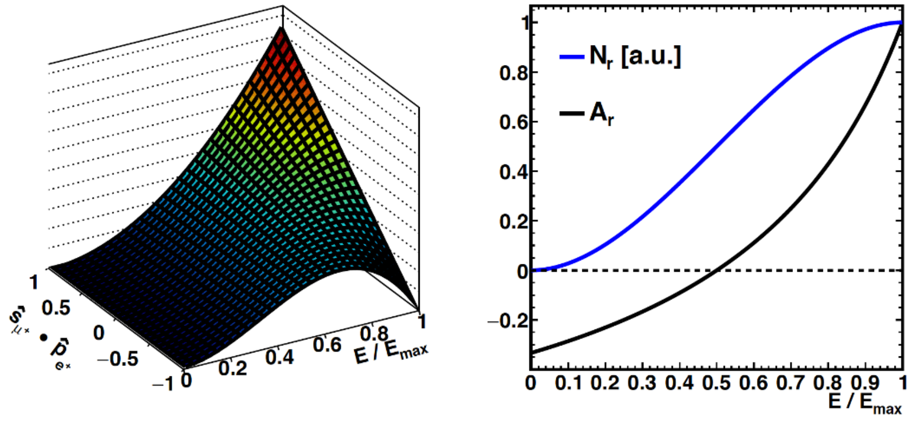

Figure 1.6: Parity violation in the weak decay of the muon. The left plot shows the relative number of muon decays as a function of the relative alignment of the muon spin and decay positron momentum in the muon's rest frame. $E_{\max}$ is the maximum energy possible for the decay positron, nearly $m_\mu c^2/2$. The right plot shows the relative number of decays N_r and asymmetry A_r vs. decay positron energy in the rest frame of the muon. The parameter A_r is a measure of how correlated the positron energy and muon spin are. Image reproduced from [43].

Due to the parity violation present in the muon decay, the rate of the highest energy decay positrons will oscillate as the spin rotates about the momentum. The highest energy decays are emitted preferentially in the direction of the muon's spin in its rest frame; as the spin aligns and anti-aligns with the momentum, the energy of the decay positrons will be boosted forward and backward. Up to other effects, the number of high energy decays in a particular direction will go as

$$N(t) = N_0 \exp(-t/\gamma\tau_\mu)\left(1 + A\cos(\omega_a t - \phi)\right), \tag{1.22}$$

where N_0 is a normalization, $\gamma\tau_\mu$ is the Lorentz-boosted lifetime of the muon, A is the asymmetry, ω_a is the anomalous precession frequency, and ϕ is a phase. In practice this average phase ϕ depends on several effects, some of which are positron energy dependent: the phase of the muon spins at injection, the drift time from muon decay to positron detection, and muon decay position within the storage region.

The simplified expression for the magnetic anomaly a_μ is:

$$a_\mu = \frac{mc}{eB}\omega_a. \tag{1.23}$$

On top of the ω_a measurement technique above, the strength of the magnetic field B must be determined. To do so, pulsed nuclear magnetic resonance (pNMR) probes will measure the magnetic field through the response of a shielded proton:

$$B = \frac{\hbar\omega_p'(T)}{2\mu_p'(T)} = \frac{\hbar\omega_p'(T)}{2}\frac{\mu_e(H)}{\mu_p'(T)}\frac{\mu_e}{\mu_e(H)}\frac{1}{\mu_e}. \tag{1.24}$$

The ratio $\frac{\mu_e(H)}{\mu_p'(T)}$ is the ratio of the magnetic moments of an electron bound in hydrogen to the magnetic moment of a proton shielded in a water sample measured at a temperature $T_r = 34.7°\mathrm{C}$ [44]. The term $\frac{\mu_e}{\mu_e(H)}$ is the ratio of the magnetic moment of the electron bound in hydrogen to that of the free electron [45]. Using $\mu_e = \frac{g_e}{2}\frac{e}{m_e c}\frac{\hbar}{2}$:

$$a_\mu = \frac{\omega_a}{\tilde{\omega}_p'(T_r)}\frac{m_\mu}{m_e}\frac{\mu_e(H)}{\mu_p'(T_r)}\frac{\mu_e}{\mu_e(H)}\frac{g_e}{2}. \tag{1.25}$$

The mass ratio $\frac{m_\mu}{m_e}$ is known from measurements of the hyperfine splitting of muonium [46]. The g factor of the electron is known from quantum cyclotron measurements [39]. A summary of these constants can be found in Table 1.2.

Here, the notation $\tilde{\omega}_p'(T_r)$ has been used to indicate that the pNMR measurements of the magnetic field are taken as a weighted temporal and spatial average of when/where the muons are stored for data-taking. The ratio given by $\mathcal{R}_\mu' \equiv \frac{\omega_a}{\tilde{\omega}_p'(T_r)}$ encapsulates the measurements produced in the Muon $g-2$ experiment.

In practice, the measured frequencies ω_a^m and ω_p must be corrected for a variety of effects to be used to calculate a_μ. With the corrections in place, $\mathcal{R}_\mu'$ becomes

$$\mathcal{R}_\mu' = \frac{f_{\mathrm{clock}}\ \omega_a^m\,(1 + C_e + C_p + C_{ml} + C_{pa})}{f_{\mathrm{calib}}\ \langle\omega_p(x,y,\phi) \times M(x,y,\phi)\rangle(1 + B_k + B_q)}. \tag{1.26}$$

These terms will be addressed individually in the upcoming sections.

Table 1.2: Experimental values that factor into a_μ through Equation 1.25. The associated errors of each are also listed to compare to the errors on the values measured in the Muon $g-2$ experiment. Notice that the error on a_μ is heavily dominated by these locally measured values.

Parameter	Relative error (ppb)	Source
$\mu_e/\mu_e(H)$		CODATA. [45]
g_e	0.000 26	Quantum cyclotron. [39]
$\mu_e(H)/\mu'_p(T_r)$	10.5	Hydrogen maser. [44]
m_μ/m_e	22	Hyperfine splitting in muonium. [46]
ω_a	434	Muon $g-2$ Run $-$ 1 result. [47]
$\tilde{\omega}'_p(T_r)$	114	Muon $g-2$ Run $-$ 1 result. [48]

1.5 The shielded proton precession frequency $\tilde{\omega}'_p(T_r)$

$$\tilde{\omega}'_p(T_r) = f_{\text{calib}} \, \langle \omega_p(x,y,\phi) \times M(x,y,\phi) \rangle (1 + B_k + B_q) \tag{1.27}$$

To begin with, the measurement of the magnetic field by the pNMR probes $\omega_p(x,y,\phi)$ is made as a function of position within the storage region. This is this weighted by the density of the muon beam at each position $M(x,y,\phi)$. The two of these together provide the average field experienced by muons while precessing. The term f_{calib} relates the measurements of the pNMR probes (with petroleum jelly as an active volume) to that of a well-known water probe and includes corrections for temperature, shape, and material. The two terms B_k and B_q are transient fields created by systems used to inject and store the muon beam. The sizes of these parameters can be seen in Table 1.3.

1.6 The anomalous precession frequency ω_a

Now considering the numerator of Equation 1.26, corrections must be made for the assumptions that went into the Equation 1.21 (C_e and C_p) and for the complications of measuring ω_a (C_{pa} and C_{ml}).

1.6.1 Corrections for non-zero $\vec{E}$ and $\vec{\beta} \cdot \vec{B}$ terms

Without the assumptions that $\vec{E} = \vec{\beta} \cdot \vec{B} = 0$, the expression for ω_a is:

$$\omega_a = \frac{e}{mc}\left(a_\mu \vec{B} - a_\mu \frac{\gamma}{\gamma+1}\left(\vec{\beta} \cdot \vec{B}\right)\vec{\beta} - \left(a_\mu - \frac{1}{\gamma^2-1}\right)\vec{\beta} \times \vec{E}\right). \tag{1.28}$$

This expression combines the spin rotation shown in Equation 1.9 with the proper time-changing momentum from Equation 1.8. The corrections to the measured ω_a value owing to non-zero $\vec{E}$ and $\vec{\beta} \cdot \vec{B}$ terms will be addressed in the upcoming sections.

1.6.2 E-field (C_e)

Without vertical containment, muons would spiral out of the magnetic field, so an electric field provides soft vertical focussing in the storage ring. By choosing a muon momentum such that

$$\frac{g-2}{2} - \frac{1}{\gamma^2-1} = 0,$$

the electric field contribution will go to 0. The "magic momentum" is $p \approx 3.1\,\mathrm{GeV}/\mathrm{c}$ or $\gamma \approx 29.3$. Since in practice the beam will have a finite momentum spread, this term will not completely vanish.

The electric field acts as a restoring force in the vertical direction (aligned with the magnetic field) and defocuses in the radial direction. Beam oscillations in both axes result. The frequencies of the oscillations depend on a quantity called the field index n:

$$n = \frac{R_0}{vB_0}\frac{\partial \mathcal{E}_y}{\partial y} \tag{1.29}$$

where R_0 is the ideal storage radius, v is the velocity of the muons, B_0 is the magnetic field, and the electric field gradient is calculated from the electrostatic quadrupole operating

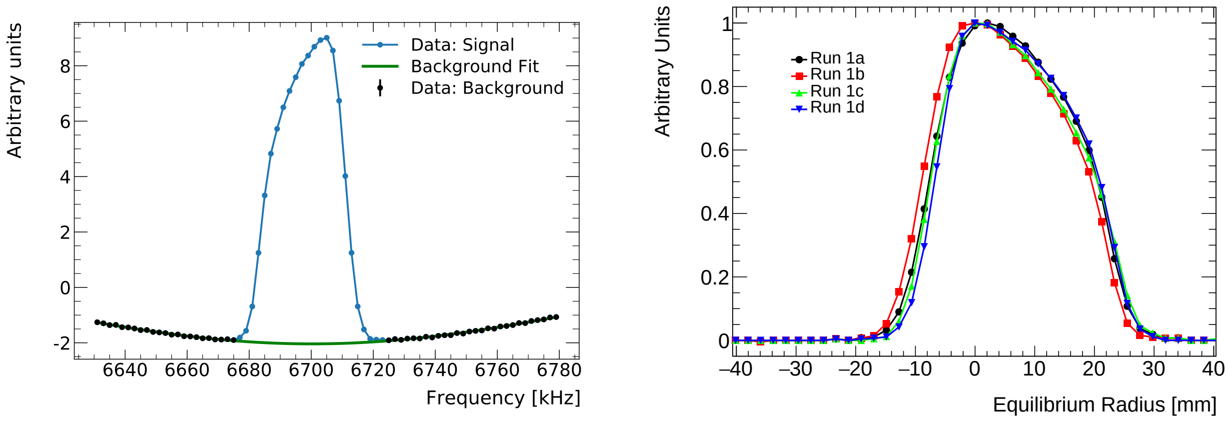

(a) Fourier transform of detector data

(b) Storage radius distribution

Figure 1.7: Cyclotron frequency distribution (left) and the stored radial distribution for each Run $-$ 1 sub-dataset (right). The correction to ω_a^m depends on the containing electric field experienced by the muons as they precess. Necessary to that calculation is knowledge of the radial distribution of the store muon beam, equivalent to the momentum distribution. Shown here are (left) a Fourier transform of detected decay positrons to determine the momentum spread and (right) the resultant radial distributions for the four Run$-$1 datasets. Reproduced from [49].

voltage and geometry. The frequency of oscillations in the $x-$ (radial) direction is $\omega_x = \omega_c\sqrt{1-n}$ where ω_c is the cyclotron frequency. In the $y-$ (vertical) direction, the oscillation frequency is $\omega_y = \omega_c\sqrt{n}$.

The mean radial electric field, the component that will contribute, experienced by a muon in the quadrupole field is given by [49]:

$$\langle E_r \rangle = \frac{n\beta c B_y}{R_0} x_e, \tag{1.30}$$

where x_e is the storage radius for that muon (rather than the ideal radius). The total electric field correction is given by

$$C_e \approx 2n(1-n)\beta_0^2 \frac{\langle x_e^2 \rangle}{R_0^2}. \tag{1.31}$$

To determine the distribution of x_e, a Fourier transform is taken of fixed-location detector data. As the beam circulates, the higher momentum muons will have a larger radius. Having a longer travel path with the same speed, the cyclotron frequency will be lower. This effect is known as the "fast rotation." The Fourier transform of positron decays provides the range of frequencies and converts to equilibrium radii through knowledge of the field strength. The Fourier transform and resultant radial distribution can be seen in Figure 1.7. The adjustment and associated error the electric field correction C_e can be seen in Table 1.3.

1.6.3 Pitch correction (C_p)

As with the E-field correction, the pitch correction stems from the approximation that a term in Equation1.9 vanishes. The term

$$\frac{g-2}{2}\frac{\gamma}{\gamma+1}\left(\vec{\beta}\cdot\vec{B}\right)$$

will necessarily gain a contribution from the vertical motion created by the restoring electric force, since there will necessarily be a spread of vertical angles upon injection into the magnetic field. The so-called pitch correction goes as: [49]

$$C_p \approx \frac{n}{2}\frac{\langle y^2\rangle}{R_0^2} = \frac{n}{4}\frac{\langle A^2\rangle}{R_0^2} \tag{1.32}$$

where A is the amplitude of vertical oscillation, R_0 is the ideal storage radius, and n is the field index defined above. In the experiment, the distribution of A values is measured directly by tracking detectors.

1.6.4 Early-to-late effects

$N(t) = N_0 \exp(-t/\gamma\tau_\mu)\,(1 + A\cos(\omega_a t - \phi))$, Equation 1.22 showed a model for the number of high-energy decay positrons detected in terms of the frequency of interest ω_a. Introducing a ϕ that changes throughout the measurement period:

$$\begin{aligned}\cos\left(\omega_a t - \phi\left(E\right)\right) &= \cos\left(\omega_a t - \left(\phi_0 + \frac{d\phi}{dt}t + ...\right)\right) && (1.33)\\ &= \cos\left(\left(\omega_a + \frac{d\phi}{dt}\right)t - \phi_0\right) && (1.34)\\ &= \cos\left(\omega_a' t - \phi_0\right). && (1.35)\end{aligned}$$

Any number of effects can pull effect the measured value of ϕ. The phase acceptance C_{pa} and muon loss C_{ml} terms are two of them.

Phase acceptance (C_{pa})

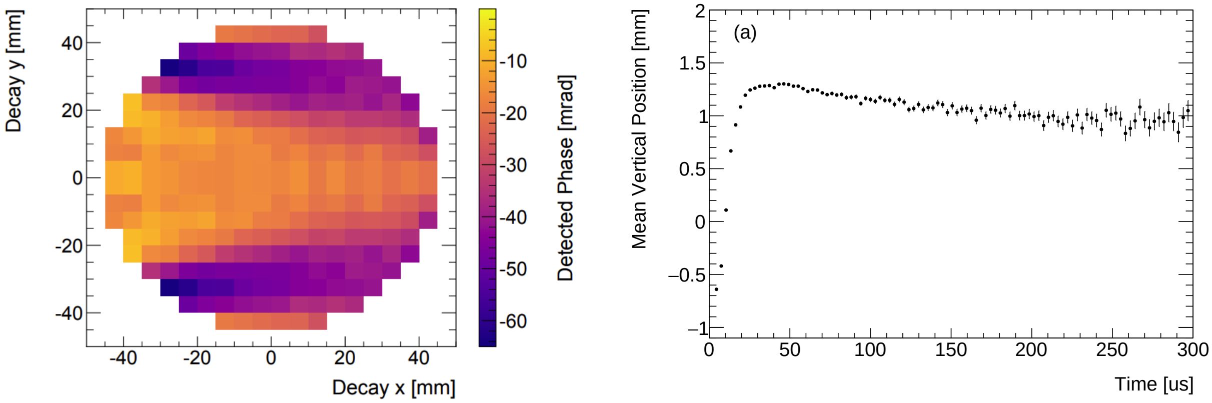

(a) Detected ϕ vs. decay position

(b) Change in vertical position through fill

Figure 1.8: Detected phase as a function of decay position (left) and mean vertical position through the fill (right). The left plot shows the detected phase ϕ as a function of radial and vertical decay position. When coupled with the drifting mean vertical position shown in the right-hand plot, the changing phase through the fill constitutes an early-to-late effect that pulls the measured value of ω_a. Reproduced from [49].

The measured phase ϕ depends on the decay position within the storage region, as seen in Figure 1.8. The acceptance of a decay also depends on the decay position. If there's coherent beam motion across the measurement period, the changing acceptance coupled to the changing phase would create an early-to-late pull on ω_a^m as shown above.

Combining the information from Figure 1.8a with extensive knowledge of the beam position and motion (Figure 1.8b), the average value for C_{pa} was determined. Extension simulation campaigns aided in extrapolating data from the fixed-location tracking detectors to the entire storage region. The C_{pa} correction and associated error is seen in Table 1.3.

Muon loss (C_{ml})

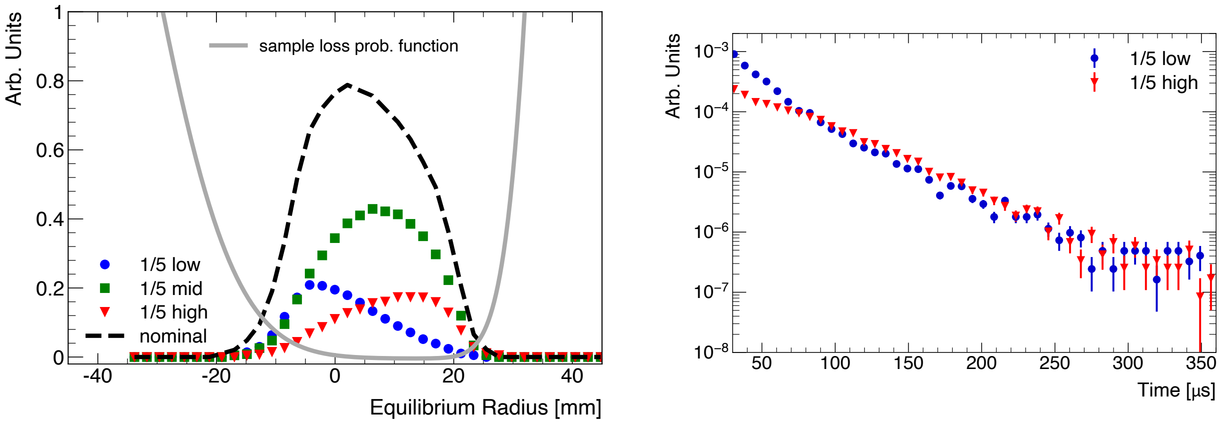

(a) Tuned momentum distributions

(b) Loss rate

Figure 1.9: Radial distributions for high- and low-momentum studies (left) and the relative muon loss rate for each (right). The left figure shows the fast rotation distributions as the beam momentum was tuned within the nominal storage range during physics data-taking. The right plot shows the loss rates for the most extreme of the momenta distributions created. To form a correction to the measured value of ω_a, the measured phase vs. muon momentum function was also determined experimentally in a separate study. Reproduced from [49].

Muons that are lost before decaying create a changing normalization factor N that must be accounted for when modeling ω_a behavior. If the muons are preferentially lost on one side of the momentum spread, a difference in measured phase ϕ for different momenta would lead to a pull on ω_a.

By tuning the magnetic field, tests were done at the nominal momentum and that ±0.6%. The measured difference in the average detected phase between the three datasets matched simulation data quite well. The loss rate as a function of the stored momenta spread was determined by taking eight datasets with average momentum scanned. Example of this process can be seen in Figure 1.9. The correction to ω_a^m was very small from this effect, as seen in Table 1.3.

Table 1.3: Values and uncertainties of the correction terms in $\mathcal{R}'_\mu$, defined in Equation 1.26. Also shown is the total uncertainty due to the constants in Equation 1.25. Positive C_i increase a_μ, while positive B_i decrease a_μ. Reproduced from [1].

Quantity	Correction Terms (ppb)	Uncertainty (ppb)
ω_a^m (statistical)	–	434
ω_a^m (systematic)	–	56
C_e	489	53
C_p	180	13
C_{ml}	-11	5
C_{pa}	-158	75
$f_{\text{calib}}\langle\omega_p(x,y,\phi) \times M(x,y,\phi)\rangle$	–	56
B_k	-27	37
B_q	-17	92
Total systematic	–	157
Total fundamental factors	–	25
Totals	544	462

1.7 Measured anomalous precession frequency ω_a^m

The majority of this document will focus on the hardware systems, data reconstruction, and analysis techniques employed to measure a value for ω_a. The overall clock frequency is blinded, leading to the adjustment term f_{clock}. Chapter 2 will discuss the Muon $g - 2$ experiment at Fermilab in its entirety, highlighting the necessary systems to make the measurements as described above. Chapter 3 presents the reconstruction chain for calorimeter data, from digitized pulses to decay positron energies and hit times. The necessity and manner of correcting for detector hardware effects is detailed. In Chapter 4, the function used to fit positron data to find ω_a^m is expanded upon and sample systematic studies are presented. Chapter 5 explores the implications of the 4.2σ deviation between the experimental measured a_μ and the calculated value a_μ^{SM}, followed by a discussion of upcoming improvements to both values. The intriguing Muon $g - 2$ result is used to explore other exciting opportunities in precision physics.

Chapter 2

THE MUON $g-2$ EXPERIMENT AT FERMILAB, E989

After a discussion of the broad experimental technique in the last chapter, this chapter will address the specific hardware used in the Fermilab experiment. Beginning with the accelerator complex, the polarized muon beam is created from smashing protons into a target and allowing the resulting pions to decay in flight. The muons are stored in the magnetic storage region using a highly specialized magnetic kicker and contained through the use of electrostatic quadrupoles.

To measure the magnetic field, a suite of pulsed nuclear magnetic resonance (pNMR) probes continuously measure the field outside the storage region. Periodically, a trolley containing the same pNMR probes traverses the azimuth of the storage region to map the response of the probes outside the storage region to the field the muons experience. Approximately once per year, the trolley probes are calibrated to a very well-known water probe through an ingenious procedure of swapping the positions of each repeatedly. The average magnetic field within the storage region must be averaged based on the location of the muon beam, provided by two stations of straw tracking detectors.

As the muons decay, the lower mass positrons spiral inward to the calorimeters stationed around the inner radius of the storage ring. The calorimeters report a hit time and energy for each particle that can be used to measure the precession frequency. An extensive laser system measures the stability of the calorimeter energy response and helps correct for things such as temperature drifts.

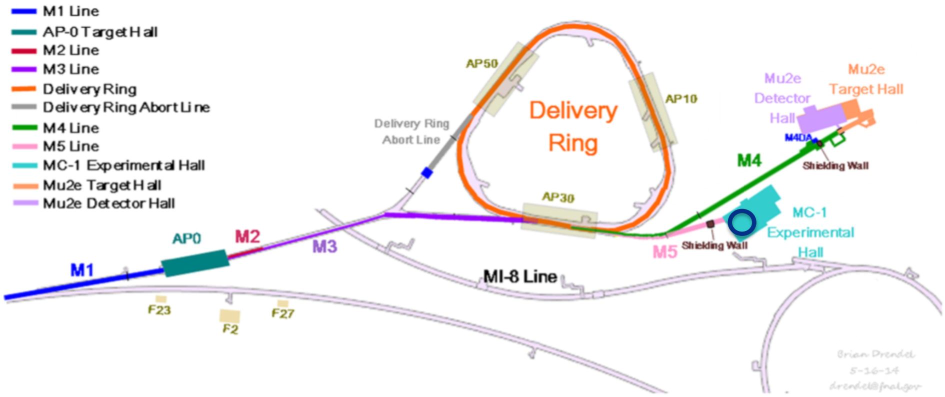

Figure 2.1: Schematic of the Fermilab beamline leading to the Muon $g-2$ storage ring. Bunches of protons at 8 GeV from the M1 line in the lower left are smashed into a target in the AP-0 Target Hall. Positive particles (pions, protons, positrons) at 3.1 GeV / c are selected; the pions decay in flight into polarized positive muons. After four turns in the delivery ring, the time-of-flight difference between the muons and protons is large enough to extract the muons down the M5 line to the $g-2$ storage ring and kick the protons into a beam dump. Image reproduced from [50].

2.1 *The accelerator complex*

Starting with a bottle of hydrogen gas, single protons are accelerated to 400 MeV using a linear accelerator. The protons are then accelerated in the Fermilab booster ring to 8 GeV. The 8 GeV protons are formed into bunches of roughly 10^{12} particles each using RF (radiofrequency) cavities in the recycler ring. The bunches are extracted one at a time and smashed into an Inconel target, chosen for properties such as its density, cooling ability, and resistance to corrosion. A secondary beam is collected using a lithium lens. The resulting positive particles at 3.1 GeV / c ($\pm 10\%$), primarily positive pions (π^+), are selected using a bending magnet. The branching ratio of positive pions to positive muons $\mathrm{Br}\left[\pi^+ \to \mu^+ \nu_\mu\right] = 0.9998770(4)$ and the lifetime is remarkably short ($\tau_{\pi^+} = 26.033(5)\,\mathrm{ns}$ at rest) [51]. When accounting for the relativistic $\gamma \approx 22.2$, the decay length of the pions is 173.35 m. About 80% of the pions decay in the 280 m beamline between the target station and the delivery ring. Because the pion decay violates parity, the decay muons' spins are very tightly aligned with the direction of travel. Once in the delivery ring, four turns around the 505 m circumference allows any remaining π^+ to decay and time of flight separation to occur between muons and the heavier protons present. The positive muons (with some other positive particle contamination, primarily positrons and deuterons) are extracted to the M5 beamline towards the $g-2$ experimental hall, and a kicker directs the protons into a beam dump. A schematic of the Fermilab accelerator complex can be seen in Figure 2.1.

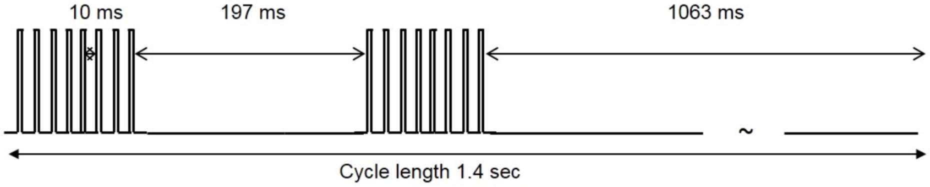

Figure 2.2: Temporal pattern of injection into the $g-2$ storage ring. This pattern is repeated every 1.4 s, for an average rate of ≈ 12 Hz. The time between fills, more 150× the boosted muon lifetime, is sufficient for all muons to decay before subsequent injection. Due to the RF technique used to create the bunches, the first and second sets of 8 fills are nearly identical. Image reproduced from [52].

This whole process happens 16 times every 1.4 s, an average rate of 11.4 Hz. The timeline of this process can be seen in Figure 2.2. Because of the characteristics of the RF cavities, the first 8 bunches are identical to their counterparts in bunches 9-16. The gap between each fill within a group of eight is 10 ms, and the gap between the two batches of 8 fills is 197 ms [52]. When the Fermilab test beam facility is running, the beam is redirected away from the muon campus for approximately 3 s every minute. Each bunch consists of approximately 10^{12} protons-on-target, resulting in approximately 7.8×10^5 μ^+ at the entrance to the storage ring [50]. Accounting for $\approx 1\%$ injection efficiency, $\approx 10\%$ of e^+ decays falling within the desired energy range, and the restricted physics window of $[30, 650]\,\mu$s after injection, the experiment might expect on the order of 500 high-energy decay e^+ per fill (colloquially called CTAGs). This is in good agreement with the number observed in the experiment.

2.2 Muon beam storage

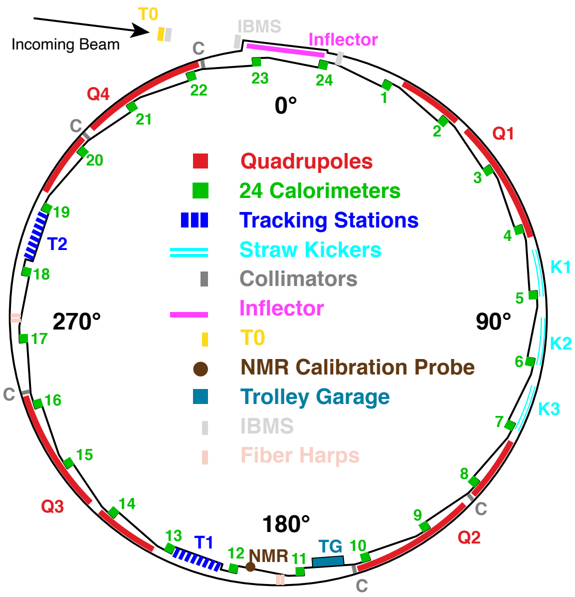

Figure 2.3: Diagram of the storage ring and associated equipment, as seen from above. The pulsed magnetic kickers, electrostatic quadrupoles, and inflector all contribute to storage of the muon beam. The calorimeters, T0 detector, IBMS detectors, and straw tracking stations are different particle detectors that provide information about the beam and decay products [49].

Upon arriving at the ring, the particle bunch first encounters a series of detectors that report arrival time and temporal/physical beam shape, as seen in Figure 2.4. The first, called the T0 detector, is a $15 \times 10\,\mathrm{mm}^2$ piece of 1.0 mm thick scintillating plastic read out by two independent photomultiplier tubes (PMTs). The PMTs each have a neutral density filter between the light guide and the front face. Laser fibers are run from two calorimeters to provide the same laser pulses to the T0 detector that are used to synchronize the calorimeter channels. The laser light bypasses the neutral density filters. In principle, the two PMTs could utilize different density filters and run at different voltages, though the choice was made to use a Thorlabs NE20B-A filter ($\approx 1\%$) and each run at roughly 950 V (provided by a CAEN high voltage NIM unit) for redundancy. This allowed the proper dynamic range to measure the timing of the incoming muon beam with good precision and avoid any saturation or clipping effects in digitization while also providing the necessary gain for the PMT to see the incoming laser pulse (used for timing synchronization with the calorimeter system). Using a less dense filter and running at a higher PMT voltage in one of the PMTs would allow a more precise determination of the incoming beam size at the cost of timing the laser pulse well. The number of incoming particles in each bunch (on the order of 10^6 particles) can be estimated based on some single source calibration done in bench tests, though it was found that the relationship of response to the voltage applied to the PMT was inconsistent between the two situations. As such, the magnitude of the beam as seen by the T0 detector is used only as a very loose condition in data quality. The PMT pulses are digitized by a spare waveform digitizer in one of the calorimeter's electronics crates.

After the T0 detector, muons travel through a passage created in the storage ring magnet and then through a 18 mm-wide × 56 mm-high × 1700 mm-long nearly magnetic field-free region created by the superconducting inflector magnet, taken from the Brookhaven experiment. Prior to the inflector, the beam encounters the first two detectors of the Inflector Beam Monitoring System (IBMS). These two detectors each comprise a 16×16 grid of 0.5 mm-diameter scintillating fibers read out by $1 \times 1\,\text{mm}^2$ SiPMs. The SiPM pulses are digitized in a $1\,\mu$s window by a CAEN digitizer (most of the detectors in the experiment are digitized in a $700-800\,\mu$s window by custom built digitizers). These detectors provide the spatial beam profile and location in a crucial location, allowing feedback on how to better steer through the narrow inflector channel. A similar detector can be inserted at the downstream face of the inflector for more information about beam injection, though the IBMS3 detector only provides information about the horizontal distribution of the beam — it only has a set of 16 fibers running vertically. As this is a destructive measurement, it is only performed in

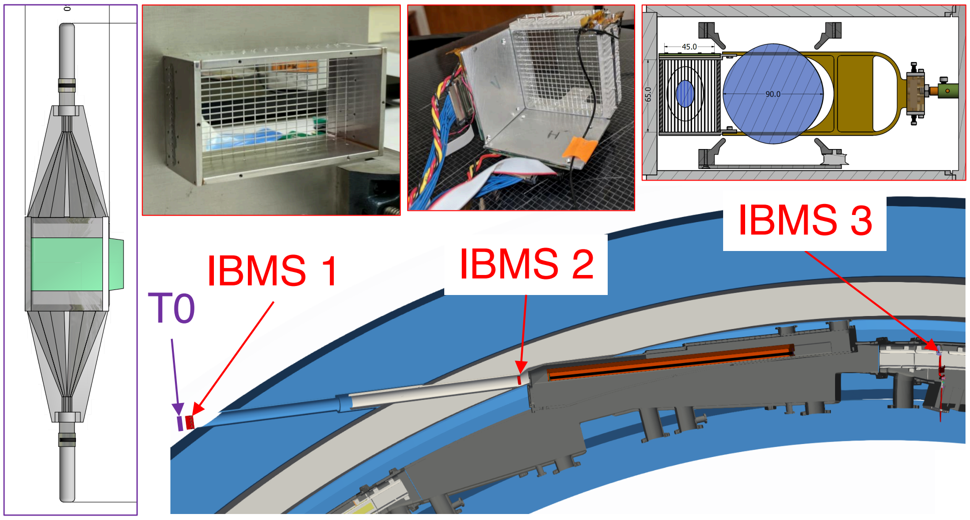

Figure 2.4: Beam entrance detectors. The four beam entrance detectors are shown alongside a graphic showing their locations along the beam line. The T0 detector is a piece of scintillating plastic coupled to two photomultiplier tubes through light guides. The three inflector beam monitoring system (IBMS) detectors comprise scintillating fibers read out by silicon photomultipliers. The third in the system (IBMS 3) is retracted during normal data-taking as it is destructive both for the incoming and stored beam. The other 3 detectors' impact on the beam is minimal given the ≈1% efficient transport through the inflector channel and subsequent storage.

special situations [49].

By necessity, the beam is injected 77 mm off the nominal storage orbit. A set of three 1.27 m-long high-powered pulsed magnetic kickers deflect the beam outward by applying a very short (temporally) vertical field, opposing the storage ring magnetic field, as the incoming beam passes through. The location of the kickers can be seen in Figure 2.3, approximately 90° from the injection point. The intention is to deflect the muon beam ≈ 10 mrad radially outward [53]. Ideally, the kickers would turn on prior to injection and turn off before the beam returns to the same area of the ring. This is immensely challenging to achieve, considering the 149.2 ns cyclotron period and the temporal width of the injected

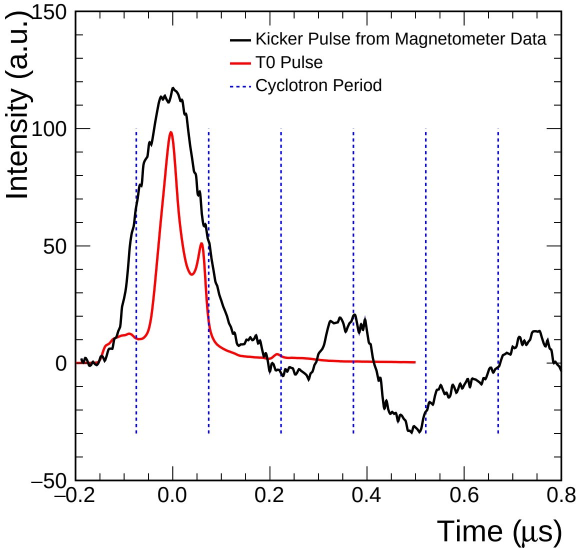

Figure 2.5: Kicker pulse and temporal structure of injected beam. Shown is a plot of the voltage supplied to the magnetic kicker plates overlaid with a temporal profile of the injected muon beam, as measured by the T0. Vertical lines indicate the cyclotron period in the storage ring. The goal of the kicker is to apply a 10 mrad deflection only the first time the muon beam encounters that section of the ring. Image reproduced from [49].

beam. A visual representation of this can be seen in Figure 2.5. The total "kick" required to achieve full deflection is on the order of 155 kV; in data-taking Run − 1 and Run − 2, the kicker system was unable to reach the optimal voltage to center the beam radially within the storage region. The off-center beam increases the necessary electric field correction described in the previous chapter and also increases the magnitude of the betatron oscillations that complicate the ω_a analysis. Partway through Run − 3, the kicker system achieved stability at the design voltage, and the beam was finally centered radially within the storage volume.

Once stored, a system of electrostatic quadrupole plates provides vertical focusing and containment for the beam. At the same time, the applied electric quadrupole field defocusses in the horizontal direction; the main magnetic field provides horizontal focussing. The locations of the quadrupole plates relative to the beam and the resultant field lines can be seen in Figure 2.6. Four stations, each comprising a long section and short section independently powered, subtend 43% of the circumference of the storage region. The locations of these can be seen in Figure 2.3. During the data-taking period of each fill, the quadrupole plates are held at 18.2 kV, corresponding to a field index of $n = 0.107$, as defined in Chapter 1.

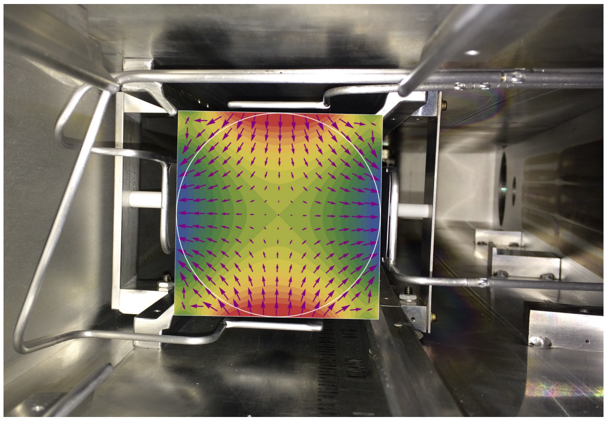

Figure 2.6: End view of electrostatic quadrupoles with field lines overlaid. To provide vertical containment within the storage ring, a system of pulsed electrostatic quadrupoles are powered with +18.2 kV on the top/bottom plates and -18.2 kV on the radially inner/outer plates. The end of one set of quadrupole plates is shown here with corresponding field lines; the total system covers approximately 43% of the circumference of the storage region. From this vantage, muons would travel into the page.

Two of the sections are asymmetrically powered following injection, a process called scraping: opposite plates within a section are held at voltages 5 kV apart. The asymmetric electric field displaces the beam downward and radially inward on one side of the ring and radially outward on the opposite side. This serves to collide any particles at the edge of the possible phase space for storage into one of the five 9 cm-inner diameter × 1 mm-wide × 1 cm-thick copper collimators that define the storage region. This reduces the number of muons lost from the storage region through non-decay modes during the nominal ω_a fit window within a fill. Roughly 7 μs after injection, the plates are brought back to full power, with a time constant of $\tau = 5\,\mu$s. Residual effects are present at the nominal ω_a fit start time of $t = 30\,\mu\text{s}$ after injection, but these effects are small and correctable through knowledge gained by the tracking detector system [49].

2.3 Knowledge of the magnetic field

The central component of the experiment is the 7.112 m radius superconducting magnet, moved from Brookhaven National Laboratory to its current location at Fermilab. A cross-section is shown in Figure 2.7. The magnet produces a nearly uniform dipole magnetic field with a strength of 1.45 T at an operating current of $\approx 5200\,\text{A}$. The field stability is driven by thermal expansion of the steel, and improvements to the temperature stability of the experimental hall have yielded impressive temporal stability. Several levels of tune-ability were designed into the magnet to facilitate the process of making the field uniform, also called "shimming."

The overall field strength is driven by the precise position of the 72 (upper and lower for each of 36 locations) pole pieces relative to the storage volume. The angle of the pole pieces relative to the volume determines field gradients. Additional pieces of iron were placed between the pole pieces and the storage volume (edge shims), between the pole pieces and the yoke (wedges), and on the top and bottom of the 24 separate yoke pieces (top hats) to fine-tune the average field as a function of azimuthal position and to adjust gradients across the storage volume. Additionally, a set of 100 current-carrying wires both above and below the storage volume allows an additional level of fine-tuning. Current distributions for the wires were created to minimize variations in the azimuthally averaged magnetic field across to the storage volume to below 1 ppm. The current distribution employed is updated periodically to better reflect the changing condition of the magnetic field [48]. For more information about shimming the magnetic field, the reader would be directed to References [54] and [55].

In order to precisely measure the magnetic field, a trolley carrying an array of 17 pulsed nuclear magnetic resonance (NMR) probes maps the magnetic field in the storage region by periodically circling the circumference of the ring, sampling ≈ 9000 azimuthal locations on each circumnavigation. See Figure 2.8 the placement of the NMR probes within the storage

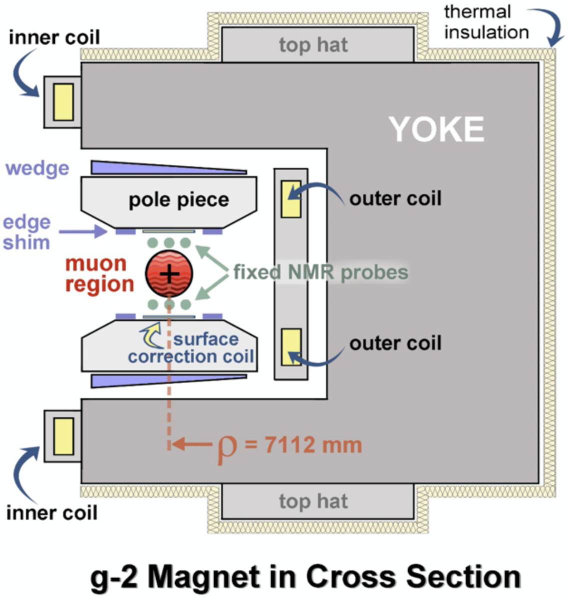

Figure 2.7: Cross-section of the storage ring magnet. The 9 cm diameter storage region is situated between the two pole pieces, where the magnetic field is nearly uniformly up. The built in hardware to shim the magnet is as follows (from most coarse adjustments to most fine): the top hats, wedges, edge shims, and surface correction coils. The labeled "fixed probes" are NMR probes used to track the magnetic field between periodic mappings. In this view of the storage region, the muon beam would be traveling out of the page. Image reproduced from [56].

volume. Since this would be destructive to the muon beam, this measurement cannot be concurrent with the precession frequency data-taking. Trolley runs are performed roughly every 3 days during data-taking. The main component of the magnetic field analysis is the digitized free induction decay (FID) signals from the NMR probes [57]. Extracting the frequency of the FIDs gives a measure of the average field within the active volume of each probe. The multipole moments of the field as measured by the trolley are used to take a weighted average with the corresponding moments of the muon distribution.

For the time between the trolley runs, a system of 378 NMR probes are embedded in the

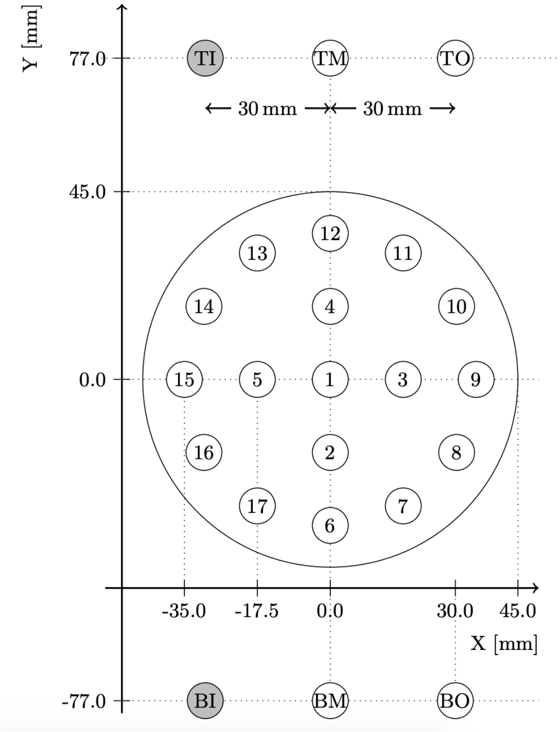

Figure 2.8: End view of the trolley used to measure the magnetic field. The array of 17 NMR probes samples the storage region as the trolley, designed to match the curvature of the magnet, is moved around the storage ring. Also shown are the locations of the fixed probes, used to track the magnetic field during normal data-taking operations. Image reproduced from [48].

vacuum chambers surrounding the storage region at 36 discrete azimuthal locations tracks the magnetic field between trolley runs in a non-invasive way. The "fixed" and trolley NMR probes, built at the University of Washington, use samples of petroleum jelly to measure the response of a shielded proton. These probes are also used in a power supply feedback loop to slightly increase or decrease the current within the ring as necessary to maintain a consistent dipole field. Because the probes are arranged spatially around the storage volume, the lower level multipole moments can be tracked in time [48].

A precise probe containing a cylindrical sample of pure water is used to calibrate the response of the trolley probes. To achieve the accuracy needed, effects such as temperature, material effects, and specific field conditions during measurement were taken into account when assessing the calibration probe's response. The response was cross-checked against probes containing a spherical sample of a water and ^{3}He [58]. The calibration probe's measurement of the magnetic field is better than 20 ppb [48].

Once a three-dimensional map of the magnetic field is formed, a weighted average is formed using knowledge of where the muons were stored within the storage region. This is the $M(x, y, \phi)$ term in Equation 1.26. To provide this information, two stations of straw trackers extrapolate the tracks of decay positrons to their decay points. Each station contains 8 modules; each module consists of 4 layers of 32 straws. Two of the layers in each module are at a 7.5° angle relative to the other two layers.

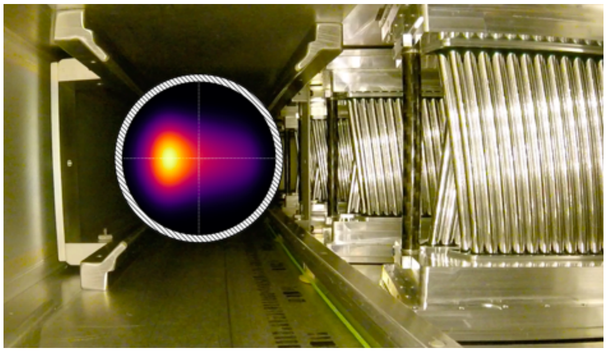

Figure 2.9: A straw tracker, as seen from the storage volume. When positrons (or muons) traverse the straws, electrons from the ionized argon ethane drift across the high voltage field held between a wire in the center of the straw and the outer surface. Tracks, reconstructed from multiple straw hits, are extrapolated to the point of decay to determine the muon beam's position. The reconstructed Run − 1 beam position is shown with the position of collimators shown by the thick border. In this image, the muon beam is moving into the page.

When a decay positron passes through the straws, it ionizes the argon ethane gas mixture inside and the freed electrons drift towards a sense wire. The timing and location of hits in the various straws/modules combined with a model of the magnetic field allow reconstruction of particle trajectories. One of the straw tracker stations can be seen in Figure 2.9; the reconstructed beam position also shown. In additional to providing information about the location of the beam for the average magnetic field, the trackers provide crucial knowledge about the beam dynamics for things such as the pitch correction term C_p (defined in Chapter 1) and adjustments to the ω_a fit function.

2.4 *Measuring the precession frequency*

The energy spectra of the decay positrons serves as a proxy for the spin directions of the parent muons. By measuring the energy and arrival time of particles at fixed locations near the storage region, the spin precession can be measured very precisely.

2.4.1 Electromagnetic calorimeters

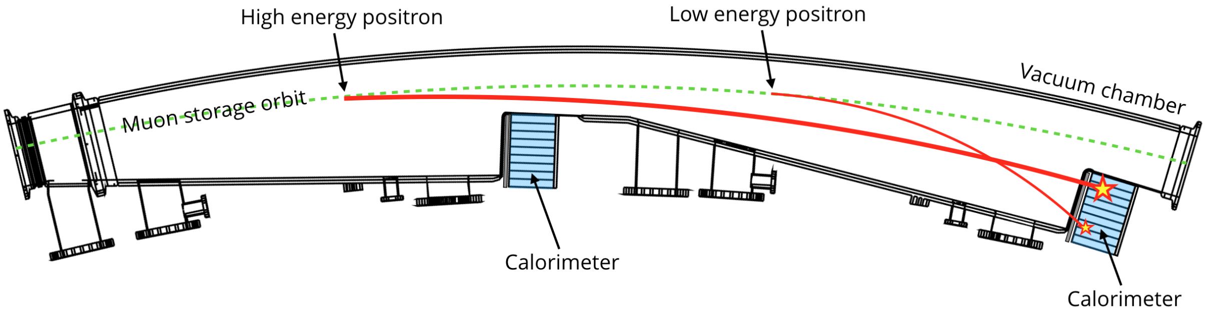

Figure 2.10: Calorimeter position with respect to muon storage region. As the muons decay, daughter positrons curl inward towards the calorimeters that record the hit times and energies. Example positron paths are shown. The vacuum chamber wall directly in front of the calorimeter face is thinner to reduce pre-showering. Image reproduced from [59].

The main tool for observing the spin precession of the muons in the storage ring is the oscillating rate of decay positrons of a particular energy. In this experiment, twenty-four calorimeters are stationed on the inner radius of the storage region, as seen in Figure 2.10.

As the muons decay, the lower mass positrons will curl inward in the magnetic field and strike a calorimeter.

Each calorimeter station consists of the active volume of fifty-four lead fluoride (PbF_2) and silicon photomultiplier (SiPM) readout, electronics to power the SiPMs, calorimeter enclosure cooling fan and power supply, digitizer crate, and laser distribution to the fifty-four channels.

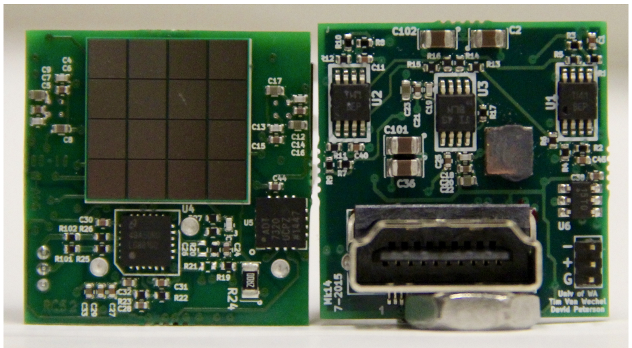

Figure 2.11: SiPM and on-board electronics. The 16 independent channels can be seen in the left panel. In operation, that portion is glued to the face of a PbF_2 crystal and the channels are summed to provide a single output signal per SiPM. The HDMI connection is used to provide bias voltage for the SiPM, power for on-board amplifiers, and communication with the programmable gain amplifier and temperature readout. Image reproduced from [60].

The $25 \times 25 \times 140\,\mathrm{mm}^3$ lead fluoride crystals produce pure Cherenkov light (index of refraction $n \approx 1.8$), a necessity for the stringent timing resolution required of the detectors. SiPMs were chosen as a readout tool for their compact nature and ability to function near the $1.45\,\mathrm{T}$ storage ring field without perturbing it. The SiPMs feature 16 independent channels consisting of ≈ 4000 pixels each that are then summed to produce one output signal per SiPM. On the SiPM electronics board are several stages of pulse amplification and shaping, including a programmable gain amplifier. Appendix A shows the circuit diagrams in detail. The gain amplifier is used to match the behavior of the SiPMs on separate crystals, a process described in Chapter 3.3. The SiPM and on-board electronics can be seen in Figure 2.11.

SiPMs are a form of avalanche photodiode. The Muon $g-2$ SiPMs are operated in Geiger mode, meaning that an individual pixel's response is binary: either it has fired or not. Each pixel is essentially a capacitor. Once fired, there is a characteristic recovery time (of order nanoseconds) before the pixel can re-fire. At a macroscopic scale, this means that two pulses of the same size in quick succession in one channel will produce output signals of differing sizes. Measuring and correcting for this effect is important when analyzing the calorimeter data.

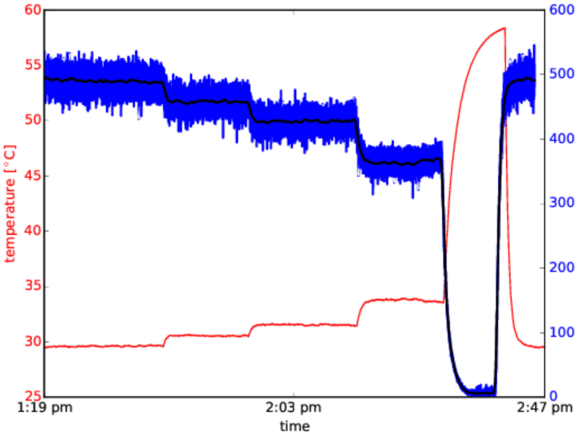

Figure 2.12: Changing SiPM response with temperature variations. As the temperature of the SiPM changes, the breakdown voltage also changes. According to lab tests, the breakdown voltage of the SiPMs employed is +0.07 V / °C. If the bias voltage applied to the SiPM is constant, the gain will go down as the temperature is raised as demonstrated here. Image reproduced from [61].

The output signal of a SiPM depends on two quantities per photoelectron: the bias voltage applied and the SiPM's temperature-dependent breakdown voltage. Signal size scales linearly with the difference between the two, or over-voltage $V_{ov} = V_{bias} - V_{breakdown}$. Each calorimeters' fifty-four SiPMs' bias voltages were provided by a set of four BK Precision 9124 power supplies running in constant voltage mode (at roughly 67.5-68 V) with a current limit of 0.02 V. This will protect the circuitry in the off-chance of a light leak in a detector.

It limits the rate at which the bias voltage can be turned on for SiPMs, but that rate is also limited by charging the capacitors used to quickly replenish SiPMs after large loads. Communication with the power supplies is achieved via USB from the two BeagleBone Black microcomputers installed in the electronics rack of each detector. The serial numbers of the four power supplies are stored on the BeagleBone to ensure repeatably. One early failure mode of the BeagleBones has been traced to using a USB hub that relied upon the BeagleBone for power. After installing independently powered USB hubs, the failure mode has ceased. Since the applied bias voltage was constant throughout data-taking, a correction to the gain was applied for the variable over-voltage due to the temperature-dependence of the breakdown voltage.

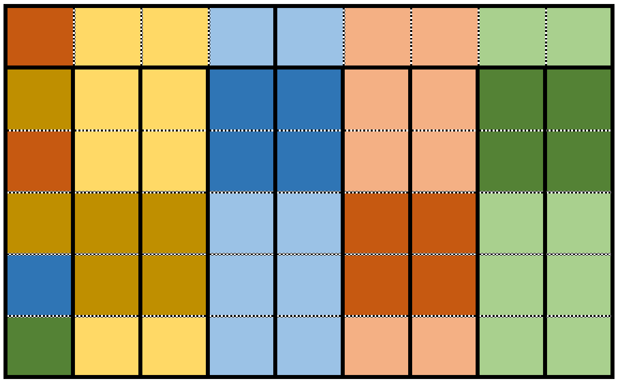

Figure 2.13: Digitizer and power supply connections within a calorimeter. The back view of a calorimeter (where beam is on the right of the detector coming out of the page) is shown. The solid black rectangles denote crystals that share a 5-channel digitizer. The colors (green, orange, blue, yellow) correspond to the 4 different bias voltage supplies. The darkened crystals are on separate voltage distribution circuits than the lighter crystals and therefore have a different characteristic voltage recovery time after a series of large impulses, such as the splash at injection. Cooling air flows from the left side of this schematic across a heatsink attached to each SiPM's electronics board, creating a gradient of temperatures from left (coolest) to right (warmest).

The SiPMs were arranged into groups when they arrived from the manufacturer Hamamatsu based on their production run. See Figure 2.13 for the power supply organization scheme. These production groups had similar breakdown voltages and were therefore grouped onto the same power supply within a detector station. The different power supplies' SiPMs were organized into columns following the temperature gradient created by the airflow within the detector enclosure, as the breakdown voltage depends on temperature (see Figure 2.12).

To arrange the 1300 crystals into twenty-four calorimeters, the crystals were sorted by their dimensions. The major constraints were to match the width of each crystal within a column and to have the top of the third row in each calorimeter (i.e. the midpoint) be level. The measurements of the dimensions were repeatable to within $\approx 0.02\,\mathrm{mm}$, based on the author's experience.

Before gluing the SiPMs to the crystals, each underwent a quality control procedure. For the SiPMs, this involved a collimated LED being flashed onto each of the sixteen channels to verify the existence and shape of the resulting electronics pulse. For the crystals, the dimensions of each was taken and large outliers were excluded from assignment to a detector. Obvious defects such as a failed polish or crystalline defect were identified visually. The transmittance of each crystal was measured [60].

The crystals were glued to the SiPMs using a two-part epoxy, ZEISS OK2030. The epoxy (index of refraction $n = 1.65$ at $\lambda = 400\,\mathrm{nm}$) was chosen as an intermediary between the lead fluoride ($n = 1.8$) and the SiPMs ($n \approx 3$). The hardening agent H950 was mixed in a ratio of 10 parts OK2030 to 1.8 parts H950. Exceptional care was taken to avoid any air bubbles in the layer of glue. The crystals were wrapped in $50\,\mu$m-thick black Tedlar to help with pulse shape and isolate the light generated in each crystal.

Important to the reconstruction technique (described in the upcoming chapter), the black wrapping maintains the same SiPM pulse shape independent of particle impact position on the crystal. Figure 2.14 shows the pattern for wrapping the sides of the crystals, intended to minimize the space between crystals within the active volume. A black absorptive wrapping was chosen as opposed to a white reflective wrapping as a tradeoff between maximizing light

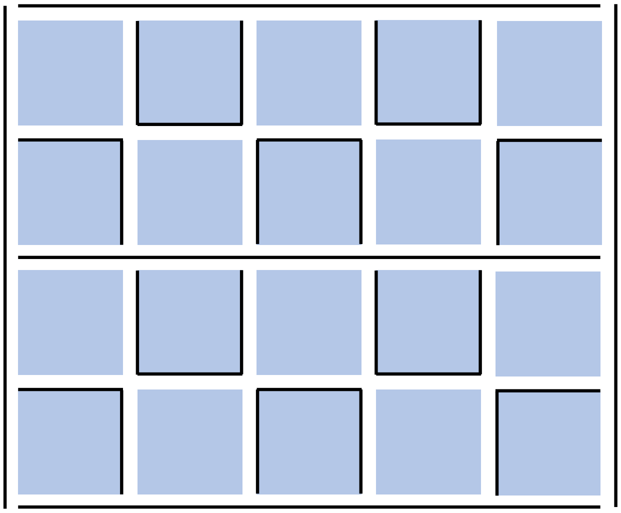

Figure 2.14: Crystal wrapping pattern. To minimize the amount of space between crystals within the calorimeters, the wrapping scheme shown was employed. Large, flat pieces of Tedlar are laid on the bottom of the calorimeter and on top of every second row. Similar pieces are put upright on the radially inward and outward edges. Within the pairs of rows, alternating 'U'-shaped pieces are folded around the crystals. At the edges, 'L'-shaped pieces are employed. The diagram only shows a 5×4 grid, but the pattern can easily be extrapolated to the true dimension of 9×6.

collection for energy resolution or optimizing pulse shape for timing resolution [62]. Between the front face of the crystal and the Delrin laser distribution panel is a thin layer of black plastic, used to isolate the crystal from glue in the laser panel that had previously blemished the front face of some crystals.

On board each SiPM is a series of pulse amplifiers and shapers. One of those is a LMH6881 variable gain amplifier with a range of amplification between $2-20\times$, a range of 20 dB [60]. To set the values of the variable gain amplifiers for the 54 SiPMs in each calorimeter, two BeagleBone Black microcomputers use an SPI data bus. The same communication protocol is utilized to periodically read out the on-board temperature sensor on each SiPM's printed circuit board. Additionally, a built-in EEPROM memory unit on each SiPM stores an identifying number that can be checked against a database. The power for the SiPM communication and pulse amplification is provided by two Acopian W5NT2500 regulated switching power supplies. An HDMI cable provides communication capability, bias voltage, and $\pm 5\,\mathrm{V}$ power for the electronics to the SiPM. HDMI was chosen for the number of pins it provides (19) within the $25\,\mathrm{mm}$ space constraint.

Because of the proximity to the storage region, the HDMI cables, SiPMs and on-board electronics, and signal cables all needed to be designed as non-magnetic to avoid perturbing the field. This was most challenging for the HDMI cables; all commercial non-magnetic cables explored were found to be too bulky to fit within the spatial constraints of the calorimeter housing. The main component of the cable that is typically magnetic is the ground surrounding the connector on either end, presumably because of some nickel content. The unfortunate result was the decision to remove the grounding clip from each end of every HDMI cable for the calorimeters. The process is not too difficult, but mortality was high as pins within the connecter would get disconnected. In order that the HDMI cables still fit snugly within a port, Kapton tape was wrapped around the end of the connector to replicate the width of before demagnetization. Before using the modified cables, the cables were tested with a simple apparatus that ran current from a $9\,\mathrm{V}$ battery through each pin to light up an LED. Should any pin fail or be unstable against small wiggles, the cable would be thrown out.

As measured during test beam runs involving prototype designs of the calorimeter, each detector has an energy resolution of 3.1% at $2\,\mathrm{GeV}$. This is better than the necessary 5% as quoted in the TDR. The single-channel timing resolution for an electron with energy of $3\,\mathrm{GeV}$ is $\approx 40\,\mathrm{ps}$. The energy response is stable vs. impact position on the calorimeter at about the 2% level. It is important to note that these results were obtained using a preliminary version of the reconstruction software and may have improved slightly with better pulse fitting and clustering techniques, but no significant change is expected.

The output signals of the SiPMs are digitized by custom waveform digitizers built at Cornell University. The digitizers sample at a rate of 800 mega-samples per second at a depth of 12-bits utilizing two interleaved analog-to-digital converters [47]. The blinded clock used derives from a 40 MHz clock, driven by a GPS-stabilized 10 MHz rubidium clock source. The 40 MHz clock is detuned to a frequency in the range $[39.997, 39.999]$ MHz, at a value chosen by someone outside of the experiment and kept secret until the collaboration decides to unblind. A second blinded clock is mixed with the master clock to monitor stability over time without knowing the blinded clock offset [47]. The digitizers use the same μTCA technology (as developed for the LHC experiments) with AMC13 readout. Each calorimeter station has a dedicated frontend and GPU to read in the data and reduce data size online. The primary tool is called "island chopping" whereby a window of the digitized traces is saved only if a channel within that detector crosses a 60 ADC threshold above a calculated baseline. This window is typically ≈ 30 ns in length. In commissioning and the first data-taking run, the full calorimeter's traces would be saved for each island. Subsequently this was reduced to the channel and any physically bordering channels, with the full detector readout being saved on a prescaled basis.

2.4.2 Laser system

As discussed above, the laser system has two main purposes: to provide a method of synchronizing the 1296 calorimeter channels employed and to measure and correct for different modes of gain fluctuations in each channel. The general layout of the laser system can be seen in Figure 2.15. Six Picoquant LDH-P-C-405M laser diode heads provide light for four calorimeters each. The laser produces pulses of width < 90 ps (FWHM) at a wavelength of 405 nm. A remotely controllable "filter wheel" (Thorlabs FW212CWNEB) can be used to modulate the light coming from each laser head, with 12 possible positions ranging from $< 0.01\%$ to 100% transmission. The default position has $\approx 22\%$ transmission. The silica fibers are protected by plastic casing and run under a false floor in the center of the storage ring to avoid damage to the fibers themselves. Along the optics path, light is siphoned off and redirected to three sets of stability monitors.

The first, the "source monitor," provides a measure of the laser power drift over time due to things such as temperature fluctuations along the optical path. Roughly 30% of the light

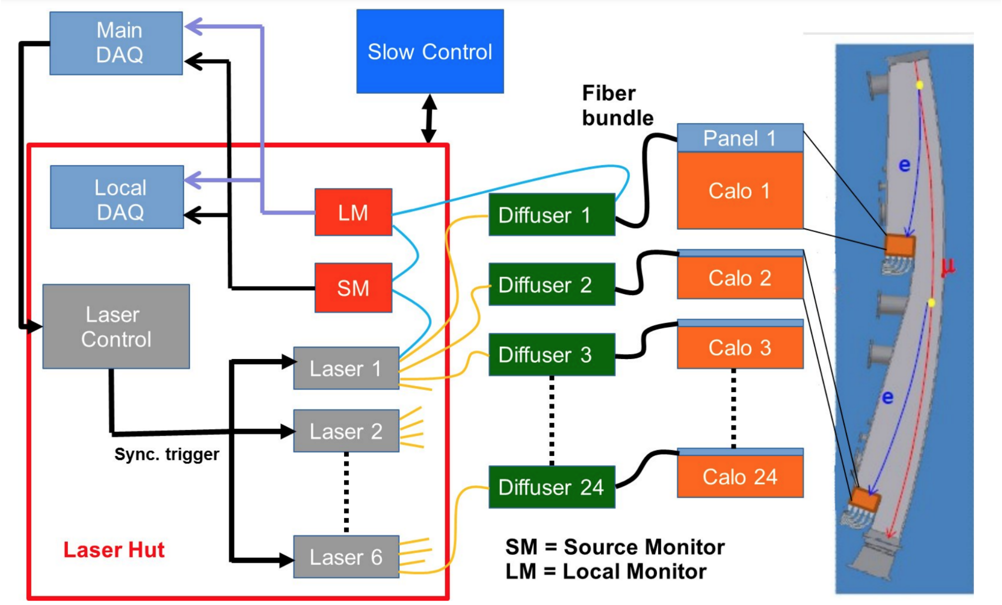

Figure 2.15: Schematic of the laser distribution system and associated stability monitors. Six laser heads provide the pulses for four calorimeters each. For each laser there is a PIN diode, called the source monitor (SM), to track the laser intensity. For each calorimeter there is a PMT, called the local monitor (LM), that receives a pulse before and after distribution to the calorimeter itself. Image reproduced from [63].

from each laser head's path is diverted to the source monitor using a beam splitter placed just after the laser head. The light is injected into an integrating sphere that disperses the light to two PIN diodes (PINs) and a photomultiplier tube (PMT). The PINs receive on the order of 10^6 photoelectrons per laser pulse, corresponding to a precision of 0.1% per shot. Variations in the PIN's response come primarily from temperature fluctuations, which is unfortunate given the poor temperature stability achieved in the experimental hall during the first two runs of $g - 2$ at Fermilab. Data from the source monitor detectors is taken similarly to that from the calorimeters, as discussed previously.

The second set of stability monitors measures the consistency of the distribution path to each calorimeter. Forty-eight (two per calorimeter for redundancy) PMTs (Hamamatsu

R1924A) receive two pulses per laser shot. The first is light taken from the integrating sphere in the source monitor, used as a stability check when compared to the source monitor detectors' response. The second is light redirected from the diffusers within the calorimeter laser panel back to the room with the laser components and optical table. The two pulses appear ≈ 240 ns apart in the digitized traces due to the travel time within the 50 m of fiber.

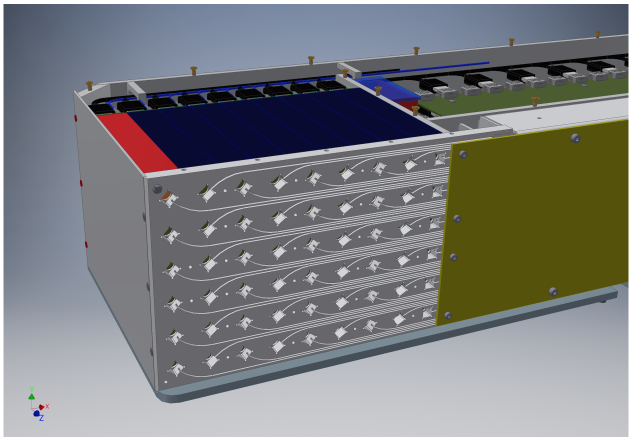

Figure 2.16: Laser light distribution to the 54 channels in a calorimeter. A computer model of the calorimeter shows the positions of the crystals (shown in dark blue) relative to the laser distribution panel on the upstream side of the detector. Each channel receives an almost simultaneous pulse of light from a very stable laser; these pulses are used to measure and correct for gain in each channel as well as for aligning the channels in time. The visible differences in the laser fiber lengths will become important when aligning the channels in time. The SiPMs are glued onto the back face of the crystals relative to the laser panel. Image reproduced from [59].

At the calorimeter, the transport fiber encounters a diffuser that spreads light uniformly across 60 fibers. Two of those are used to redirect light back to the local monitor and 54 are used to distribute the light to the upstream face of the crystals. The fibers terminate at $10 \times 10\,\mathrm{mm}^2$ right-angle prisms that redirect the light into the lead fluoride crystals (see Figure 2.16). The fibers and prisms are embedded in a 12 mm wide piece of Delrin, corresponding to a radiation length of $0.044\,X_0$. The light cone is sufficiently wide to uniformly illuminate the back face of the crystal (and therefore the SiPM attached). Three lengths of fibers are used to distribute light from the diffuser to the crystals, a fact that will become important when discussing the time synchronization of different channels within a calorimeter.

A third set of stability monitors comes in the form of isolated SiPMs of the same construction as those installed in the calorimeters, dubbed the super local monitors (SLMs). For six calorimeters (one receiving light from each laser head), a bare SiPM is enclosed in a light-tight box outside the calorimeter housing. The light comes from an unused fiber in the bundle of the laser panel. It receives bias voltage from one of the four power supplies discussed, but in a way that bypasses some electronics components internal to the calorimeter. Communication with the Beaglebone and slow controls system is done in the same way as the SiPMs internal to the calorimeter. This small detector is useful to study how the shared nature of the bias voltage may impact the particle data with a set of laser-only data. Since the same digitization scheme is used, it's also a tool to track the timing stability of the other SiPMs within that calorimeter.

Chapter 3

ELECTROMAGNETIC CALORIMETER SYSTEM

The calorimeter system is the primary tool for measuring the anomalous precession frequency ω_a in the Muon $g-2$ experiment. As discussed previously, the calorimeter system comprises 24 individual stations with 54 channels each. The channels' signals are digitized individually, rather than summing the output of the full detector. Each channel has a characteristic pulse shape that is constructed and then used to calculate the pulse integral and peak time of the pulse.

For a usable detector, the desired physics signals must lie well within the dynamic range of the SiPMs. Balancing the needs of calibrating to the 170 MeV peak, reducing electronics noise, and maintaining resolution in the face of pileup drives the choice of operation point. Using the laser system, the operating bias voltage and programmable gain of each channel can be optimally chosen.

To combine data from the various detectors and channels, the times of all channels in each fill must have a common start. The synchronization of these channels is done using muons that travel between detectors at a known speed and positrons whose energy is shared amongst several channels within a calorimeter. An energy scale calibration is performed on a channel-by-channel basis using the same muons utilized to calculate the timing offsets between detectors. To avoid biases in the ω_a measurement, the gain of each channel must be incredibly stable across a fill. The laser system provides a method to measure and correct for gain fluctuations.

Once the timing and energy of the pulse fits from each channel have been performed, the various crystal hits are clustered into a usable physics object. This chapter will detail the procedure to go from digitized trace to a positron hit with associated time and energy.

3.1 Data acquisition

The digitizers are only active during 800 μs-long fixed windows, triggered by a signal from the accelerator complex that beam will be arriving. Sixteen bunches come every 1.4 seconds. For each fill, the digitizers are activated approximately 100 μs before beam arrival. The first laser pulse ('sync' pulse) is fired 5 μs after digitization begins. For each channel, a pedestal (baseline) value is calculated outside the digitization window. A GPU farm compares the incoming digitized signals to the pedestal; whenever a channel's digitized signal-minus-pedestal is above a configurable threshold ($\approx$ 60 MeV), a window of time is saved for that channel and the physically bordering channels within the same calorimeter. A typical length of this island is 8 samples before the signal peak and 20 samples following. If another digitized sample within that island is above threshold, the island will be extended. A typical island for a single calorimeter can be seen in Figure 3.1.

Three levels of the reconstruction are performed and used for different levels of analysis. A very pared-down version is run in almost real time to provide diagnostics about how the experiment is operating through the Data Quality Monitoring page. A more robust, but still incomplete, version of data reconstruction called "nearline processing" is run within hours of data being taken. This allows for trends related to physics to be studied and monitored as well as quick turnaround analysis to be performed. The most complete version of the reconstruction is run on the grid processing nodes once chunks of data (of about a week in length) have been taken and identified.

3.2 Pulse fitting

The reconstruction of calorimeter data begins with pulse fitting. The method used for the analyses shown in this thesis fits individual digitized pulses to a characteristic shape (template) that is allowed to slide in time and stretch in amplitude in order to optimally fit. The pulse integral scales with the number of photons arriving at the SiPM. Studies have been done to show that templates built from different energies of pulses or at different times in the fill are identical. However, the templates do slowly vary on long time scales. Additionally, the difference in light collection between a laser pulse and a pulse arising from Čerenkov light requires separate templates for the two scenarios. This adds a minor complication of identifying when a laser pulse has been found in a calorimeter.

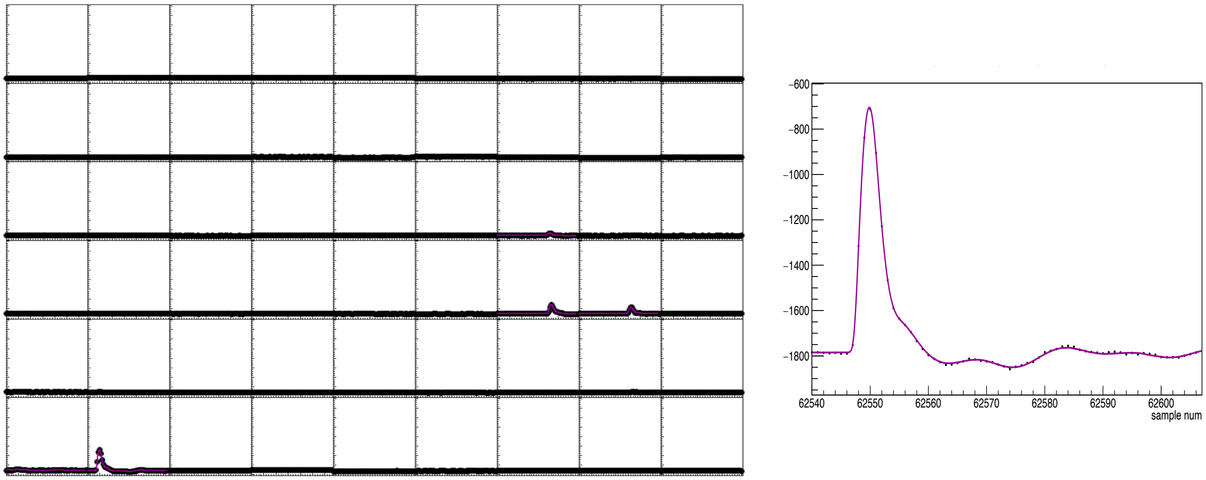

(b) Single channel

(a) Full calorimeter readout

Figure 3.1: Digitized readout of a full calorimeter (left) and a single crystal (right). Two positron hits can clearly be seen in a full calorimeter's digitized pulses at left. The large pulse, roughly 1 GeV, from the lower left of the calorimeter is shown at right with its template fit overlaid. The samples are separated by 1.25 ns.

Despite the short windows of time being saved under normal conditions, the characteristic pulse shape for each channel extends much further out to properly account for baseline perturbations on trailing pulses, should there be any. To that point, taking data to build the templates requires configuring the data acquisition to record wider windows. For this work, two sets of data were taken. The first set of runs were taken with 20 pre-samples and 250 post-samples. Due to the baseline calculation necessary for the template building process, data with 75 pre-samples was taken in a second set of runs. To reduce noise and keep the data rate manageable, the island chopping threshold is raised above the characteristic lost muon peak (to be explained later). The templates described here will be 200 samples long, with 18 of those being presamples.

Once the traces are collected, the template building procedure begins. In each digitizer, the two interleaved ADCs can in principle have different baseline values. See Figure 3.2 for a visible example of this effect. If this effect is not corrected, the characteristic pulse shape will be biased and ultimately impact the pulse fitting results. To calculate the effect, the pre-samples (samples before the peak of a pulse) are taken from the digitized traces. The clock and control system works in such a way that the even numbered samples from a digitized trace will always be from the same ADC, making this work conceptually easier. The "quiet" sample values are histogrammed for each of the even and odd sample numbers individually, as well as the appropriately signed differences between consecutive values. This allows for a determination of the difference between the two ADCs; the difference in a single digitized pulse can be seen in Figure 3.2. The values have been found to be stable over long time periods, but this procedure is still performed at the beginning of each data-taking period. Digitizers including channels with particularly large even/odd pedestal differences were swapped between Run $-$ 1 and Run $-$ 2 for spares with more favorable behavior.

For each collected pulse, the even/odd digitizer offset is corrected. Before building an average pulse shape, each individual pulse's baseline is calculated based on a 60-sample window ending six samples before the peak sample. For the traces with fewer pre-samples, the number of points used for the baseline calculation is 12. A baseline-corrected integral is taken ±2 samples around the peak sample to determine whether to include the pulse. Larger pulses are chosen as they are generally cleaner. The integral cut for this study was 2800 ADC·ct and reduced to 2400 ADC·ct for channels with particularly low statistics. The size of this window (5 samples) was chosen to roughly correspond to the period of intrusive 160 MHz electronics noise; requiring a large integral over this range reduces the chance that a pure noise pulse will be included.

Once a pulse has passed the quality checks, the calculated baseline is subtracted from each

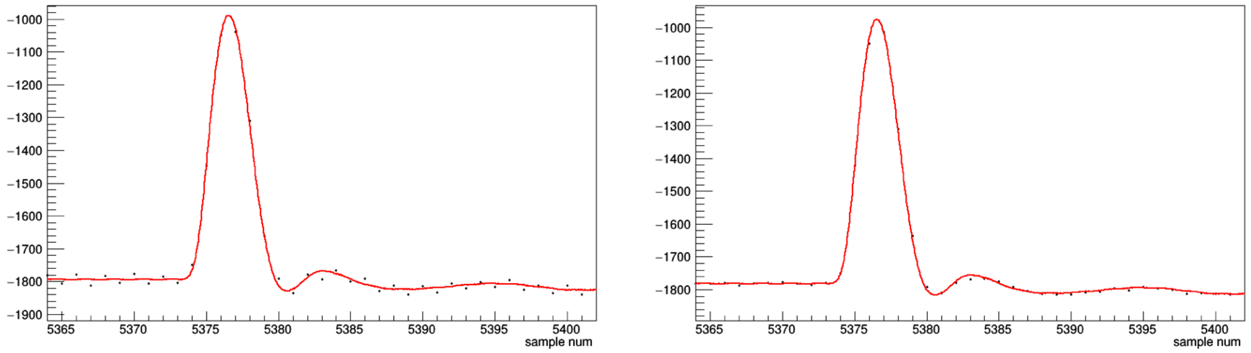

Figure 3.2: Digitized pulse before (left) and after (right) even/odd pedestal correction. The two interleaved ADCs in each digitizer channel can, in principle, have different baseline values. Shown here is the same trace and associated pulse fit before (left) and after (right) correcting for the even/odd ADC effect. This correction is especially important when constructing the average pulse shape.

sample, and the pulse integral is normalized. The pulse's "pseudotime" is then calculated. This is a measure of where the true time position of a trace should be, based on the relative heights of the maximum sample and the samples on either side. Pseudotime t_p is defined as

$$t_p = \frac{2}{\pi} \arctan\left(\frac{s_0 - s_-}{s_0 - s_+}\right)$$

where s_0 is the peak sample and $s_\pm$ are the samples directly preceding and following. Notice that when the "true" peak is as far to the right from the peak sample, $s_0 \to s_+$ and $t_p \to 1$. Alternatively, when the true peak is far to the left, $s_0 \to s_1$ and $t_p \to 0$. To get a useful transformation for our pulses, the pseudotime is mapped onto another quantity called "realtime":

$$t_r\left(t_p\right) = \int_0^{t_p} \rho\left(t'_p\right) dt'_p$$

where $\rho\left(t_p\right)$ is the normalized probability distribution function (PDF) of pseudotimes. Defining realtime in this way allows centering of the true peak of a pulse by placing its peak sample at $t = 0.5 - t_r$.

Once a pulse is peak-aligned and normalized, each sample is added to a 2D-histogram

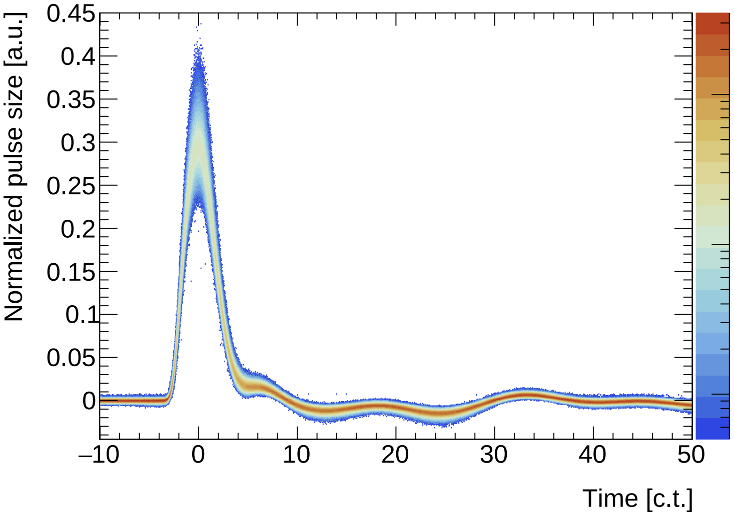

Figure 3.3: Overlaid SiPM pulses to extract fit template. Thousands of pulses are normalized and overlaid to create an average pulse shape for each channel. A template is created both for laser pulses and for pulses arising from charged particles impacting the crystals. The average pulse shapes are then used to fit digitized signals.

called the “fuzzy template.” An example is shown in Figure 3.3. A spline is taken of the y-profile of this histogram and then smoothed. The template in a ROOT file for use in the $g-2$ software.

The whole process is repeated using a dataset consisting of laser pulses to form the average pulse shape for the laser pulses in that channel. Additionally, special studies require a second laser template from the neighboring laser from which light is redirected. A comparison of the three average pulse shapes can be seen in Figure 3.4.

Taking the template as a function $T(t)$, pulse fitting characterizes the digitized pulses to a function $f(t) = E * T(t - t_0) + P$. E is the fitted pulse integral, since the template is normalized, t_0 represents the peak time of the pulse, and P is the digitizer pedestal. When multiple pulses are simultaneously fit, a shared pedestal is used. A χ^2 minimization technique is employed that uses a modified Newton’s method to find the optimal placement

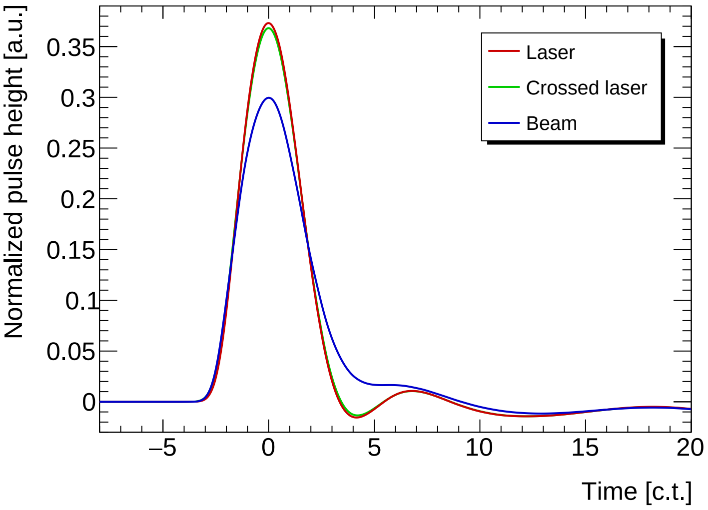

Figure 3.4: Comparison of the three pulse templates for a single channel. For each channel, three templates are created for fitting purposes. Incident charged particles produce Čerenkov light that arrives at the SiPM with a larger temporal spread than the injected laser pulses, so the normalized pulse shape is wider. The "crossed" laser refers to light redirected from a neighboring laser head, used to study the calorimeter response to impulses.

and scaling of the templates to match a digitized pulse. For more detailed information on the pulse fitting procedure, see Aaron Fienberg's thesis [43].

To identify laser pulses within a calorimeter, the reconstruction software scans the chopped island for the number of channels that have a pulse over threshold. Since not every calorimeter has all 54 channels live throughout Run $-$ 2, the number of requisite channels to be classified as a laser pulse is variable, though typically 50 above-threshold pulses are required in an island to be classified a laser pulse. For example: during a week of data-taking, two digitizers in one calorimeter were out of sync with the other nine. The traces from those nine channels ended up temporally separated from the rest of the detector, even for normally-safe situations like laser pulses. Given the small impact and short time of this unfortunate event, the data from those nine SiPMs was discarded for that time period.

When pulses are close enough together that their templates overlap, the fit is re-done on multiple pulses as a “chain fit” to properly account for the tails of previous hits. If this situation arises while fitting a laser pulse in a channel (as identified in the manner described above), the fitter will try using the laser template for each identified pulse present. The situation with the best fitted χ^2 value selects which fits to report. The muon fills during which lasers are fired to measure the SiPM gains are currently only used to construction a correction function, not as additional data in the ω_a fits. This is the primary situation in which laser and muon/positron hits within a SiPM would need to be reconstructed within a single fit. As the laser data is the primary goal of such fills, the accuracy of the laser pulse fits is very important.

Some odd situations arise when fitting large amounts of data that must be incorporated into the reconstruction chain. One main source of trouble for the fitter is the presence of electronics noise at 161 MHz. Attempts were made to reduce the impact of the noise on the hardware side by disconnecting a redundant ground from the signal cable, but the noise persisted. Unfortunately this frequency matches very well the pulse width of real physics events, so the noise passes the frequency filters on the SiPM board. This confuses the fitter that tries to fit a chain of pulses into the peaks of the sinusoidal signal. By using known properties of how the fitter treats such noise, the pulses can be flagged and later removed in the data-processing.

3.3 Gain matching

Calibration of individual channels within a calorimeter is done by equalizing the energy peak of positive muons that pass through the lead-fluoride while depositing a minimum amount of energy. This peak lies at an effective energy of 170 MeV, as will be discussed later in the section on energy calibration. In order to efficiently chop these muon-only islands, the gain of the muon-struck channel needs to be high enough that the peak is well-defined above the island chopping threshold. Conversely, the gain of a particular channel should be low enough that a channel’s signal should remain within the dynamic range of the digitizer when struck by the highest energy decay positrons or, occasionally, multiple positrons. The endpoint of the positron spectrum is at the energy of stored muons, or 3.1 GeV. Assuming a pedestal of -1800 and a maximum readout of 2048, each digitizer has a dynamic range of 3848 ADC counts.

It is necessary to coarsely equalize the gain of individual channels at the hardware level, but more precise energy scale alignment will happen in the energy calibration step. Each channel has two knobs that can be used to tune the gain: the bias voltage and the programmable gain amplifier. SiPMs have a characteristic breakdown voltage; applying a bias voltage higher than the SiPM's breakdown voltage allows operation in Geiger mode. The SiPM's gain scales linearly with overvoltage, or bias voltage minus breakdown voltage [60]. Four BK Precision power supplies in each calorimeter supply the necessary voltage (≈68 V) to the SiPMs. Care was taken during calorimeter construction to group SiPMs with similar breakdown voltages onto the same power supplies. The programmable gain amplifier on each SiPM electronics chip has a range of 2-20× signal amplification. The programmable gain setting ranges from 0-80. The signal amplification r as a function of this setting s is given by

$$r(s) = 10^{\frac{1}{20}(26-s/4)}.$$

This function can be inverted to find the appropriate setting for a desired amplification. Adjusting the output voltage of one power supply affects up to 16 SiPMs, but the programmable gain amplifier can be tuned for each channel.

Similar gain equalization work to that presented here was performed prior to $\text{Run}-1$. The need to re-do the work before $\text{Run}-2$ was threefold: one calorimeter was completely rebuilt after an incident that left the detector water damaged, temperature stability of the experimental hall was found to be less than optimal, and single-frequency electronic noise was pervasive in the data. The rebuilt calorimeter drew from a combination of crystal/SiPM pairs that had been used in $\text{Run}-1$ and those that hadn't. Those pairs that had been used in $\text{Run}-1$ were re-arranged and grouped on the bias voltage supplies differently, thus requiring a new equalization campaign.

To mitigate temperature effects and single frequency noise, the bias voltage of all SiPMs was increased by about 0.2 V. A temperature drift of 1° C corresponds to a change in SiPM breakdown voltage of ≈0.07 V according to lab tests. If we assume the hall temperature could change by ±2° C and the overvoltage is 0.9 V (close to the Run − 1 set points), this is a gain variation of ±15.6 % from nominal. By raising the applied voltage and therefore the overvoltage by 0.2 V, the same temperature variation results in a gain change of ±12.7 % from nominal. Of course the laser system provides a way to adequately track these gain changes, but reducing the magnitude at the hardware level is preferable.

Now, moving on to the single-frequency electronics noise. The 161 MHz noise in question has been shown multiple times in the context of attempted hardware and software solutions. The noise in multiple channels is large enough to trigger island chopping, and that's the main limitation on lowering the island chopping threshold. Once the noise pulses are in the data, they manifest as many small, consecutive pulses. The frequency is such that each peak is about the same width as the templates used to fit – this is why the noise passes the filters in the SiPMs' electronics chain. Originally the noise was assumed to enter through the shielding of the signal cable for each channel. However, the noise persisted after disconnecting the shielding from the SiPM signal chain. Since then, it has been found that the noise passes through the programmable amplifier of the SiPM through an as-yet unknown mechanism. Lowering the programmable gain of a channel can reduce the size of the noise pulses to be below the island chopping threshold. By raising the bias voltage and lowering the programmable gain of all channels, noise can be reduced while still balancing the effects described in the motivation for gain equalization.

The gain of each calorimeter channel is calculated by using series of filters with varying transmittance (arranged on a remotely adjustable wheel) attenuate the amount of light arriving at the SiPM. The number of interest is pulse integral per number of photoelectrons incident on a SiPM (= pixels fired). By changing the filters in the path of the laser light, the number of photons arriving at a particular SiPM changes in a controlled fashion. The statistics explaining why this works is included in Appendix B.

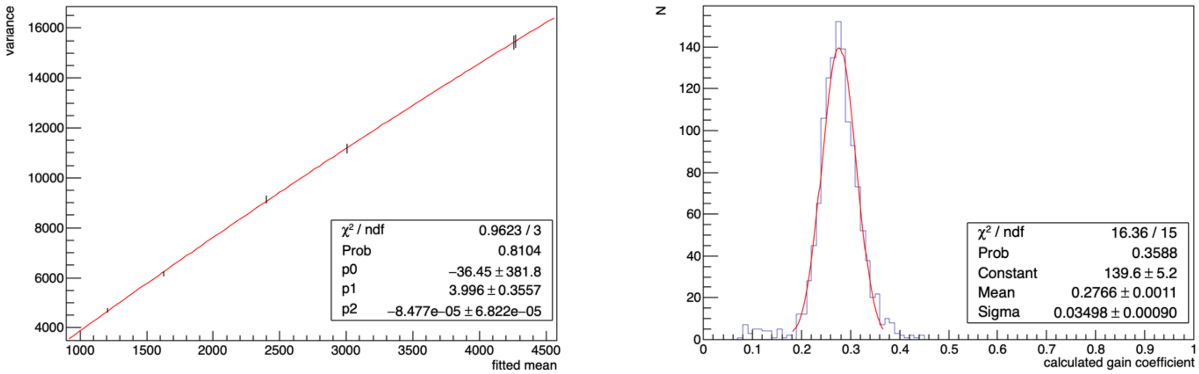

(a) Laser filter wheel scan

(b) Laser filter wheel scan

Figure 3.5: Example variance of pulse integral values vs. mean for a filter wheel scan. At left, the amount of laser light distributed to the calorimeters is varied to determine the gain of a SiPM, as described in the text. The gains of each channel are extracted as the coefficient of the first-order term of a second degree polynomial fit, as shown. The gain coefficients (1 / gain) of the 1300 channels after alignment are shown at right; a Gaussian is fit to the distribution. The values in the tails of the distribution reached hardware limits on the on-board amplifier before full alignment was achieved.

The filter wheel scans were performed using the sequencer function of MIDAS, the data acquisition system used. The filter wheels for each of the six lasers, providing light to four calorimeters each, were changed simultaneously. This allowed for gain calculations of all 1296 channels using the same set of data. Data was taken at filter wheel positions 6, 7, 8, 9, 10, 11, and back to 6. As a reminder, a higher filter wheel position number corresponds to less attenuation of the light. As described above, it is not important to know the relative transmittance of the filters. The two points at position 6 were to verify the drift across the time of the scan, roughly 20 minutes, was reasonably small. The data was taken in this fashion: four laser pulses per fill, 12 fills per second, and 120 seconds at each filter wheel setting. For each channel, the resulting pulse integral values were fit to a Gaussian. The mean and variance were extracted and plotted. Then, each channel's variance vs. mean graph was fit to a second order polynomial in the range $\mu \in (100, 10000)$ where μ is the mean pulse integral from a particular filter wheel setting. This fit range was chosen to avoid fitter problems at the low end and saturation effects at the high end. See Figure 3.5a for an example of the variance vs. mean curve for one channel. The coefficient of the first-order term is identified as the gain of that channel. Some plots shown will involve the "gain coefficient" or 1 / gain. This puts the constant in terms of photoelectrons / pulse integral. It also allows for convenient plotting – all gain coefficients will be in the range (0,1).

Following the extraction of the gain coefficients for each channel, new programmable gain settings needed to be calculated. For Run $-$ 2, a gain coefficient of 0.28 was chosen as the target value [62]:

$$\frac{0.28\text{ photoelectrons}}{1\text{ pulse integral}}\,\frac{1\text{ pulse integral}}{0.35\text{ ADC}}\,\frac{1\text{ MeV}}{0.8\text{ photoelectrons}} = 1\frac{\text{MeV}}{\text{ADC}}. \tag{3.1}$$

Following this, the muon peak of 170 MeV should lie far enough above the island chopping threshold of 60 ADC to be useful, and the dynamic range of the digitizers should accommodate large decay positron pulses. A few iterations of filter wheel scan → extract gain coefficients → calculate new programmable gain settings → filter wheel scan were necessary to adequately equalize the SiPM gains. Upon finding the current gain coefficient c_g of a SiPM, its new gain setting s_{new} can be calculated based on its old setting s_{old} and the goal gain coefficient $c_{\text{goal}} = 0.28$ through:

$$s_{\text{new}}\left(s_{\text{old}}, c_g, c_{\text{goal}}\right) = s\left(\frac{c_g}{c_{\text{goal}}} \times r\left(s_{\text{old}}\right)\right). \tag{3.2}$$

This uses the equations for the amplification as a function of gain setting $r(s)$ and the gain setting as a function of the amplification $s(r)$:

$$r(s) = 10^{\frac{1}{20}(26 - s/4)} \tag{3.3}$$

$$s(r) = 104 - 80\log_{10} r. \tag{3.4}$$

The gain settings were limited to be $s \geq 20$ to avoid clipping pulses. Available hardware settings further constrains $s \leq 80$. After several iterations, the final distribution of gain coefficients can be seen in Figure 3.5b. The SiPMs in the tails on either side are those whose "desired" settings are outside the limits mentioned.

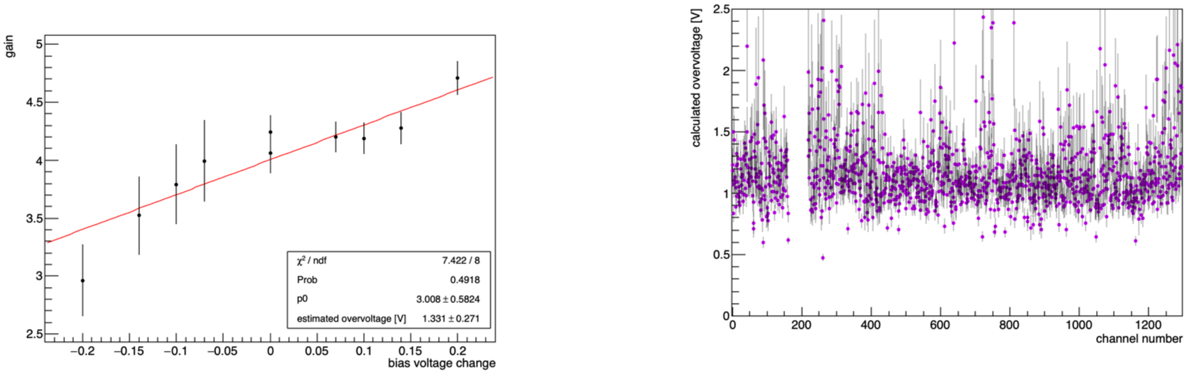

(a) Gain vs. voltage change

(b) Approximate overvoltage for most channels

Figure 3.6: SiPM gain as the applied bias voltage is changed (left) and the approximate overvoltage for the 1300 SiPM channels (right). By adjusting the bias voltage set-point of a SiPM and measuring the gain at different settings, the approximate overvoltage at nominal set-point can be estimated by fitting to a line and extrapolating to the voltage where the gain would drop to 0: $G(V) = p0(\Delta\mathrm{BV} - \mathrm{OV}_{\mathrm{est}})$. This was done for every channel prior to $\mathrm{Run} - 2$, except for a calorimeter that was unusable at the time of the study.

This process of measuring the gain of a channel (or set of channels) using the laser filter wheel can be repeated at different bias voltages. By plotting the gain vs. bias voltage adjustment, the overvoltage of each channel can be approximated by determining at what voltage the gain goes to 0, as seen in Figure 3.6. Knowing the approximate overvoltage for each channel allows estimations to be made about the level of gain fluctuations to expect. As the operating voltage is held constant but the breakdown voltage changes due to temperature fluctuations, the change in gain due to temperature drifts can be approximated from the overvoltage at a known time/temperature. For example, a typical measured overvoltage in

Figure 3.6 might be 1.1 V. A ±1°C change would change the breakdown voltage by ±0.07 (per Figure 2.12), resulting in a gain change of $\mp 0.07/1.1 = \mp 6\%$.

3.4 Timing alignment

To simultaneously use data from all detectors when extracting ω_a, the detectors must have a common reference point for time. The T0 detector, true to its name, provides this reference point. The intensity-weighted average time of the injected beam (as seen by the T0) is taken as the time of injection. The corrected pulse times in a particular channel are given by

$$t' = t - t_{\text{sync}} - \left(t_{\text{T0, beam}} - t_{\text{T0, sync}}\right) - \delta_{\text{SiPM}} - \delta_{\text{calo}}, \tag{3.5}$$

where δ_{SiPM} and δ_{calo} are constant offsets specific to that channel and its calorimeter related to laser fiber length differences and arbitrary electronics delays.

3.4.1 Intra-calorimeter

The electromagnetic showers of decay positrons provide a method of aligning neighboring channels in time. This alignment accounts for fiber length differences leading to the laser distribution panel on the front face of the calorimeter. The time difference between crystals within the same cluster are plotted if each has an energy of over 200 MeV. A system of equations is constructed similar to that described for the lost muon events. In particular, for each set of adjacent (vertically and horizontally, not diagonally) crystals, $\Delta_{ij} = \delta_i - \delta_j$. This assumes that the time difference should be 0. Constructing that equation for each pair of crystals yields 93 equations. The system is overdetermined, so it becomes a least squares problem. Each equation is weighted by the inverse of ϵ_{ij}, the fitted error on the mean of Δ_{ij}:

$$\Delta_{ij}/\epsilon_{ij} = \left(\delta_i - \delta_j\right)/\epsilon_{ij}.$$

One crystal's offset is set to zero as a reference. The constants are produced with the condition that $\delta_0 = 0$. The system was solved using a singular value decomposition in the

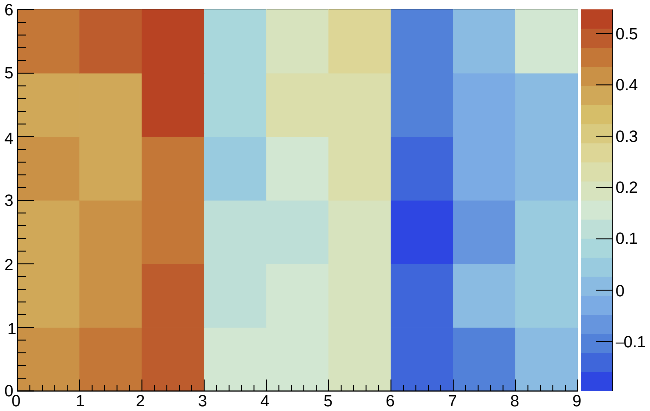

Figure 3.7: Calculated timing offsets within a single calorimeter. To account for the different laser fiber lengths and electronics delays, a series of equations is solved to set the time difference between pulses in neighboring crystals from a common source to be 0. In this image, the beam is on the left-hand side going into the page. Each bin represents a crystal, and the z-value is that crystal's calculated time-correction. The 10 cm differences between the 3 fiber lengths used in the construction of the laser panel are immediately recognizable as groupings approximately 0.3 ct apart.

linear algebra software package Eigen and crosschecked with other software. For a consistency check, the crystal designated to be $\delta_i = 0$ can be changed and the difference between two arbitrary offsets can be compared. It was found that this changes at the level of 10^{-6} ct. To put it simpler, the calculated offsets are completely consistent no matter which crystal is fixed. An example of the offsets found through this process can be seen in Figure 3.7. The pattern seen is reflective of the three different lengths of laser fiber in the distribution panel, as seen in Figure 2.16.

3.4.2 Inter-calorimeter

Lost muon events provide a sharp peak for the energy scale of individual crystals and also provide a way to align separate calorimeters in time. Figure 3.8 shows a sample trajectory of a muon lost before decaying and the resultant energy and time-of-flight peaks. A destructive beam monitoring detector system was inserted into the storage ring at two locations for a set of data where more muons are scattered than in nominal data.

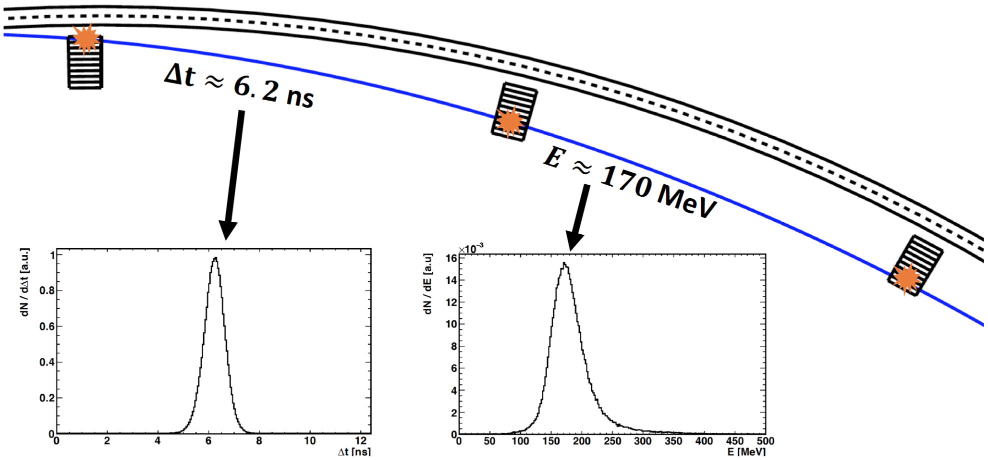

Figure 3.8: Example trajectory of a muon lost from the storage region before decaying. When a muon is lost from the storage ring before decaying, it can travel through the active volume of several calorimeters. The well-defined time-of-flight and energy peaks from these events are useful for aligning and calibrating the channels. Image reproduced from [43].

To begin this study, coincidences were found in three consecutive calorimeters. Timing alignment was done at the level of aligning all sync pulses to one another. Very loose cuts were applied to the energy of the clusters and time separation of successive hits. Additionally, noise cuts in the fitting procedure were loosened as they had not yet been optimized. The time separation Δ_{ji} vs. calorimeter j can be seen in Figure 3.9a. The large jumps every four calorimeters correspond to different laser heads, controlled independently. The six lasers are ostensibly fired simultaneously, but laser-dependent time differences could exist in the triggering sequence, light distribution path, or elsewhere.

Assume that each calorimeter i reports a time $t'_i = t_i + \delta_i$, where t is the true hit time and δ_i is a constant offset combining factors such as cable delays and laser light propagation time. Now assume that the travel time α for a muon traveling from one calorimeter to the next is constant for all pairs of consecutive calorimeters. This leads to a system of equations of the form:

$$\begin{aligned} \Delta_{ji} &= t'_j - t'_i && (3.6) \\ &= t_j + \delta_j - t_i - \delta_i && (3.7) \\ &= \delta_j - \delta_i + \alpha. && (3.8) \end{aligned}$$

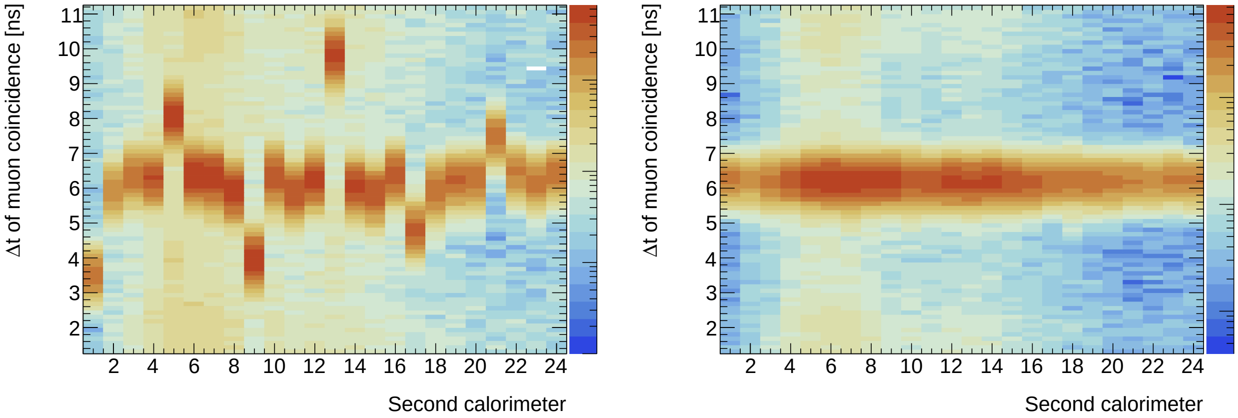

(a) Before correction.

(b) After correction.

Figure 3.9: Measured time-of-flight before (left) and after (right) calorimeter timing alignment. Triple coincidences from consecutive calorimeters are used to identify muons lost from the storage region. The left panel shows the measured time-of-flight prior to aligning the twenty-four calorimeters. The large jumps seen are where the two calorimeters are fed laser light from different laser heads. The plot on the right shows the agreement in the time-of-flight after alignment at the expected value of 6.25 ns.

By setting the offset $\delta_{22} = 0$, the system can be solved exactly. Calorimeter 22 is chosen because that electronics crate also digitizes the signal from the T0 entrance counter. Using this method prevents introducing an error based on a chosen value for the time-of-flight between calorimeters and offers a straightforward sanity check. When this system is solved using the data shown in Figure 3.9a, $\alpha = 6.25\,\text{ns}$. This is consistent the value of $T_c/24 = 149.2\,\text{ns}/24$. After applying the constants, the situation becomes that seen in Figure 3.9b.

3.5 Energy calibration

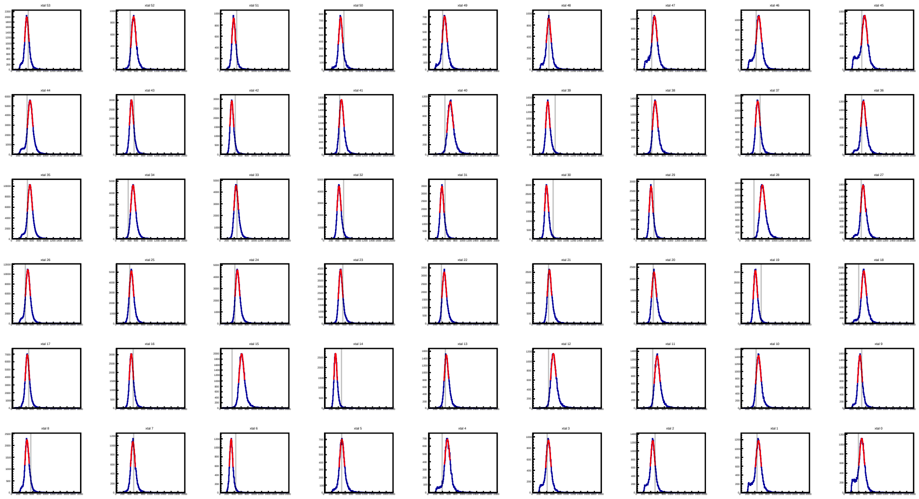

Figure 3.10: Histograms of fitted pulse integral values from lost muon events for each SiPM in a calorimeter. The 54 channels of a calorimeter are shown with the pulse integral from identified lost muon events, with the beam on the left-hand side traveling into the page. A lost muon acts as a minimum ionizing particle (MIP) in each channel; a Gaussian fit to the peak of this pulse integral distribution from each crystal/SiPM pair is used to align the energy scales of the nearly 1300 channels in the calorimeter system.

With a tighter coincidence time window in use, the pulse integral for clusters including only one crystal for the second two hits in a triple coincidence is plotted for each channel. The reason for only using the second two hits is to reduce possible background events that arise more frequently on edges of a coincidence. Normally this would mean only using the second event in a triple coincidence; however, the second hit in a lost muon coincidence very rarely hits the top and bottom rows, and furthest radially inward column of crystals in a calorimeter. Using only the second and third hits in triplets allows for a balance of statistics and background. Once plotted, the peak was fit to a Gaussian in the 1.5σ range around the peak. See Figure 3.10.

The calibration constant for each channel is defined as $170\,\mathrm{MeV}$ / fitted peak (pulse integral). The value for 170 MeV could be anything at this point in the process – it is just used to align all channels.[1] The energy spectrum after applying the calculated calibration constants can be seen compared to the uncalibrated spectrum in Figure 3.11.

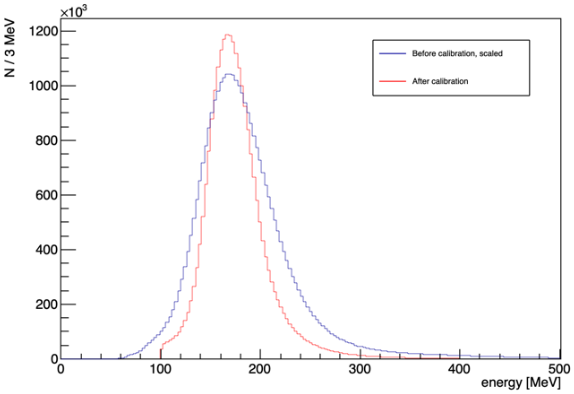

Figure 3.11: Lost muon energy spectrum before and after calibration. The peak of the fitted pulse energies is scaled appropriately to compare the energy spectra of identified lost muons before and after calibration. The peak sharpens, as expected.

---

[1] Using information from the Particle Data Group, the estimated energy loss of a muon with momentum $p = 3.1\,\mathrm{GeV} / \mathrm{c}$ passing through one of the calorimeter's crystals would be [51]:

$$1.551\,\mathrm{MeV\,cm}^2 / \mathrm{g} \times 7.77\,\mathrm{g} / \mathrm{cm}^3 \times 14\,\mathrm{cm} = 168.7\,\mathrm{MeV}.$$

See *Muons in lead fluoride (PbF$_2$)*. This calculation does not take into account the production of Čerenkov light in the crystals, particularly when relating the light produced from a positron shower to that produced from muon ionization. The close match should be viewed as a coincidence, not a concrete reason to choose this value for the lost muon peak.

3.6 *Gain corrections*

The energy scales of each SiPM can drift independently. Should they do this in a coherent way that couples to a early-to-late effect (for example: a gradient of relative gains across the face of the calorimeter that couples to a time-dependent hit position), the fitted value of ω_a could be pulled from its true value. The corrected gains of the nearly 1300 channels must be stable to better than 10^{-3} to meet the requirements laid out in the TDR [52].

3.6.1 Long term [hours]

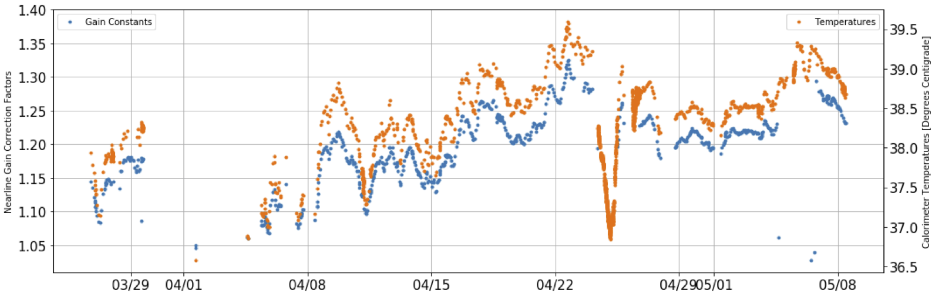

Figure 3.12: Average gain correction and temperature across weeks for a single calorimeter. The gain of the SiPMs changes with temperature. The average temperature of the SiPMs in a calorimeter is plotted with the necessary gain correction found by comparing the size of a laser pulse in each channel to the size at an earlier reference time. The temperature and correction factor track each other, as expected. Image credit: Josh Labounty.

As discussed in Chapter 2, the gains of the SiPMs are highly temperature-dependent. The temperature stability in the experimental hall during $\mathrm{Run}-1$ and $\mathrm{Run}-2$ was sub-optimal, leading to large drifts in the uncorrected energy that the calorimeters would report. Long term energy-scale drifts do not couple directly to the measurement of ω_a, but the power of the fit is reduced when the energy scale is smeared. As such, the gain is tracked by firing a simultaneous laser pulse into every one of the calorimeter channels four times between each beam injection. The ability of the calculated gain correction to track temperature drifts can be seen in Figure 3.12.

The reconstruction of the between-fill laser pulses is a simplified version of the pulse fitting described above, utilizing only the laser template. The size of the laser pulse in each channel is compared to the average value over the course of the calibration run to pin the energy scale appropriately. Independent stability monitors provide a measure of how the laser power compares to itself at the reference time.

3.6.2 In-fill gain recovery [μs]

Quite unlike the effect described in the section above, this effect does couple directly to ω_a. Since $d\phi(E_T)/dE_T$ is non-zero, as discussed in Chapter 1, a changing energy scale throughout the measurement period directly pulls the fitted value of ω_a.

In normal muon fills, the laser is fired twice: once at the beginning of the fill before injection and once at the end of the fill after almost all muons would have decayed. The "sync" pulse at the beginning of the fill is used as a reference point to align the times of all calorimeter channels. Between the muon beam injections, four laser pulses are fired to track the gain of the nearly 1300 channels across long time periods and correct for any drifts. Once every 11 fills (a number chosen to be intentionally relatively prime to the number of bunches), four laser pulses are fired 185 μs apart during the normal ω_a fit period. These fills are not used in the extraction of ω_a, but the laser pulses are used to track any insidious gain effects present during the fit time and provide a correction function. Each time the so-called "in-fill" laser pulses are shot, they are shifted in time 2.5 μs across a range of 230 μs. There is intentional overlap between where the first and second in-fill pulses will be fired as a consistency check. After 93×11 muon fills, the pattern repeats.

After a SiPM discharges, the voltage across drops as the charge is recovered. When a large set of pulses almost completely depletes a SiPM, the recovery to full bias voltage is driven by how fast charge can be returned to the device. This happens most egregiously in the channels closest to the inflector, where beam splashes off the yoke when the beam is injected into the ring. The difference between the calorimeters close to and away from the

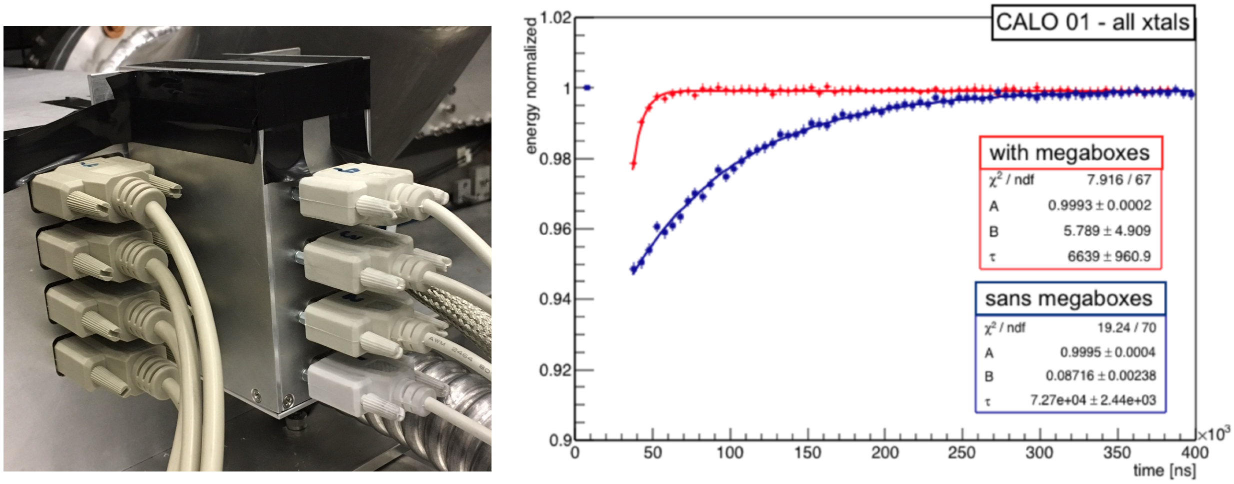

(a) Hardware modification.

(b) IFG before/after correction.

Figure 3.13: Hardware modification (left) to improve the in-fill behavior of the SiPMs (right). When the beam is injected into the storage ring, almost 99% of particles are not stored. This splash of particles creates a blast of light in the crystals, depleting the SiPMs. Recovery is driven mostly by the on-board capacitors. Shown at right is the gain drop and recovery following injection with and without the hardware modification (left). Both the amplitude of the drop and recovery time improved. The time shown is relative to the start of digitization; injection occurs at approximately 30 000 ns, and fits to extract ω_a begin at approximately 60 000 ns. The fit function employed is $G(t) = A - B\exp(-t/\tau)$, where τ is shown in units of nanoseconds. Plot reproduced from [64].

beam injection site can be seen in Figure 3.14. Capacitors on the SiPMs' electronics boards provide this most quickly, but as was discovered early in the commissioning of the detectors, the small on-board capacitors do not provide enough charge for stable gain throughout the measurement period. Lab tests were performed on how to mitigate this, and an additional set of capacitors, packaged in an easily installable standalone housing, were added to each bias voltage distribution circuit. This improvement can be seen in Figure 3.13.

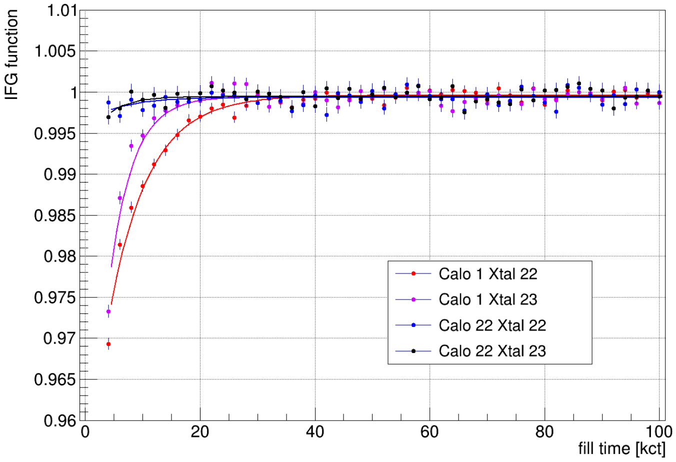

Figure 3.14: In-fill gain recovery functions for four crystals. The same two crystals are shown for a calorimeter heavily impacted by the injection splash (calorimeter 1; red and pink curves) and one barely affected (calorimeter 22; blue and black curves). Also shown is the difference in recovery time depending on the two circuits of each breakout board: red has a fitted recovery time of 8.6 μs whereas pink has a recovery time of 5.1 μs. Typical fits to extract ω_a begin at 24 kct. Image credit: Paolo Girotti.

3.6.3 Short term [ns]

A nanosecond-level gain effect enters into the ω_a measurement as a rate-dependent effect. Since the effect only occurs in pileup-type events, the occurrence rate changes throughout the measurement period. The effect is driven by the fact that SiPM pixel response is binary: either a pixel fires or it does not. Once it has fired, a pixel must recover charge in order to re-fire. The statistics of the number of pixels available once a number of them have fired drives an effective gain drop in the nanoseconds following a pulse.

In special campaigns during each data-taking run, the laser system is used to determine this nanosecond level response of a SiPM to repeated pulses. In each laser head's optical path is an optional mirror that can be lowered such that the light is redirected into its complementary laser head's optical path. That means that the laser pulses from two laser heads can be fired into four calorimeters at once while the other four calorimeters receive no laser light. The delay and relative light levels between the sets of pulses can be adjusted to emulate the effect of two positrons of varying energies hitting the same channel within nanoseconds. Starting in Run $-$ 2, an abbreviated version of this study was performed during each periodic trolley run to track this effect as the temperature in the experimental hall changed. That became less important in Run $-$ 3, as the ambient temperature was almost completely stabilized. The resulting correction function for one crystal, measured at one set of light intensities, is shown in Figure 3.15.

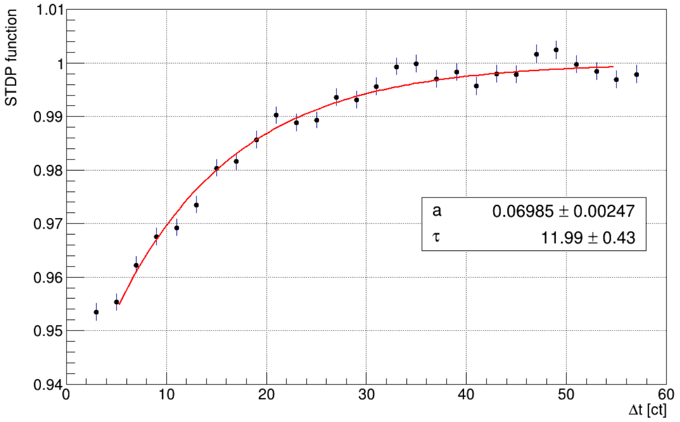

Figure 3.15: SiPM gain recovery at the nanosecond level. Following a large pulse, the relative gain of a SiPM is lowered relative to the number of fired pixels. Using the laser system, two pulses can be fired in quick succession to measure this effect and correct for it in data. The fit shown has the form $G(t) = 1 - a\exp(-t/\tau)$, where τ is measured in clock ticks. Image credit: Paolo Girotti.

A naïve model to understand the short-term correction is that the number of pixels that will fire for some incident amount of light will scale linearly with the number of available pixels. This approximation works only in the case of $N_{\text{total}} \gg N_{\text{fired}}$. The SiPMs employed have 57 300 pixels, so the approximation should hold relatively well for 3 GeV positrons (≈ 2400 pixels fired), though the model can be expanded to include saturation effects. That would mean the relative gain at time t after a pulse would be

$$G(t) = 1 - \frac{N_{\text{fired}}}{N_{\text{total}}} \exp(-t/\tau_{\text{pixel}}), \tag{3.9}$$

where τ_{pixel} is the pixel recovery time.

Confusingly, the relative gain drop measured by the laser system did not align with the *a priori* expectation given by the model above. There are no indications that the laser system is providing an incorrect measurement, so a few tests were performed on individual SiPMs to better understand the effect. In these, a single laser pulse was split into two beams, one of which was delayed, that were then diffused across the SiPM face. At the most basic level, this would verify the relationship between number of pixels fired and relative gain drop as seen in the true experimental setup. A second goal was testing the premise that the short-term effect is driven primarily by pixel saturation effects. To that end, tests were performed where the two split beams were shone on distinct parts of the SiPM face through use of a physical barrier. The orientation of the SiPM was turned 90° for an additional set of tests, to verify the summing logic on the SiPM electronics board was not contributing.

Some results from this study can be seen in Figure 3.16. The size of the gain drop in the lab tests was comparable to that in the experimental setup. The linearity with number of pixels fired in the first pulse is evident, giving credence to the model in Equation 3.9. The relative gain drop is quite extreme, which could be explained by having a smaller number of total pixels N_{total}.

It is clear that when the two laser pulses are kept separate on the face of the SiPM, the drop in relative gain for the second pulse is negligible compared to when the two pulses share the full face. Additionally, no effect due to the on-board summing logic was noticeable. No resolution was reached on why the gain drop at short times is larger than was expected, but an additional layer of confidence was added to the *in situ* laser measurements.

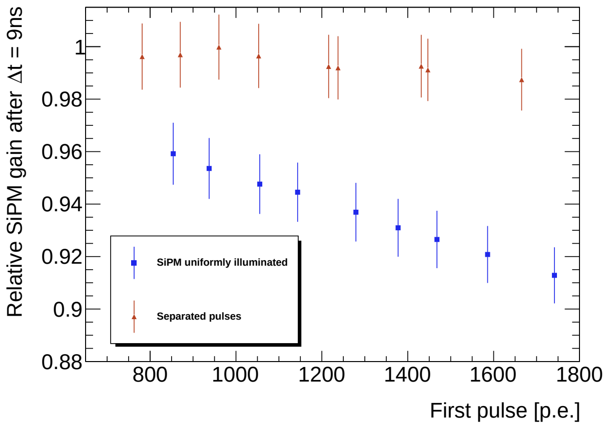

Figure 3.16: SiPM gain drop with physically distinct or overlapping laser pulses. Lab tests indicated that the nanosecond-level gain drop is due to pixel saturation in the SiPMs. A test was devised to split a laser beam into two pulses, one of which was delayed $\approx 9\,\text{ns}$ relative to the other. The two pulses could either both illuminate the full face of the SiPM or they could be directed to physically distinct parts of the same SiPM. The test as shown indicates that pixel saturation and recovery is the driving force in the short-term gain effect, not any electronics considerations.

3.7 Clustering

When a positron hits a calorimeter, the electromagnetic shower is shared between several crystals. For a particle with energy 2.5 GeV normally incident upon the center of a crystal, roughly 80% of the energy is contained within the crystal impacted. The rest of the energy is shared mostly by the surrounding crystals. Each channel will produce a pulse that is then fit and the proper corrections are applied. Once through that process, a clustering algorithm is used to combine separate pulses into a final physics object.

Several clustering algorithms have been explored, but the one used for this thesis relies entirely upon the temporal information of the crystal hits. The typical spread of fit times within a single positron hit is less than 0.5 ns. A deadtime of 2.5 ns is applied, so all hits within that window are combined. Their energies are summed, and the time of the event is the time of the largest (by energy) crystal hit within the cluster. A softer bound exists at 3.75 ns: any hit with $2.5\,\mathrm{ns} < \delta t \leq 3.75\,\mathrm{ns}$ will check before and after for the closer of the two clusters, as determined by energy-weighted time. Once clustered, a logarithmic weighting technique is used to determine position information of the impact. Using only temporal clustering allows for the pileup subtraction method that will be described in the next chapter. More information about the clustering routine can be found in [43].

3.8 Data flow

To recap how the data is handled:

- Data acquisition
    - Windows of data are saved whenever a channel exceeds threshold.
- Pulse fitting
    - Pulses are fit using an average pulse shape to extract a hit time and energy.
- Time corrections
    - The synchronization laser pulse is aligned in all channels.
    - The beam crossing the T0 detector is taken as the injection time, $t = 0$.
    - Using the known travel time of lost muons between calorimeters, electronics and laser delays between detector stations are accounted for.
    - By setting the time between pulses in adjacent channels from shared positron showers to 0, the laser fiber length and electronics delay differences within a calorimeter are removed.
- Energy calibration and gain corrections

  - A special set of runs designed to scatter muons from the storage region allows alignment of the MIP peak in every channel.
  - Tracking the size of laser pulses across the length of the run corrects for temperature drifts.
  - Laser studies correct for the gain drop due to the splash of unstored particles at injection.
  - Pulses within 100 ns after another pulse are corrected for SiPM pixel recovery effects.

- Clustering
  - Crystal hits in the same calorimeter within 3 ct are combined into one physics object with an associated time and energy.

Chapter 4

EXTRACTING ω_a

Fitting to find the anomalous precession frequency ω_a uses the energies and times of hits reported by the calorimeter system, following the data chain described in the previous chapter. In this chapter, the $\mathrm{Run}-2\mathrm{D}$ dataset to be analyzed will be described, including some discussion of the other data from $\mathrm{Run}-2$ that is not included. After correcting the data for pileup events, the simplest form of the fit used to find ω_a will be presented. Following that, corrections will be introduced to account for muons that are lost from the storage region before decaying and for acceptance changes as a result of beam motion. The optimal weighting scheme to maximize precision on the fitted ω_a is demonstrated. Some consistency checks will be shown, and a new analysis tool to mitigate an effect in per-calorimeter fits will be described. Some detector-related systematics will be presented along with an estimation of the impact on ω_a.

4.1 ***Run*** $-\,2$

Data-taking for $\mathrm{Run}-2$ of the E989 experiment took place in March - June 2019. Unlike the previous year's run, consistent electrostatic quadrupole voltages were used throughout the run. As discussed previously, the operating voltage of the quadrupole system impacts the frequencies at which the beam oscillates within the storage region. These oscillations play a large part in the fit for ω_a, so maintaining a consistent operating point allows, in theory, the full dataset to be fit simultaneously using the upcoming techniques. The magnetic kicker voltage, also highly impactful to the beam dynamics, was also held steady throughout the running period.

In practice, there have been some miscues within the production chain for the $\mathrm{Run}-2$ data that have rendered much of the data unfit at the time of writing. A challenge when fitting the full $\text{Run}-2$ dataset simultaneously will be the sporadic nature the data was collected. For budgetary reasons, the accelerator campus at Fermilab was put into a 5 days on / 9 days off mode during the spring of 2019. These disparate segments of data each has their own oddities that may prove complicating. Such effects could arise from the instability of the temperature within the experimental hall, meaning that equipment sensitive to temperature (notably the storage ring magnet and calorimeter SiPMs) may behave differently among the different datasets. The corrections for this effect normally rely on a smooth transition that may break down with the gaps in data. As of the time of this writing, the collaboration plans to combine the $\text{Run}-2$ and $\text{Run}-3$ data into one dataset.

4.1.1 Data quality

To ensure the quality of the data to be fit, the data undergoes a series of filters to select the best fills. Fills with a low $T0$ integral, a measure of the size of the incoming beam, are cut. When an electrostatic quadrupole or magnetic kicker plate sparks, a window of data is removed surrounding the failure and recovery. Cuts are in place to guarantee all SiPMs are receiving laser pulses for proper gain corrections. Finally, periods of time around instability in the measurement of the magnetic field are removed. When combined together, the data quality cuts removed $\approx 25\%$ of the data from the $\text{Run}-2\text{D}$ dataset.

4.2 Histogram construction

To start off, only the threshold-based fitting will be discussed. In this method, a lower threshold is placed on the positron hit energy. All hits above that energy are given a weight of 1. In general, different weightings (typically as a function of energy) can be applied to the hits in order to better leverage the information from high-energy decays better.

The data is binned in 20 MeV segments by $0.1492\,/\,6\,\mu$s bins, as seen in Figure 4.1. This corresponds to the cyclotron period in the storage ring divided into smaller sections of time to facilitate the pileup correction. Once the correction is complete, the histogram is rebinned to the full-cyclotron period bins to reduce impacts of the so-called fast rotation as the beam de-bunches following injection. A small impact of the cyclotron frequency remains, and one mitigation strategy will be discussed later in this chapter.

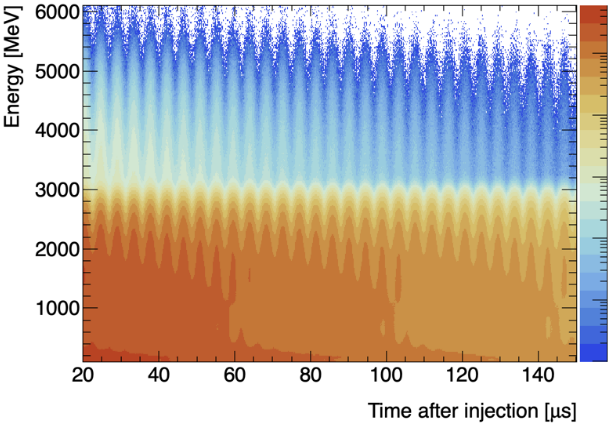

Figure 4.1: Hit times and energies of decay positrons. The oscillation in intensity at ω_a is clearly visible. Entries with energy above the beam energy of 3.1 GeV are necessarily pileup contamination.

4.3 Pileup correction

The path length for a higher-energy positron to a calorimeter for detection is longer than that of lower-energy positron. This means that if the two positrons are detected at the same time, they had different spin phases when emitted. If two lower-energy positrons were to hit a detector within the effective deadtime (meaning from both hardware and reconstruction effects), one positron with a higher energy would be reported. Unfortunately, an ensemble of these events would bias the phase from the true value. The amount of pileup is changing throughout the fill due to the changing number of decays, so this constitutes an early-to-late effect that could bias the measurement of ω_a and must be accounted for before fitting.

The pileup correction utilized for this work was designed and written by Aaron Fienberg [43]. The basic method will be explained here, but for more detail the reader is directed to

Aaron Fienberg's thesis, Reference [43].

The expression $\rho(E,T)$ will be used to express what a calorimeter would measure without any pileup contamination. The parameter Δt is the effective deadtime of the calorimeter, which includes effects both from pulse fitting and from any clustering algorithm employed. The contamination from double pileup (two positrons hitting a detector almost simultaneously, resulting in a single reported positron hit) goes as:

$$\begin{aligned}\delta\rho_{pu,d}(E,T) &\propto \Delta t\left[\int \rho(E-E_2,t)\cdot\rho(E_2,t)dE_2 - 2\rho(E,t)\int\rho(E_2,t)dE_2\right] &(4.1)\\ &= \Delta t\left[\rho_{d+}(E,t) - 2\rho(E,t)\int\rho(E_2,t)dE_2\right]. &(4.2)\end{aligned}$$

This expression is showing that the true measured (E,t) spectrum will include positive contributions of separate hits with a summed energy of E and negative contributions of the two true counts lost as a result. The redefinition of the integral in the first term will be convenient as the process to calculate the contribution is described.

By taking the measured $N(E,t)$ spectrum as a first guess at the unperturbed spectrum $\rho(E,t)$, the term $\delta\rho(E,t) = \rho_{d+}(E,t) - 2\rho(E,t)\int\rho(E_2,t)dE_2$ can be calculated numerically. This term is then scaled by a parameter A such that:

$$\int_{3500\,\mathrm{MeV}}^{6000\,\mathrm{MeV}} \left(N(E,t) - A\,\delta\rho(E,t)\right)\,dE\,dt = 0. \tag{4.3}$$

This equation is stating that the integral of the corrected spectrum in the unphysical region of $E \in [3500, 6000]\,\mathrm{MeV}$ should be 0 after correcting for pileup. By then taking $\rho(E,t) = N(E,t) - A\,\delta\rho(E,t)$, this process can be iterated to get a more precise pileup correction. Additionally, the method has been extended to account for triple pileup (3 simultaneous positron hits that result in either 1 or 2 reported counts). The success of the pileup correction can be seen in Figure 4.2, where almost all of the reported hits with unphysical energy $E > 3.1\,\mathrm{GeV}$ have been removed.

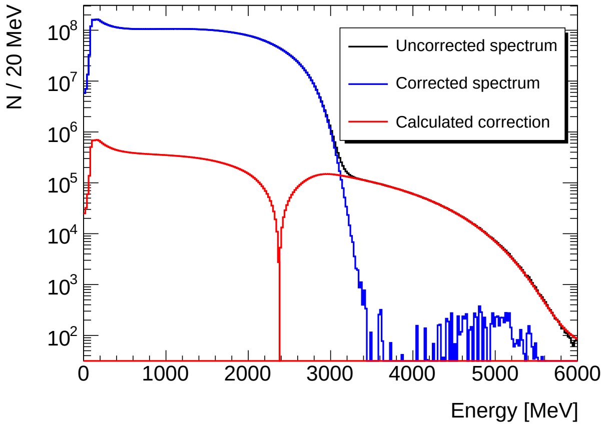

Figure 4.2: Positron energy spectrum before and after pileup correction. The Run − 2D energy spectra before (black) and after (blue) pileup correction are shown for the fit range of $t \in [30, 650]\,\mu$s. The absolute value of the pileup correction is shown in red; entries with $E > 2.4\,$GeV are subtracted from and entries with $E < 2.4\,$GeV are added to the uncorrected spectrum.

4.4 Functional form

The simplest form of the ω_a fit is a sinusoid modulating the exponential decay. This is often referred to the "the 5-parameter fit:"

$$N(t) = N_0 \exp\left(-t/\gamma\tau_\mu\right)\left(1 + A\cos\left(\omega_a t - \phi\right)\right), \tag{4.4}$$

where N_0 is some normalization factor, $\gamma\tau_\mu$ is the Lorentz-boosted lifetime, ω_a is the frequency in question, ϕ is the average spin phase at injection, and A is the asymmetry. The first alteration to this form comes from a software blinding that is specific to this analyzer. A hash of secret string is used to produce a constant offset that is unknown to the analyzer to avoid bias. This offset ΔR can be in the range of $[-24, 24]\,$ppm, much larger than the $< 1.0\,$ppm statistical uncertainty on the fitted ω_a in this data. This transforms the fit to:

$$N(t) = N_0 \exp\left(-t/\gamma\tau_\mu\right)\left(1 + A\cos\left(\omega_{a,ref}\left(1 - (R - \Delta R) * 10^{-6}\right)t - \phi\right)\right), \tag{4.5}$$

where ΔR is a fixed offset from the frequency $\omega_{a,ref} = 2\pi \times 0.2291\,\mathrm{MHz}$ [47]. The parameter R is then fit to the data using a χ^2 minimization. The fitted value on R is inconsequential, but the fitted error is equivalent to the statistical error on the resulting measurement of ω_a. As such, values and errors of R will be used as a proxy for the anomalous precession frequency throughout this chapter. A sample fit using the form listed above can be seen in Figure 4.3.

To determine the optimal threshold above which to fit the data, the lower energy cut can be scanned using the 5-parameter fit. The resultant error on R is minimized when selecting positron energies above 1.7 GeV, as seen in Figure 4.3.

The fit start time for all fits shown below will be $30.14\,\mu$s after muon beam injection. This particular value is chosen to be after the electrostatic quadrupole system has reached the desired voltage and the detector gains have stabilized. The lost muon and pileup rates also decrease through the fill, so those rates are reduced by starting later. These effects are balanced against the gain in statistical by starting the fit earlier. A zero-crossing is chosen to reduce the impact of slow terms on the measurement of ω_a.

4.5 Incorporating beam motion effects

The Fourier transform of the residuals from the 5-parameter fit can be seen in Figure 4.3. The peaks that are visible arise from the motion of the beam within the storage ring. Figure 4.4 shows the acceptance vs. decay position for the 24 calorimeters. One driver of differences between the calorimeters is the material that a positron must travel through between the point of decay and the calorimeter. Figure 2.3 gives an indication of these regions. Calorimeters 13 and 19 sit behind the straw tracker stations, limiting the vertical acceptance of the calorimeters. Calorimeters 2-4, 9, 10, 14-16, and 20-22 sit behind the electrostatic quadrupole plates. The magnetic kicker plates subtend the region in front of calorimeters 5-8. The vacuum chamber used near the inflector region, where calorimeters 23 and 24 are, is different from the others and impacts the acceptance differently. The other calorimeters (1, 11, 12, 17, and 18) are behind relatively free space and thus have a higher overall acceptance [49].

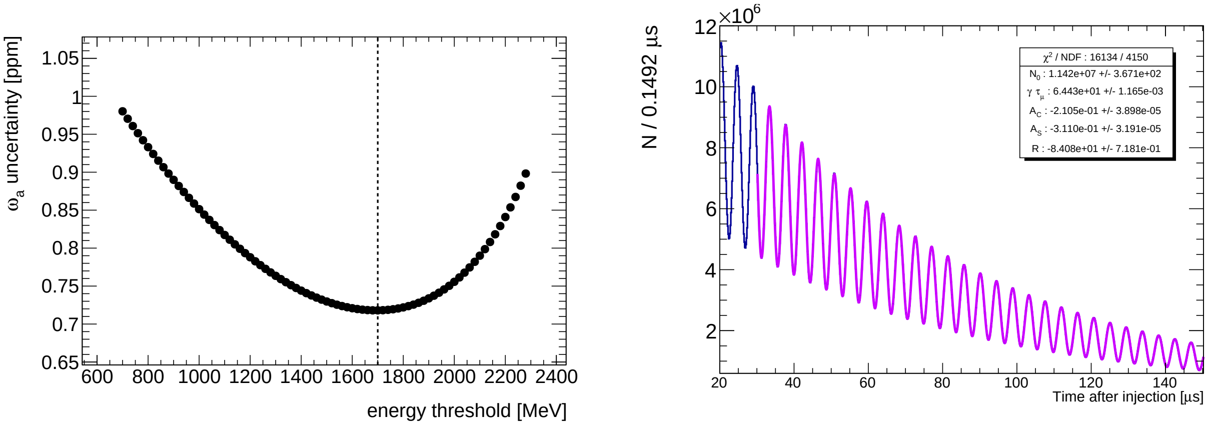

(a) Lower energy threshold scan

(b) Sample 5-parameter fit

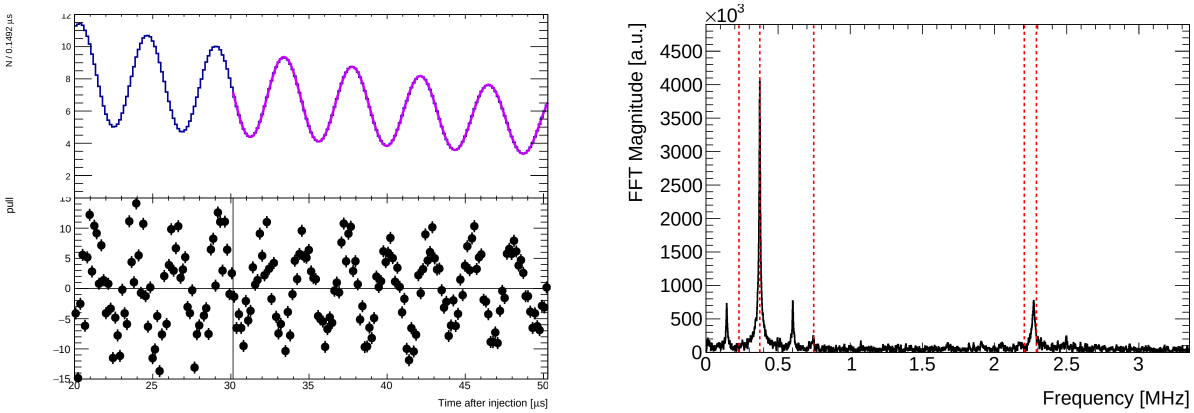

(c) Fit and pulls for the 5-parameter fit

(d) FFT of residuals from the 5-parameter fit

Figure 4.3: Determining the optimal threshold for fitting, a 5-parameter ω_a fit, and the resulting pulls and FFT of residuals. The result of scanning the energy threshold is shown in (a). Hits above 1.7 GeV are then histogrammed and fit, as shown in (b). In (c), the pulls are shown for bins early in the fit. A clear oscillation is visible. This oscillation is reflected in the Fourier transform of the residuals, (d). The red dashed lines indicate frequencies that will be incorporated into the fit function. They are, from left to right: f_a, f_{CBO}, $2f_{CBO}$, f_y, and f_{VW}. Peaks at the frequencies $f_{CBO} \pm f_a$ are also visible in the FFT spectrum.

The oscillatory beam motion then changes the acceptance and must be incorporated into the fit function. The main peak is the horizontal (radially inward and outward) motion of the beam, related the the electrostatic quadrupole voltage by the equation in Chapter 1, aliased at the cyclotron frequency, called the coherent betatron oscillation (CBO). Due to how close the period of the $x-$oscillation is to the cyclotron frequency, this effect is vastly reduced when fitting data from all calorimeters. Differences in acceptances and the imperfect cancellation both necessitate including it in the full fit function, and fits to individual calorimeters require it even more strongly. This CBO term is the only contribution to the fit function that appreciably changes the fitted value of R when included.

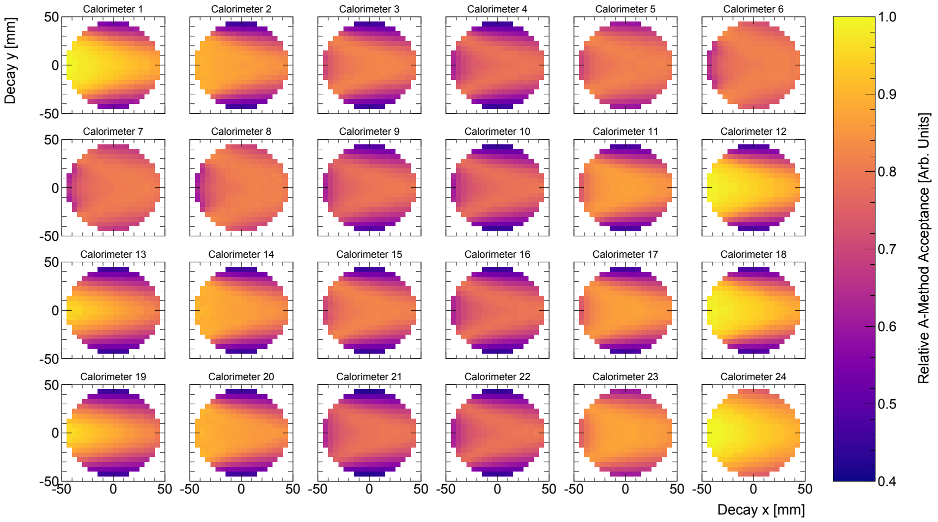

Figure 4.4: Acceptance as a function of decay position. Features in these plots are heavily dependent on the material between the beam and the calorimeter at each location, such as for calorimeters 13 and 19 where the straw trackers limit the acceptance from the vertical extrema of the beam. The non-uniformity across the storage region necessitates modifying the normalization term of the fit function to account for the different modes of beam motion. Image reproduced from [47].

The frequencies of note are listed in Table 4.1; the ones to be included in the fit function are the radial coherent betatron oscillation (CBO), vertical waist (VW), and $y-$oscillation.

The separate beam frequencies are related through the field index as described in Chapter 1. When fitting, these can be coupled or allowed to fit separately. In this formulation of the fit function, they will be fit separately. Another term is included for the 2CBO harmonic. Each of the frequencies has a decoherence lifetime, but the lifetimes of the width oscillations are half of the respective CBO and $y-$oscillation lifetimes ($\frac{1}{2}\tau_{CBO} = \tau_{2CBO}$ and $\frac{1}{2}\tau_y = \tau_{VW}$). The measured phase ϕ and asymmetry A are similarly dependent on decay position, but at a much smaller level. As such, only a CBO term is included as a correction to these parameters.

The CBO frequency depends on the field index, which depends on the strength of the electrostatic quadrupoles. After scraping, the voltage on the quadrupoles approaches the nominal value with a lifetime of $\approx 6.5\,\mu$s. While nearly at nominal voltage at the time of fit start, the CBO frequency is still changing early in the fill. The straw tracking detectors are able to characterize this changing frequency, which is then used as in the fits to calorimeter data. This step was much more important in the $\mathrm{Run}-1$ data when the quadrupole system had several broken resistors that introduced a longer relaxation lifetime.

Table 4.1: Beam motion frequencies. Frequencies to be included in the fit function are colored red, while those that are above the Nyquist frequency $(f_c/2)$ are blue.

Name	Expression	Frequency [MHz]	Period [μs]
Anomalous precession	f_a	0.229	4.365
Cyclotron	f_c	6.702	0.1492
Horizontal betatron	f_x	6.33	0.158
Vertical betatron	f_y	2.19	0.46
Coherent betatron oscillation (CBO)	$f_{CBO} = f_c - f_x$	0.37	2.7
Vertical waist	$f_{VW} = f_c - 2f_y$	2.27	0.44
2CBO	$f_{2CBO} = 2\,(f_c - f_x)$	0.74	1.35

4.6 Muon losses

One correction that must be made to the N parameter when fitting is to account for muons that exit the storage region without decaying. This phenomenon is dubbed "muon losses." The functional form is derived from the rate of losses as measured using the calorimeters. Since the muons act as as minimum ionizing particles, coincidences can be identified in consecutive calorimeters using the very sharp energy (170 MeV) and time-of-flight (6.2 ns) signature. When using triple coincidences with cuts of $E \in [100, 250]\,\mathrm{MeV}$ and $\Delta t \in [5.5, 7]\,\mathrm{ns}$, a problem arises with contamination within the lost muon sample.

4.6.1 Deuteron contamination

By constructing the coincidences with larger travel time and hit energy windows, the contaminant can be shown to have a time-of-flight of 7.4 ns and an energy of 300 MeV. These events are common enough that they appear to dominant in the low energy region late into the muon fills. The frequency at which they appear is very consistently 0.5 MHz, as can be seen in the middle plot of Figure 4.5. If this frequency is assumed to be a travel time around the 505 m-long delivery ring, it would correspond to $\beta = 0.863$. This speed looks like a particle of mass $\approx 1.8\,\mathrm{GeV}$, at the design momentum of $3.1\,\mathrm{GeV}/\mathrm{c}$. All of this information indicates the particle could be a deuteron, so it will be referred to as such. Detected deuterons are localized to detectors near the injection site and the magnetic kickers, so it appears as though the deuterons leak into the ring from the delivery ring and exit the ring quickly.

While both the deuterons and lost muons have very sharp energy and time-of-flight peaks, a small constant background of what appears to be deuterons has been observed in the lost muon spectrum constructed with triple coincidences. To reduce contamination from the deuterons, two measures were taken. First, quadruple coincidences were used instead of triple coincidences, using the same energy and time windows mentioned above. Secondly, the time of the fourth hit in the quadruple coincidence minus the time of the first hit is

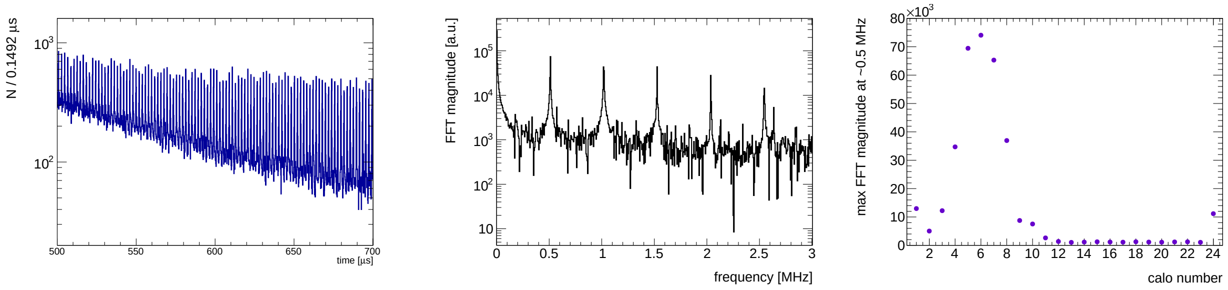

(a) Low energy events late in fill (b) FFT of low energy events (c) 0.5 MHz peak size vs. calo

Figure 4.5: Evidence for deuteron leakage into the storage ring. Events with energies between 100 MeV and 500 MeV are histogrammed late in the fill (a). A clear peaked structure is seen; a frequency of ≈ 0.5 MHz (and harmonics) is clearly visible in the FFT (b). The magnitude of the 0.5 MHz peak in the FFT of this spectrum is used as a proxy for the number of this type of particles hitting each calorimeter (c). The localization to the injection site and near the kickers indicate this is some particle leaking into the storage ring periodically throughout the fill.

restricted to be $t_4 - t_1 \leq 3\delta t$ where $\delta t = 6.2\,\mathrm{ns}$ is the nominal time of flight between two calorimeters. This reduces the statistics used to construct the lost muon function but is overall a cleaner sample of just lost muons; the measured losses from the two methods can be seen in Figure 4.6a. The fit has been performed using the lost muon function constructed from both triple and quadruple coincidences with both the symmetric and asymmetric time cut each. No difference was found in the resulting value for ω_a.

4.6.2 Altering the fit function to account for lost muons

Calling the measured loss rate from the quadruple coincidences $L(t)$, the change in the number of particles N is

$$dN = -\frac{N}{\tau_\mu}dt - \frac{1}{\epsilon}L(t)dt, \tag{4.6}$$

where ϵ is some efficiency of detection that is assumed to be constant through the fill. This leads to the number of particles stored as a function of time

$$N_s(t) = N_0 \left(1 - \frac{1}{\epsilon N_0} \int_0^t L(t') \exp(t'/\gamma\tau_\mu) dt'\right) \cdot e^{-t/\gamma\tau_\mu} \tag{4.7}$$

$$= N_0 \left(1 - K_{loss} \cdot \Lambda(t)\right) \cdot e^{-t/\gamma\tau_\mu} \tag{4.8}$$

$$= N_0 \cdot N_{loss}(t) \cdot e^{-t/\gamma\tau_\mu}. \tag{4.9}$$

The quantity $1/\epsilon N_0$ is defined as a new fit parameter K_{loss}. By construction, the fitted value should be $K_{loss} \geq 0$. A negative value would indicate that there is some as-yet uncorrected effect in the data, as it should not be possible for the storage ring to gain muons throughout a fill.

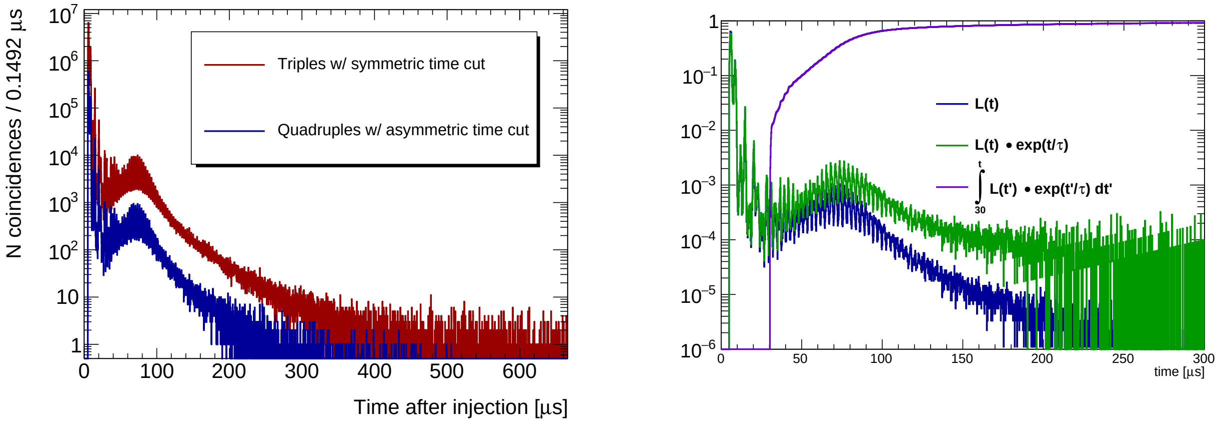

(a) Muon loss identification techniques

(b) Loss function to be incorporated into fit

Figure 4.6: Comparison of lost muon identification techniques (a) and the cumulative loss function used when fitting (b). Triple coincidences with consistent energy and time-of-flight cuts are shown with quadruple coincidences where the total time of flight is restricted to be less than 3× the nominal time of flight between calorimeters. This reduces contamination from deuterons while preserving the shape of the lost muon function. In (b), the quadruple coincidences are turned into the cumulative loss function $\Lambda(t)$ used to adjust the normalization term in the fit function.

4.7 Full fit function

When including all effects mentioned above, the fitted function becomes:

$$N(t) = N_0 \cdot N_{loss}(t) \cdot N_x(t) \cdot N_y(t) \cdot \exp\left(-t/\gamma\tau_\mu\right) \cdot \left(1 + A_0 \cdot A_x(t) \cdot \cos\left(\omega_a t - \phi_0 \cdot \phi_x(t)\right)\right). \quad (4.10)$$

$$\begin{aligned}
\omega_a &= \omega_a^{ref} \cdot \left(1 + R \cdot 10^{-6}\right) \\
N_{loss} &= 1 - K_{loss} \cdot \Lambda(t) \\
N_x(t) &= 1 + \exp\left(-t/\tau_{CBO}\right) \cdot \left(A_{CBO,N_c} \cos\left(\omega_{CBO} t\right) + A_{CBO,N_s} \sin\left(\omega_{CBO} t\right)\right) \\
&\quad + \exp\left(-2t/\tau_{CBO}\right) \cdot \left(A_{2CBO,N_c} \cos\left(2\omega_{CBO} t\right) + A_{2CBO,N_s} \sin\left(2\omega_{CBO} t\right)\right) \\
N_y(t) &= 1 + \exp\left(-t/\tau_y\right) \cdot \left(A_{y,N_c} \cos\left(\omega_y t\right) + A_{y,N_s} \sin\left(\omega_y t\right)\right) \\
&\quad + \exp\left(-2t/\tau_y\right) \cdot \left(A_{VW,N_c} \cos\left(\omega_{VW} t\right) + A_{VW,N_s} \sin\left(\omega_{VW} t\right)\right) \\
\phi_x(t) &= 1 + \exp\left(-t/\tau_{CBO}\right) \cdot \left(A_{CBO,\phi_c} \cos\left(\omega_{CBO} t\right) + A_{CBO,\phi_s} \sin\left(\omega_{CBO} t\right)\right) \\
A_x(t) &= 1 + \exp\left(-t/\tau_{CBO}\right) \cdot \left(A_{CBO,A_c} \cos\left(\omega_{CBO} t\right) + A_{CBO,A_s} \sin\left(\omega_{CBO} t\right)\right)
\end{aligned}$$

Additional terms in the $y-$ and $x-$oscillations are motivated by a beam moments-based approach but were found to be unimportant when fitting a dataset of this size. The beam oscillation terms are parametrized in the form $N_s \sin(\omega t) + N_c \cos(\omega t)$ rather than the equivalent $N \cos(\omega t - \phi)$ to stabilize the fit when these parameters are small. See Table 4.2 for the resulting values and associated fit errors when applied to the $\mathrm{Run} - 2\mathrm{D}$ dataset. Most of the parameters are not highly correlated with R apart from ϕ. A brief discussion of the relationship with ϕ can be found in Chapter 1.

The resulting $\chi^2/\mathrm{ndf} = 4248/4132 = 1.028$ is slightly outside the expected range of $1 \pm \sqrt{1/4132} = 1 \pm 0.016$. While no frequency peaks can still be discerned in the FFT of the residuals (see Figure 4.7), it is possible that a uncorrected slow term is impacting the goodness of fit. This could arise from an imperfect pileup correction or a changing detector gain.

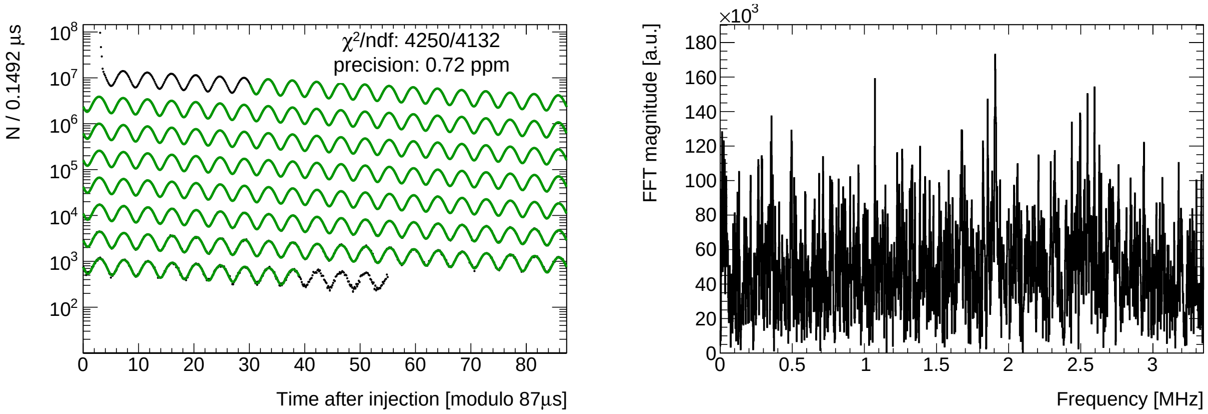

(a) Full T-method fit to the Run − 2D dataset

(b) FFT of residuals

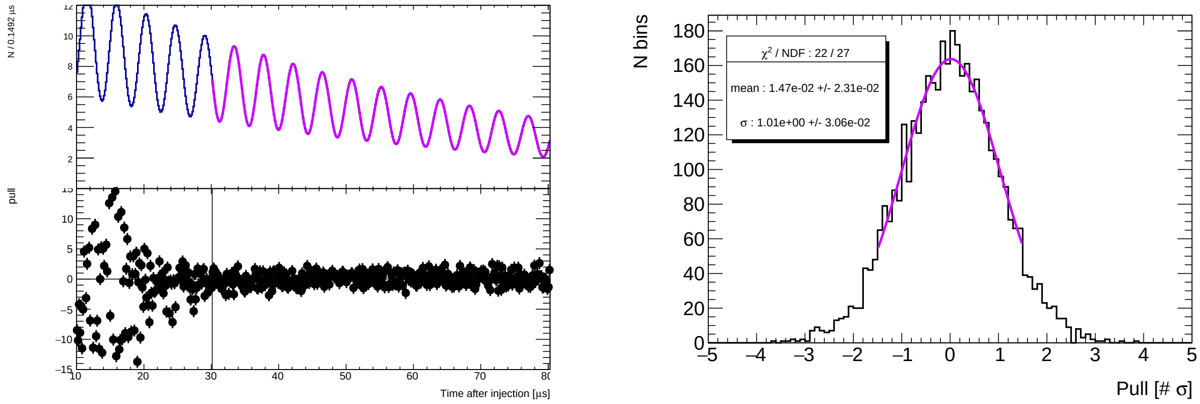

(c) Pulls at early times

(d) Distribution of pulls

Figure 4.7: Full fit to the data, resulting pulls, and FFT of the residuals. A fit to the Run − 2D dataset using the function in Equation 4.10 is shown in (a). The FFT of residuals (b) is featureless, apart from potentially a small uncorrected slow term. The pulls at early times (c) show nothing of note, and the distribution of pulls across the full fit (d) shows the expected Gaussian distribution with mean = 0 and width = 1.

Table 4.2: Parameters for the full fit function. The fitted parameter values, errors, and correlations with the parameter of interest (R) are shown for the full T-method fit to the Run $-$ 2D dataset.

Parameter	Meaning	Value	R-corr [%]
N_0	Normalization	$(1.144 \pm 0.001) \times 10^7$	-0.06
$\gamma\tau_\mu$	Boosted muon lifetime	$(64.435 \pm 0.002)\,\mu\text{s}$	-0.03
A	Precession asymmetry	(0.37558 ± 0.00002)	0.54
ϕ	Precession phase	$(-2.1654 \pm 0.0001)\,\text{rad}$	82.70
R	Blinded precession frequency	$-(82.13 \pm 0.72)\,\text{ppm}$	100.00
K_{loss}	Muon loss correction	(12.2 ± 9.0)	-0.06
$\tau_{CBO} = 2\tau_{2CBO}$	CBO decoherence lifetime	$(268.2 \pm 15.3)\,\mu\text{s}$	-2.44
ω_{CBO}	CBO frequency	$(2.3400 \pm 0.0002)\,\text{rad}\,/\,\mu\text{s}$	0.55
A_{CBO,N_c}	ω_{CBO} cosine modulation of N_0	$(3.43 \pm 0.07) \times 10^{-3}$	3.43
A_{CBO,N_s}	ω_{CBO} sine modulation of N_0	$(-1.33 \pm 0.07) \times 10^{-3}$	-0.44
A_{2CBO,N_c}	$2\omega_{CBO}$ cosine modulation of N_0	$(-2.65 \pm 0.45) \times 10^{-3}$	-0.16
A_{2CBO,N_s}	$2\omega_{CBO}$ sine modulation of N_0	$(1.11 \pm 0.45) \times 10^{-4}$	1.06
A_{CBO,A_c}	ω_{CBO} cosine modulation of A	$(6.8 \pm 4.6) \times 10^{-5}$	-0.16
A_{CBO,A_s}	ω_{CBO} sine modulation of A	$(-4.0 \pm 4.6) \times 10^{-5}$	2.69
A_{CBO,ϕ_c}	ω_{CBO} cosine modulation of ϕ	$(2.9 \pm 12.8) \times 10^{-5}$	-0.98
A_{CBO,ϕ_s}	ω_{CBO} sine modulation of ϕ	$(-5.8 \pm 12.8) \times 10^{-5}$	0.09
$\tau_y = 2\tau_{VW}$	Vertical oscillation lifetime	$(67.70 \pm 7.14)\,\mu\text{s}$	-0.33
ω_y	y-oscillation frequency	$(13.76 \pm 0.01)\,\text{rad}\,/\,\mu\text{s}$	-0.13
A_{y,N_c}	ω_y cosine modulation of N_0	$(1.24 \pm 1.04) \times 10^{-4}$	0.02
A_{y,N_s}	ω_y sine modulation of N_0	$(-1.31 \pm 0.98) \times 10^{-4}$	-0.45
ω_{VW}	Vertical Waist frequency	$(14.264 \pm 0.003)\,\text{rad}\,/\,\mu\text{s}$	-0.57
A_{VW,N_c}	ω_{VW} cosine modulation of N_0	$(-1.92 \pm 0.44) \times 10^{-3}$	0.39
A_{VW,N_s}	ω_{VW} sine modulation of N_0	$(2.87 \pm 0.51) \times 10^{-3}$	0.65
$\chi^2\,/\,\text{DOF}$		$4250\ /\ 4132 = 1.029$	

4.8 *Challenges with the Run* $-$ 2 *data*

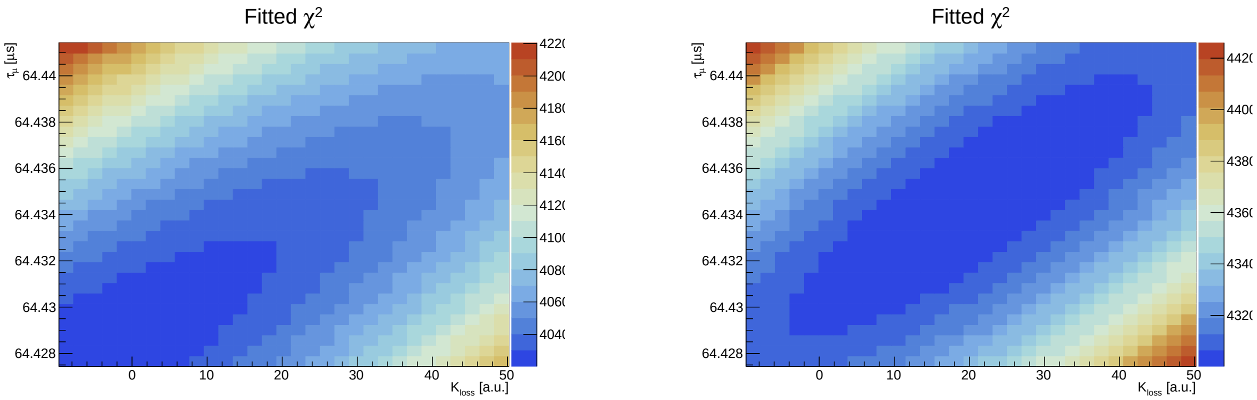

(a) χ^2 sweep in Run $-$ 2C dataset

(b) χ^2 sweep in Run $-$ 2D dataset

Figure 4.8: Resulting χ^2 values when scanning parameters for two datasets in Run $-$ 2. The two parameters $\gamma\tau_\mu$ and K_{loss} are manually scanned around the values preferred by a fit to the Run $-$ 2D dataset. The resulting χ^2 values are shown, for $\approx$ 4150 degrees of freedom. The preferred value of K_{loss} in the Run $-$ 2C data is unphysical, indicating an incomplete fit model or a problem in the production of the data. The differences between these two datasets, taken weeks apart, is unsettling and not yet explained.

Having solved the challenges involved in analyzing the Run $-$ 1 data, the hope among the collaboration was that Run $-$ 2 and Run $-$ 3 analysis would go relatively smoothly. Unfortunately, that has not been the case. One illustration of this can be seen in Figure 4.8. The χ^2-minimization, when left to its own devices, finds unphysical values for the slow terms (muon loss and lifetime) in the Run $-$ 2C dataset. This trend was true for all datasets in Run $-$ 2, apart from the Run $-$ 2D dataset. Possible sources of this inconsistency between datasets are hard to track down, as the pileup correction, calorimeter gain corrections, fitted lifetime, and muon loss terms are all highly correlated. Some errors in data production have been found that could potentially lead to a rate-dependent effect or challenges in pileup construction, but as of the time of writing, the corrected data has not yet been analyzed to determine if those problems directly impacted the fit.

4.9 Energy-binned fits and improved statistical power

Some parameters in the fit function, such as ϕ and A, are expected to change as a function of positron energy. The parameter of interest R should be compatible across energies, so fitting in separate energy bins is an important consistency check when finding ω_a.

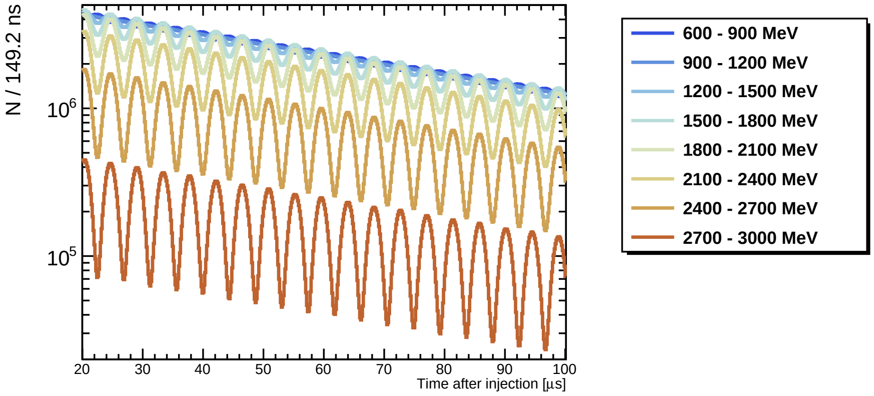

Figure 4.9: Energy-binned data. The decrease in statistics in high-energy bins is counteracted by the increased resolution of the oscillation at ω_a. By giving more weight to higher energy bins through either an energy-weighted or asymmetry-weighted analysis, the statistical precision can be improved from the nominal threshold-method.

The highest energy bins carry the most information about the spin direction (and precession) but have the fewest counts (see Figure 4.9). Rather than giving equal weight to each energy of positron, the hits could be weighted by energy to emphasize the highest energy bins. This provides a moderate improvement in precision on the ω_a fit, but a better method is available. The figure of merit for the ω_a fit goes as the value NA^2. By scaling each energy by the measured asymmetry A in that bin, the statistical precision of the fit can be maximized [65]. The improvement in precision on ω_a can be seen in Figure 4.10.

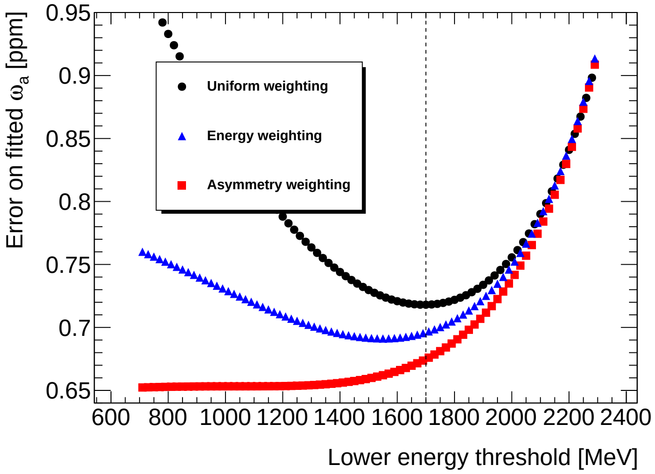

Figure 4.10: Statistical precision of three weighting methods as a function of energy threshold. The asymmetry-weighted method uses the data most efficiently [65].

4.10 Alleviating fast rotation effects using the kernel method

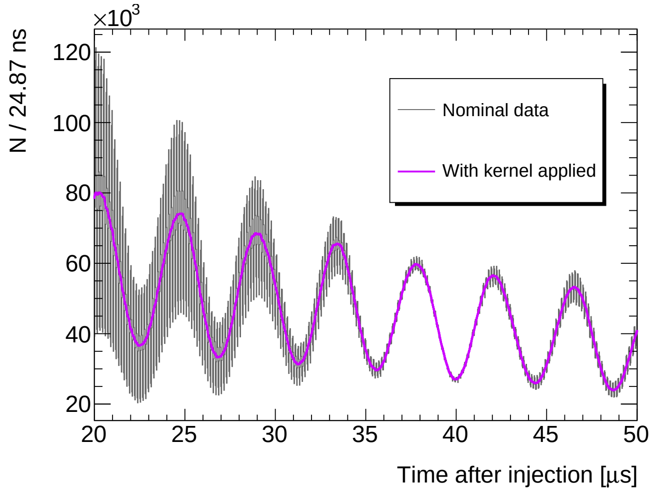

Figure 4.11: Data from a single calorimeter with and without the kernel applied. The bunched structure of the beam is very clear without use of the kernel; this slightly biases the per-calorimeter fits of ω_a. The kernel method smears each positron hit to the cyclotron period to eliminate the bunched structure.

4.10.1 Per-calorimeter fits

Similar to the energy-binned fits, fitting the data from each calorimeter individually is an important test. The CBO oscillations mostly cancel out when summing all calorimeters, so successfully fitting data from each calorimeter with much larger oscillations in the acceptance, asymmetry, and phase provides confidence in the fit model. It also provides an opportunity to catch any calorimeter stations that may have quietly malfunctioned during data-taking in a way that was not caught during reconstruction.

4.10.2 Kernel method

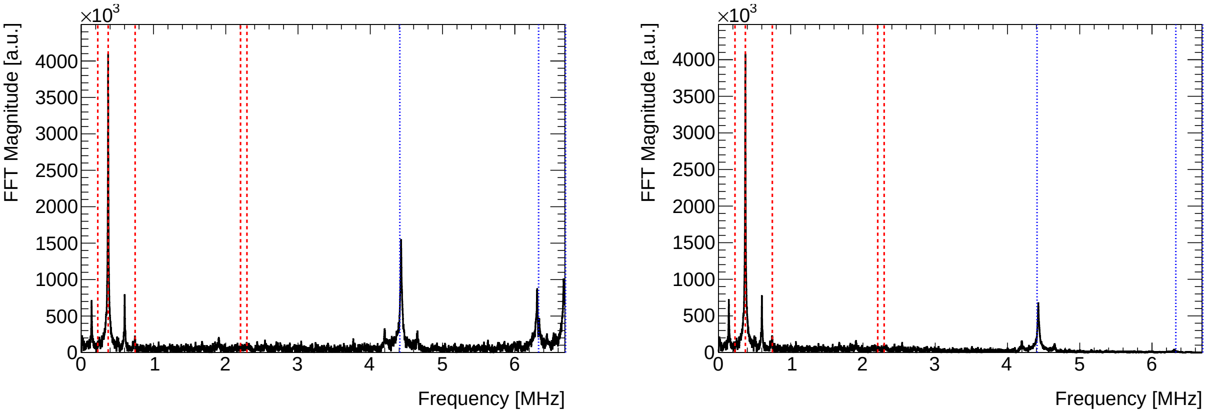

(a) FFT, nominal

(b) FFT, kernel

Figure 4.12: FFT of 5-parameter fit residuals without (left) and with (right) the kernel applied. The peak at f_c is clear in (a), whereas it has completely disappeared in (b). The frequencies in red will be incorporated into the fit; frequencies marked with blue are aliased when using T_c-width bins.

Binning data in the cyclotron period is designed to ameliorate the impact of the bunched nature of the muon beam on ω_a fits. The effect is still present at a small level when fitting per calorimeter. As can be seen in Figure 4.13a, the fitted R value per calorimeter shows a sinusoidal oscillation about the fitted value from the summed data. In the past this effect has been mitigated through randomly adjusting the times of positron hits across a full cyclotron period [47]. Since random numbers are involved, the process must be repeated a number of times and the resulting fitted R values averaged for the final quoted result. This contributed a small systematic error on the measurement of ω_a, related to the width of the R value distribution when using different seeds in the random number generator.

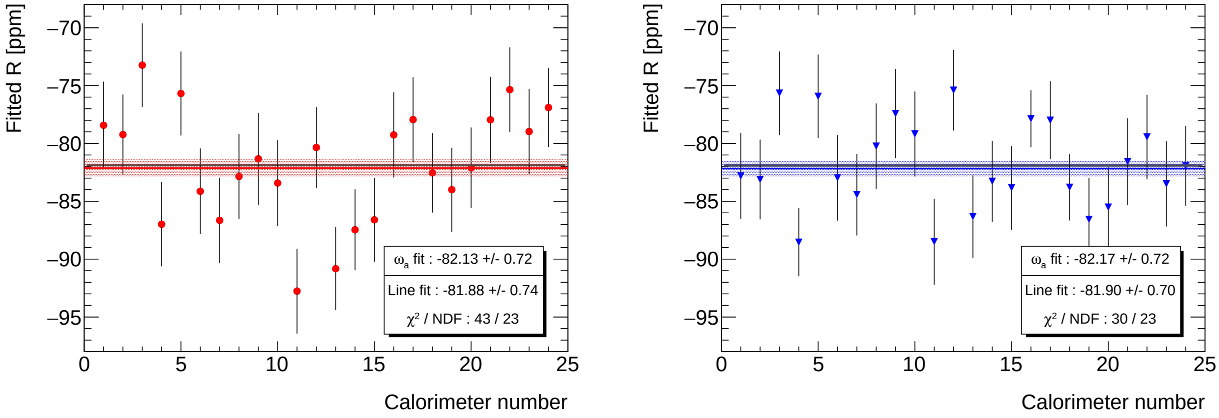

(a) Per-calo fits, nominal

(b) Per-calo fits, kernel

Figure 4.13: Per-calorimeter fits for R without (left) and with (right) the kernel applied. When performing per-calorimeter fits, the uncorrected fast rotation structure biases the resulting R-values differently around the storage ring. This can be seen in (a), where the points oscillate about the full fit, shown as the solid red line. In (b), the kernel has been applied to each positron hit. For comparison purposes, the fitted ω_a values from the all-calorimeter fit is displayed along with the result of fitting the per-calorimeter results to a single value. The spread about the R-value from a fit to all data (blue line) is consistent with statistical spread, evident by the χ^2/NDF value for the single parameter fit.

To avoid random numbers and the computational time associated with running over each dataset dozens of times, a method has been introduced to spread each positron hit uniformly over a time range of T_c. The statistics of the dataset are unchanged, but this process introduces correlations between time bins that must be accounted for when fitting.

Appendix C illustrates the kernel method and shows the calculation of the covariance matrix necessary for fitting.

Figure 4.11 shows the disappearance of the fast rotation bunching effect in the data when using the kernel method. The same can be seen in the FFT of the 5-parameter fit residuals shown in Figure 4.12, where half-cyclotron period bins were used instead of the nominal full cyclotron period bins. The disappearance of the R-wave can be seen in Figure 4.13. The per-calorimeter fits to R are inconsistent with a single value for the unmodified data but become consistent with a single value after the kernel is applied.

4.11 Systematics

Each step of the reconstruction process must be evaluated for potential impacts on the measured value of ω_a. In this section, two particular systematics related to the way the calorimeters are used will be presented. Due to the unrefined nature of the $\mathrm{Run}-2$ analysis, the studies shown will be performed on $\mathrm{Run}-1$ data. The potential effects are amplified to better ascertain how problematic they may be. Of course this is meant to be some representative studies of detector-related systematics, not a comprehensive list of what needs to be evaluated.

4.11.1 Miscalibration

When calibrating the detectors, as described in Chapter 3, typically the three hits of an identified triple coincidence are all used to form the average hit energy of a lost muon in each channel. Since the muon is necessarily curling inward, the spatial distribution of the three hits across the face of a calorimeter is different. As can be seen in Figure 4.14, the first hit in a triple coincidence (shown in black) hits primarily the column of crystals radially closest to the beam and the two central rows. With each successive hit (red, green), the average hit position on a calorimeter moves radially inward and the variance on the vertical position increases.

By using all three hits from a triple coincidence, the assumption is made that each hit

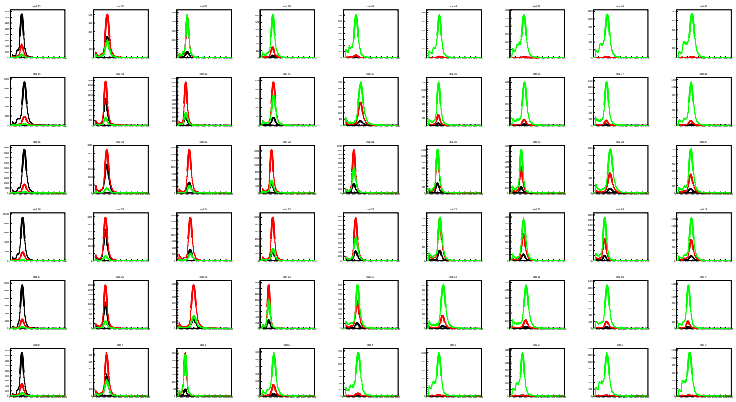

Figure 4.14: Spatial distribution of the three lost muon hits across the face of a calorimeter. When identifying lost muon triple coincidences, the geometry of the calorimeters and trajectories of the muons necessarily impacts the acceptance. Shown here is the distribution of which crystals in a calorimeter were identified as the first (black), second (red), and third (green) hits in a coincidence, where the beam would be on the left going into the page. Small differences are found when fitting the peaks of the fitted pulse integrals, the value plotted here.

has the same characteristic energy. By comparing the peaks in each crystal stemming from the first, second, or third hit, this assumption can be assessed. As seen in Figures 4.15a and 4.15b, the characteristic energy of each hit follows some sort of pattern spatially across the calorimeter. This probably arises from the average path length muons traverse in those particular crystals, while still participating in a triple coincidence.

Combining the spatial pattern of energy differences per coincidence hit with the prevalence of each hit across a calorimeter, the spread in the relative differences in calibration could be several percent. This is not an early-to-late effect by itself. When coupled to like the beam motion present in Run $-$ 1D dataset (statistical precision of 630 ppb), it could conceivably pull the measured value of ω_a. By exaggerating the effect using the relative gain values shown in Figures 4.15c and 4.15d, the impact of this potential effect is determined to be at the level of $\Delta R \approx 10\,\text{ppb}$ for a realistic effect and $\Delta R \approx 150\,\text{ppb}$ effect using the hyper-exaggerated version. It should be noted this study may amplify the phase-acceptance effect in [49], thus overestimating of the effect of mis-calibration.

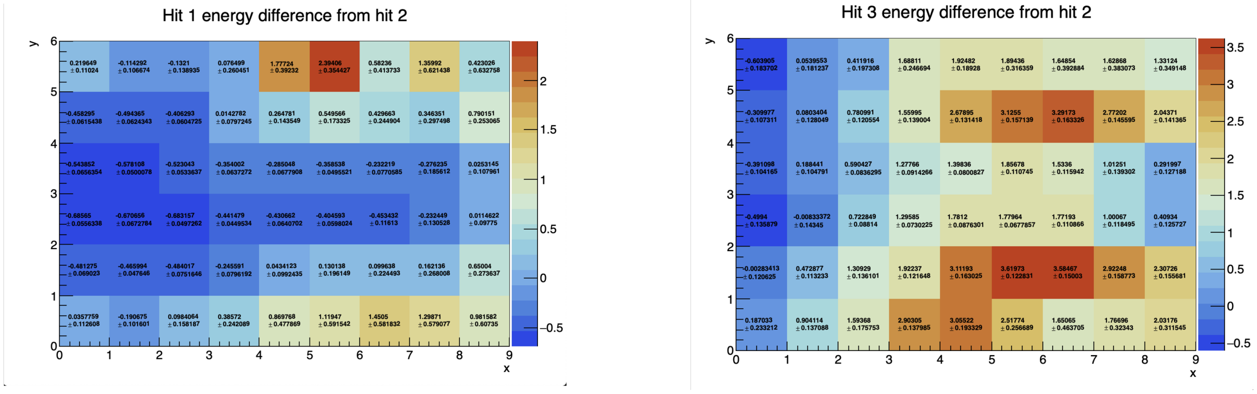
Hit 1 energy difference from hit 2

Hit 3 energy difference from hit 2

(a) First hit relative to second

(b) Third hit relative to second

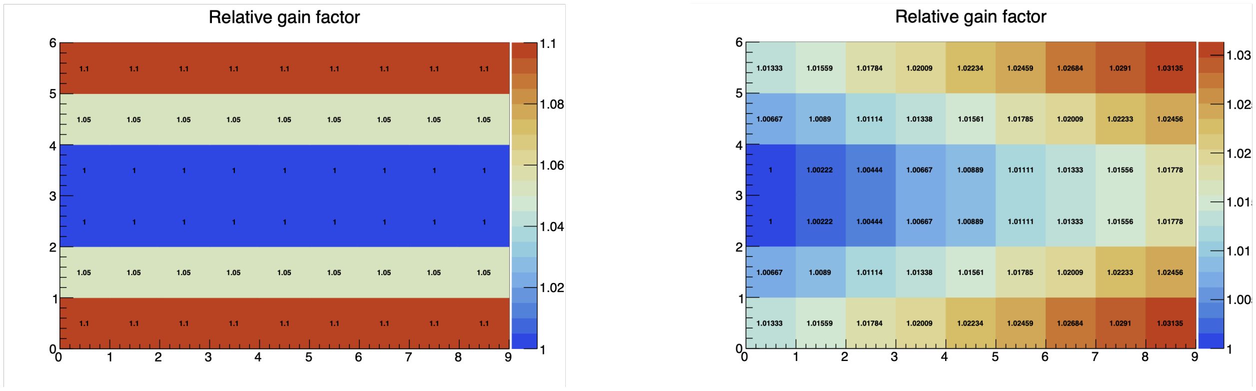
Relative gain factor

Relative gain factor

(c) Vertical gradient of energies

(d) Gradient based on distance to beam

Figure 4.15: Effects of miscalibrated calorimeters. In (a) and (b), the fitted pulse integral peaks for each of the three hits in a coincidence are compared for each crystal in a calorimeter. In general, each successive hit deposits more energy. There is clearly a spatial distribution, so the effect may arise from path length differences through the crystals. Combined with the spatial dependence of these hits in Figure 4.14, calorimeters may have a spread of relative differences in calibration across the face. In (c) and (d), two spatial maps of gain adjustments were formed to exaggerate the impact of differences in calibration. When applied to data, the differences in fitted R-values were small compared to the statistical precision of the fits.

4.11.2 Time-dependent fit threshold

An effective gain change could arise from the reconstruction software chain rather than the hardware itself. As an example, imagine that the energy required to fit a pulse changes from early-to-late in the fill. This would cause some energy to be missing from the resulting cluster, resulting in a gain-like effect. If the procedure for fitting multiple pulses in one channel at the same time (chain-fitting) impacts the effective energy scale that can be fit, a gain-like effect would be introduced at the muon decay rate.

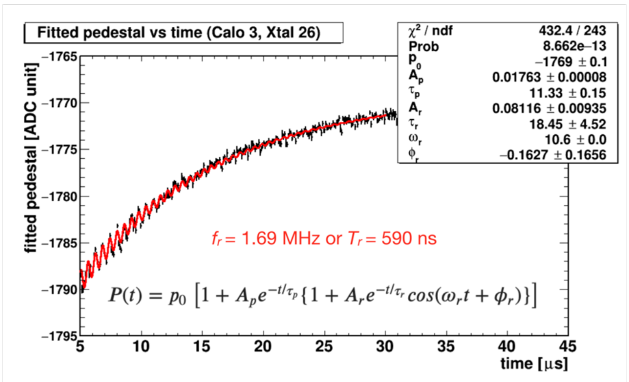

Figure 4.16: Fitted pedestal value from digitized SiPM pulses as a function of time in fill. A drop of 20 ADC counts with a recovery time of 11 μs is visible. Such an effect could be present in the data as a time-dependent island chopping threshold, as 1 ADC count $\approx$ 1 MeV.

A similar situation could arise from the island chopping, as operated in Run − 2 and beyond. An island is saved for a crystal and its surrounding neighbors whenever a sample exceeds 60 ADC counts above a pre-calculated pedestal value. This pedestal has been shown to change throughout the fill, as can be seen in Figure 4.16. The threshold to save a crystal (and neighbors) is therefore time-dependent with a lifetime of $\approx 12\mu$s. While this would not impact the central 3×3 grid of a high-energy pulse, smaller pulses on the edge of the cluster may be missed as a result. Again, this would manifest as a gain-like effect.

To study this type of effect, an artificial threshold was introduced into the clustering algorithm and fitting was performed on the Run − 1A dataset. Only crystals with energy

$E \geq 150 + \delta E \cdot \exp(-t/\tau)$ MeV would be included in clusters, where $t = 0$ is at beam injection. The value of 150 MeV was chosen to be well above the fitting and island chopping thresholds, so as to not interact with either of them. The values for δE and τ were chosen to be 10, 25 MeV and 20, 64 μs. These numbers are meant to be representative of the possible effects detailed above but exaggerated to demonstrate how robust the ω_a fitting procedure is.

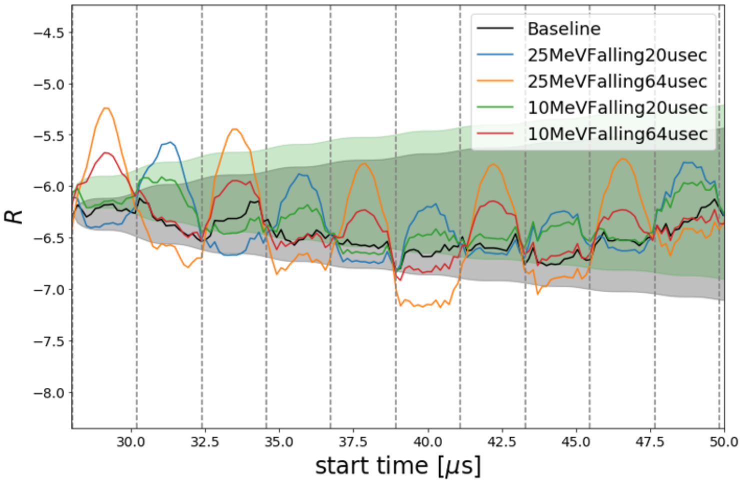

Figure 4.17: Change in fitted R as a function of fit start time for an artificial time-dependent fit threshold. The resulting R value is plotted vs. fit start time for several versions of the changing energy threshold study. The vertical gray lines indicate the zero crossings of the ω_a oscillation. The oscillatory nature of the fitted R-value is expected when a slow effect is introduced and motivates starting the fit at a zero crossing. The spread in R-values is small, so the impact of this effect is determined to be small.

The results of this study can be seen in Figure 4.17, where the fitted value of R is plotted vs. fit start time for the four situations detailed above and the nominal analysis dataset. The extreme sensitivity to gain-like effects is not unexpected. When starting at a zero-crossing of the ω_a oscillation, the sensitivity to these effects is greatly reduced and the spread in the fitted values is < 100 ppb (compared to a statistical precision of 1.35 ppm).

Chapter 5

CONCLUSIONS

5.1 Improvements to a_μ^{exp}

$$a_\mu^{exp} \propto \mathcal{R}_\mu^{'} = \frac{f_{\text{clock}}\ \omega_a^m\,(1 + C_e + C_p + C_{ml} + C_{pa})}{f_{\text{calib}}\ \langle \omega_p(x, y, \phi) \times M(x, y, \phi)\rangle (1 + B_k + B_q)} \tag{5.1}$$

Without recreating the table, the uncertainties and values of the various terms in Table 1.3 will be discussed, in the context of Equation 5.1. The largest single uncertainty remaining is the statistical uncertainty on the measurement of ω_a. The next result to be released from the Fermilab Muon $g-2$ experiment will include data from both $\text{Run}-2$ and $\text{Run}-3$, roughly 4× the amount of data included in the $\text{Run}-1$ result. This will ideally lead to a 2× reduction in the statistical error. Other points of improvement will include the errors on C_e, B_k, and B_q, where extensive measurement campaigns are underway to better understand the effects. The terms C_{pa} and C_{ml} will be greatly reduced in upcoming runs, as the stability of the beam position was improved between $\text{Run}-1$ and $\text{Run}-2$. The overall value of C_e is also expected to reduce beginning in $\text{Run}-3$, as improvements to the pulsed magnetic kicker system allowed for better centering of the muon beam within the storage region.

An additional experiment aiming to measure a_μ to 450 ppb is currently under construction at J-PARC [66]. While the precision will not be competitive with the ultimate Fermilab result, the J-PARC experiment will use a technique other than the magic momentum. An ultra-cold muon beam will be injected into a highly uniform magnetic, and each decay positron will be tracked from the point of decay [66]. The compatibility of the J-PARC and Fermilab measurements will be an important test of consistency between experimental methods.

Table 5.1: Systematics table for the measurement of ω_a. For Run $-$ 1, the result is still statistically dominated. Upcoming results will greatly reduce the statistical error while hopefully improving upon the systematic errors associated with pileup and the early-to-late effect. Reproduced from [47].

Run $-$ 1 data set	1a	1b	1c	1d
$\omega_a^m/2\pi$ (s^{-1})	229080.957	229081.274	229081.134	229081.123
$\Delta(\omega_a^m/2\pi)$ (s^{-1})	0.277	0.235	0.189	0.155
Statistical uncertainty (ppb)	1207	1022	823	675
Gain changes (ppb)	12	9	9	5
Pileup (ppb)	39	42	35	31
CBO (ppb)	42	49	32	35
Time randomization (ppb)	15	12	9	7
Early-to-late effect (ppb)	21	21	22	10
Total systematic uncertainty (ppb)	64	70	54	49
total uncertainty (ppb)	1209	1025	825	676

5.1.1 Systematic errors on ω_a

Table 5.1 shows the systematic uncertainties associated with the anomalous precession frequency measurement ω_a^m. The Run $-$ 1 measurement campaign was highly successful from a calorimeter standpoint, surpassing the systematics goals set in the TDR for gain changes [52]. This would not have been possible without hardware adjustments made early in the lifetime of the experiment to better stabilize the SiPM responses throughout the data-taking periods.

The most readily available improvement is that of the pileup correction algorithm, done in four separate ways by various analysis groups. One analysis group was able to achieve an uncertainty associated with the pileup as low as 5 ppb, so the average systematic uncertainty seen in Table 5.1 will probably be reduced in upcoming results. Similarly, using a kernel method like that described in Chapter 4 and Appendix C could negate the time randomization error. The remaining "early-to-late" effect listed refers to an as-yet un-modeled effect presenting as a gain-like effect and could be resolved in upcoming analyses.

5.2 Progress on Standard Model calculation

As the experimental precision improves, so will that of the calculated SM model. The Muon $g-2$ Theory Initiative will continue to update the recommended value for a_μ^{SM} as updates become available. The effort moving forward is largely centered on the HVP terms, as that has the largest error. Upcoming realeases of e^+e^- scattering data will yield improvements in the dispersion relation formulation. Developments in the lattice QCD calculations have already made progress towards achieving comparable precision. Lastly, a dedicated experiment to measure $a_\mu^{\text{HVP, LO}}$ directly has been proposed.

5.2.1 Lattice QCD evaluation of HVP contribution

On the same day the Run $-$ 1 results were announced, the first lattice QCD calculation of $a_\mu^{\text{HVP, LO}}$ with similar precision to the dispersion relations was released [67]. While still unresolved, there is tension between the values as calculated by the two methods. If the lattice result is trusted, the difference between experiment and theory would be approximately 2σ rather than greater than 4σ [67]. No other lattice groups have been able to replicate the result with similar precision, so for now the dispersion relation calculations are most reliable. However, the difference between the two approaches is intriguing.

5.2.2 The MUonE experiment

Moving from a more theoretical approach to a more data-driven approach, a technique to determine $a_\mu^{\text{HVP, LO}}$ from scattering $e\mu$ scattering data has been proposed [68, 69]. The associated experiment, MUonE, would scatter a beam of 150 GeV muons on the atomic electrons in a low-Z target. The expected achievable statistical precision on $a_\mu^{\text{HVP, LO}}$ is $\approx 0.3\%$, comparable to current dispersion relation techniques.

5.3 Beyond the Standard Model (BSM) explanations

A 4.2σ difference between a measured value and theoretically predicted value naturally leads to questions about whether the value of a_μ^{SM} truly includes all necessary interactions. As a long-standing discrepancy, many extensions to the Standard Model have been proposed to account for the difference; see [70] or [71] for a review of possible extensions to the SM. The classes of models include dark photons, two-Higgs doublet models (2HDM), and the ever-present supersymmetry (SUSY) [71]. Many models simultaneously incorporate a dark matter candidate particle. To explain such a large discrepancy (twice the size of the electroweak contribution) while avoiding constraints from other experiments requires some finesse.

5.4 Other interesting physics

5.4.1 Dipole moments

In Chapter 1, a brief note was made about the differences in the measured values of the fine structure constant α [37, 38] and the connection the g factor of the electron [39]. Models incorporating the discrepancies in a_e and a_μ and relating to current measurements of each particle's electric dipole moment (EDM) [6, 72] have been studied [73]. Both particles have a very small non-zero EDM in the SM, but experimental sensitivity has not yet reached the necessary precision to measure d_e and d_μ. Finding a value larger than the SM predicts for either would guide BSM model building. Improvements in the measurement of the electric dipole moment of the muon d_μ are expected from both the Fermilab and J-PARC experiments.

5.4.2 Charged lepton flavor violation

Following the description in Reference [73], the effective Hamiltonian for electromagnetic interactions (relevant to dipole moments) is

$$H = c_R^{l_f l_i} \bar{l}_f \sigma_{\mu\nu} P_R l_i F^{\mu\nu} + \text{h.c.}, \tag{5.2}$$

where the $l_{i/f}$ are lepton fields, P_R is the projection operator, and $c_R^{l_f l_i}$ is some coupling. In this description, the dipole moments are:

$$a_\mu = -\frac{4m_\mu}{e} \operatorname{Re} c_R^{\mu\mu}, \tag{5.3}$$

$$d_\mu = -2 \operatorname{Im} c_R^{\mu\mu}. \tag{5.4}$$

The branching ratio of the CLFV process of $\mu \to e\gamma$ would go as:

$$\text{Br}\left[\mu \to e\gamma\right] = \frac{\mu^3}{4\pi\Gamma_\mu}\left(|c_R^{e\mu}|^2 + |c_R^{\mu e}|^2\right). \tag{5.5}$$

Through measuring the three of these together, the magnitude and phase of c_R can be constrained for different classes of models.

The SM itself predicts a very small value for charged lepton violation, such as $\mu \to e\gamma$. At tree level, a non-zero value would be related directly to the values of the dipole moments (as shown above). While no evidence for charged lepton flavor violation (CLFV) has yet been found, models of BSM physics may incorporate a mechanism that leads to some flavor violation. Upcoming efforts such as Mu2e (Fermilab) [74], Mu3e (PSI) [75], and MEG-II (PSI) [76] are searching for CLFV in complementary channels. The "new physics" parameter space probed by these precision experiments would further guide models designed to account for the a_μ discrepancy.

BIBLIOGRAPHY

[1] **Muon $g-2$** Collaboration, B. Abi *et al.*, “Measurement of the positive muon anomalous magnetic moment to 0.46 ppm,” *Phys. Rev. Lett.* **126** (Apr, 2021) 141801. `https://link.aps.org/doi/10.1103/PhysRevLett.126.141801`.

[2] **Muon $(g-2)$** Collaboration, G. Bennett *et al.*, “Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL,” *Phys.Rev.* **D73** (2006) 072003, `arXiv:hep-ex/0602035 [hep-ex]`.

[3] R. L. Garwin, L. M. Lederman, and M. Weinrich, “Observations of the Failure of Conservation of Parity and Charge Conjugation in Meson Decays: The Magnetic Moment of the Free Muon,” *Phys. Rev.* **105** (1957) 1415–1417.

[4] T. Aoyama *et al.*, “The anomalous magnetic moment of the muon in the Standard Model,” *Phys. Rept.* **887** (2020) 1–166, `arXiv:2006.04822 [hep-ph]`.

[5] J. D. Jackson, *Classical Electrodynamics*. Wiley, New York, NY, 3rd ed. ed., 1999.

[6] **Muon $(g-2)$** Collaboration, G. Bennett *et al.*, “An Improved Limit on the Muon Electric Dipole Moment,” *Phys.Rev.* **D80** (2009) 052008, `arXiv:0811.1207 [hep-ex]`.

[7] L. H. Thomas, “The Kinematics of an electron with an axis,” *Phil. Mag. Ser.7* **3** (1927) 1–21.

[8] M. D. Schwartz, *Quantum Field Theory and the Standard Model*. Cambridge University Press, New York, NY, 2014.

[9] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*. Perseus Books, Reading, MA, 1995.

[10] T. Aoyama *et al.*, “Complete Tenth-Order QED Contribution to the Muon $g-2$,” *Phys. Rev. Lett.* **109** (2012) 111808, `arXiv:1205.5370 [hep-ph]`.

[11] T. Aoyama, T. Kinoshita, and M. Nio, “Theory of the Anomalous Magnetic Moment of the Electron,” *Atoms* **7** (2019) no. 1, 28.

[12] A. Czarnecki, W. J. Marciano, and A. Vainshtein, "Refinements in electroweak contributions to the muon anomalous magnetic moment," *Phys. Rev.* **D67** (2003) 073006, `arXiv:hep-ph/0212229 [hep-ph]`. [Erratum: Phys. Rev. **D73**, 119901 (2006)].

[13] C. Gnendiger, D. Stoeckinger, and H. Stoeckinger-Kim, "The electroweak contributions to (g-2)$_\mu$ after the Higgs boson mass measurement," *Phys. Rev.* **D88** (2013) 053005, `arXiv:1306.5546 [hep-ph]`.

[14] M. Davier *et al.*, "Reevaluation of the hadronic vacuum polarisation contributions to the Standard Model predictions of the muon $g-2$ and $\alpha(m_Z^2)$ using newest hadronic cross-section data," *Eur. Phys. J.* **C77** (2017) no. 12, 827, `arXiv:1706.09436 [hep-ph]`.

[15] A. Keshavarzi, D. Nomura, and T. Teubner, "Muon $g-2$ and $\alpha(M_Z^2)$: a new data-based analysis," *Phys. Rev.* **D97** (2018) no. 11, 114025, `arXiv:1802.02995 [hep-ph]`.

[16] G. Colangelo, M. Hoferichter, and P. Stoffer, "Two-pion contribution to hadronic vacuum polarization," *JHEP* **02** (2019) 006, `arXiv:1810.00007 [hep-ph]`.

[17] M. Hoferichter, B.-L. Hoid, and B. Kubis, "Three-pion contribution to hadronic vacuum polarization," *JHEP* **08** (2019) 137, `arXiv:1907.01556 [hep-ph]`.

[18] M. Davier *et al.*, "A new evaluation of the hadronic vacuum polarisation contributions to the muon anomalous magnetic moment and to $\boldsymbol{\alpha}(\mathbf{m_Z^2})$," *Eur. Phys. J. C* **80** (2020) no. 3, 241, `arXiv:1908.00921 [hep-ph]`. [Erratum: Eur.Phys.J.C 80, 410 (2020)].

[19] A. Keshavarzi, D. Nomura, and T. Teubner, "The $g-2$ of charged leptons, $\alpha(M_Z^2)$ and the hyperfine splitting of muonium," *Phys. Rev.* **D101** (2020) 014029, `arXiv:1911.00367 [hep-ph]`.

[20] A. Kurz *et al.*, "Hadronic contribution to the muon anomalous magnetic moment to next-to-next-to-leading order," *Phys. Lett.* **B734** (2014) 144–147, `arXiv:1403.6400 [hep-ph]`.

[21] K. Melnikov and A. Vainshtein, "Hadronic light-by-light scattering contribution to the muon anomalous magnetic moment revisited," *Phys. Rev.* **D70** (2004) 113006, `arXiv:hep-ph/0312226 [hep-ph]`.

[22] P. Masjuan and P. Sánchez-Puertas, "Pseudoscalar-pole contribution to the $(g_\mu - 2)$: a rational approach," *Phys. Rev.* **D95** (2017) no. 5, 054026, `arXiv:1701.05829 [hep-ph]`.

[23] G. Colangelo *et al.*, "Dispersion relation for hadronic light-by-light scattering: two-pion contributions," *JHEP* **04** (2017) 161, `arXiv:1702.07347 [hep-ph]`.

[24] M. Hoferichter *et al.*, "Dispersion relation for hadronic light-by-light scattering: pion pole," *JHEP* **10** (2018) 141, `arXiv:1808.04823 [hep-ph]`.

[25] A. Gérardin, H. B. Meyer, and A. Nyffeler, "Lattice calculation of the pion transition form factor with $N_f = 2 + 1$ Wilson quarks," *Phys. Rev.* **D100** (2019) no. 3, 034520, `arXiv:1903.09471 [hep-lat]`.

[26] J. Bijnens, N. Hermansson-Truedsson, and A. Rodríguez-Sánchez, "Short-distance constraints for the HLbL contribution to the muon anomalous magnetic moment," *Phys. Lett.* **B798** (2019) 134994, `arXiv:1908.03331 [hep-ph]`.

[27] G. Colangelo *et al.*, "Longitudinal short-distance constraints for the hadronic light-by-light contribution to (g-2)$_\mu$ with large-N_c Regge models," *JHEP* **03** (2020) 101, `arXiv:1910.13432 [hep-ph]`.

[28] V. Pauk and M. Vanderhaeghen, "Single meson contributions to the muon's anomalous magnetic moment," *Eur. Phys. J.* **C74** (2014) no. 8, 3008, `arXiv:1401.0832 [hep-ph]`.

[29] I. Danilkin and M. Vanderhaeghen, "Light-by-light scattering sum rules in light of new data," *Phys. Rev.* **D95** (2017) no. 1, 014019, `arXiv:1611.04646 [hep-ph]`.

[30] F. Jegerlehner, "The Anomalous Magnetic Moment of the Muon," *Springer Tracts Mod. Phys.* **274** (2017) 1–693.

[31] M. Knecht *et al.*, "Scalar meson contributions to a_μ from hadronic light-by-light scattering," *Phys. Lett.* **B787** (2018) 111–123, `arXiv:1808.03848 [hep-ph]`.

[32] G. Eichmann, C. S. Fischer, and R. Williams, "Kaon-box contribution to the anomalous magnetic moment of the muon," *Phys. Rev.* **D101** (2020) no. 5, 054015, `arXiv:1910.06795 [hep-ph]`.

[33] P. Roig and P. Sánchez-Puertas, "Axial-vector exchange contribution to the hadronic light-by-light piece of the muon anomalous magnetic moment," *Phys. Rev.* **D101** (2020) no. 7, 074019, `arXiv:1910.02881 [hep-ph]`.

[34] T. Blum *et al.*, "The hadronic light-by-light scattering contribution to the muon anomalous magnetic moment from lattice QCD," *Phys. Rev. Lett.* **124** (2020) no. 13, 132002, `arXiv:1911.08123 [hep-lat]`.

[35] G. Colangelo *et al.*, "Remarks on higher-order hadronic corrections to the muon $g-2$," *Phys. Lett.* **B735** (2014) 90–91, `arXiv:1403.7512 [hep-ph]`.

[36] S. Laporta, "High-precision calculation of the 4-loop contribution to the electron g-2 in qed," *Physics Letters B* **772** (2017) 232–238. `https://www.sciencedirect.com/science/article/pii/S0370269317305324`.

[37] R. H. Parker *et al.*, "Measurement of the fine-structure constant as a test of the standard model," *Science* **360** (2018) no. 6385, 191–195. `https://science.sciencemag.org/content/360/6385/191`.

[38] L. Morel *et al.*, "Determination of the fine-structure constant with an accuracy of 81 parts per trillion," *Nature* **588** (2020) no. 7836, 61–65. `https://doi.org/10.1038/s41586-020-2964-7`.

[39] D. Hanneke, S. Fogwell, and G. Gabrielse, "New Measurement of the Electron Magnetic Moment and the Fine Structure Constant," *Phys. Rev. Lett.* **100** (2008) 120801, `arXiv:0801.1134 [physics.atom-ph]`.

[40] J. Prades, E. de Rafael, and A. Vainshtein, "The hadronic light-by-light scattering contribution to the muon and electron anomalous magnetic moments," *Advanced Series on Directions in High Energy Physics* (Dec, 2009) 303?317. `http://dx.doi.org/10.1142/9789814271844_0009`.

[41] F. Jegerlehner and A. Nyffeler, "The muon g?2," *Physics Reports* **477** (Jun, 2009) 1?110. `http://dx.doi.org/10.1016/j.physrep.2009.04.003`.

[42] G. Colangelo *et al.*, "Dispersion relation for hadronic light-by-light scattering: theoretical foundations," *JHEP* **09** (2015) 074, `arXiv:1506.01386 [hep-ph]`.

[43] A. Fienberg, *Measuring the Precession Frequency in the E989 Muon $g-2$ Experiment.* PhD thesis, University of Washington, 2019.

[44] W. D. Phillips, W. E. Cooke, and D. Kleppner, "Magnetic moment of the proton in H_2O in bohr magnetons," *Metrologia* **13** (1977) no. 4, 179–195. `https://doi.org/10.1088/0026-1394/13/4/005`.

[45] P. J. Mohr, D. B. Newell, and B. N. Taylor, "CODATA Recommended Values of the Fundamental Physical Constants: 2014," *Rev. Mod. Phys.* **88** (2016) no. 3, 035009, `arXiv:1507.07956 [physics.atom-ph]`.

[46] W. 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** (1999) 711–714.

[47] **Muon** $g-2$ Collaboration, T. Albahri *et al.*, "Measurement of the anomalous precession frequency of the muon in the Fermilab Muon $g-2$ experiment," *Phys. Rev. D* **103** (Apr, 2021) 072002. `https://link.aps.org/doi/10.1103/PhysRevD.103.072002`.

[48] **Muon** $g-2$ Collaboration, T. Albahri *et al.*, "Magnetic Field Analysis for the Muon $g-2$ Experiment at Fermilab," *Phys. Rev. A* **103** (Apr, 2021) 042208. `https://link.aps.org/doi/10.1103/PhysRevA.103.042208`.

[49] **Muon** $g-2$ Collaboration, T. Albahri *et al.*, "Beam dynamics corrections to the Run $-$ 1 measurement of the muon anomalous magnetic moment at Fermilab," *Phys. Rev. Accel. Beams* **24** (Apr, 2021) 044002. `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.24.044002`.

[50] D. Stratakis *et al.*, "Accelerator performance analysis of the Fermilab Muon Campus," *Phys. Rev. Accel. Beams* **20** (2017) no. 11, 111003, `arXiv:1803.00597 [physics.acc-ph]`.

[51] **Particle Data Group** Collaboration, A. Zyla *et al.*, "Review of Particle Physics," *Progress of Theoretical and Experimental Physics* **2020** (08, 2020) , `https://academic.oup.com/ptep/article-pdf/2020/8/083C01/34673722/ptaa104.pdf`. `https://doi.org/10.1093/ptep/ptaa104`. 083C01.

[52] **Muon** $g-2$ Collaboration, J. Grange *et al.*, "Muon $(g-2)$ Technical Design Report," `arXiv:1501.06858 [physics.ins-det]`.

[53] A. P. Schreckenberger *et al.*, "The fast non-ferric kicker system for the Muon $g-2$ Experiment at Fermilab," `arXiv:2104.07805 [physics.ins-det]`.

[54] R. Osofsky, *Magnetic Field Determination for Run 1 of the Fermilab Muon $g-2$ Experiment*. PhD thesis, University of Washington, 2019.

[55] M. Smith, *Developing the Precision Magnetic Field for the E989 Muon $g-2$ Experiment*. PhD thesis, University of Washington, 2017.

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

[57] R. Hong *et al.*, "Systematic and statistical uncertainties of the hilbert-transform based high-precision fid frequency extraction method," *J. Magn. Reson* **XYZ** (2021) xyx, `arXiv:2101.08412 [physics.ins-det]`.

[58] M. Farooq *et al.*, "Absolute Magnetometry with He3," *Phys. Rev. Lett.* **124** (2020) no. 22, 223001.

[59] **Muon** $g-2$ Collaboration, K. S. Khaw *et al.*, "Performance of the Muon $g-2$ calorimeter and readout systems measured with test beam data," *Nucl. Instrum. Meth.* **A945** (2019) 162558, `arXiv:1905.04407 [physics.ins-det]`.

[60] J. Kaspar *et al.*, "Design and performance of SiPM-based readout of PbF_2 crystals for high-rate, precision timing applications," *JINST* **12** (2017) no. 01, P01009, `arXiv:1611.03180 [physics.ins-det]`.

[61] A. Fienberg, "Second temperature scan." Elog entry, December, 2015. `https://muon.npl.washington.edu/elog/g2/UWg2Lab/83`.

[62] A. T. Fienberg *et al.*, "Studies of an array of PbF_2 Cherenkov crystals with large-area SiPM readout," *Nucl. Instrum. Meth.* **A783** (2015) 12–21, `arXiv:1412.5525 [physics.ins-det]`.

[63] **Muon** $g-2$ Collaboration, A. Anastasi *et al.*, "The laser-based gain monitoring system of the calorimeters in the Muon $g-2$ experiment at Fermilab," *JINST* **14** (2019) P11025, `arXiv:1906.08432 [physics.ins-det]`.

[64] A. Driutti *et al.*, "The calorimeter gain sag(a): Chapter 3 - the tests." Muon $g-2$ internal note, August, 2018. Muon $g-2$ DocDB: 13537.

[65] **Muon (**$g-2$**)** Collaboration, G. W. Bennett *et al.*, "Statistical equations and methods applied to the precision muon $(g-2)$ experiment at BNL," *Nucl. Instrum. Meth. A* **579** (2007) 1096–1116.

[66] M. Abe *et al.*, "A new approach for measuring the muon anomalous magnetic moment and elec-
tric dipole moment,"*Progress of Theoretical and Experimental Physics* **2019** (05, 2019) , `https://academic.oup.com/ptep/article-pdf/2019/5/053C02/28746337/ptz030.pdf`. `https://doi.org/10.1093/ptep/ptz030`. 053C02.

[67] S. Borsanyi *et al.*, "Leading hadronic contribution to the muon magnetic moment from lattice qcd," *Nature* **593** (2021) no. 7857, 51–55. `https://doi.org/10.1038/s41586-021-03418-1`.

[68] G. Abbiendi *et al.*, “Measuring the leading hadronic contribution to the muon $g-2$ via μe scattering,” *The European Physical Journal C* **77** (2017) no. 3, 139. `https://doi.org/10.1140/epjc/s10052-017-4633-z`.

[69] C. M. Carloni Calame *et al.*, “A new approach to evaluate the leading hadronic corrections to the muon $g-2$,” *Physics Letters B* **746** (2015) 325–329. `https://www.sciencedirect.com/science/article/pii/S0370269315003573`.

[70] A. Czarnecki and W. J. Marciano, “The muon anomalous magnetic moment: A harbinger for ‘new physics’,” *Phys. Rev.* **D64** (2001) 013014, `arXiv:hep-ph/0102122`.

[71] P. Athron *et al.*, “New physics explanations of a_μ in light of the FNAL muon $g-2$ measurement,” `arXiv:2104.03691 [hep-ph]`.

[72] **ACME** Collaboration, V. Andreev *et al.*, “Improved limit on the electric dipole moment of the electron,” *Nature* **562** (2018) no. 7727, 355–360.

[73] A. Crivellin, M. Hoferichter, and P. Schmidt-Wellenburg, “Combined explanations of $(g-2)_{\mu,e}$ and implications for a large muon EDM,” *Phys. Rev. D* **98** (2018) no. 11, 113002, `arXiv:1807.11484 [hep-ph]`.

[74] **Mu2e** Collaboration, L. Bartoszek *et al.*, “Mu2e Technical Design Report,” `arXiv:1501.05241 [physics.ins-det]`.

[75] **Mu3e** Collaboration, K. Arndt *et al.*, “Technical design of the phase I Mu3e experiment,” `arXiv:2009.11690 [physics.ins-det]`.

[76] **MEG II** Collaboration, A. M. Baldini *et al.*, “The design of the MEG II experiment,” *Eur. Phys. J. C* **78** (2018) no. 5, 380, `arXiv:1801.04688 [physics.ins-det]`.

Appendix A

CALORIMETER DIAGRAMS

The diagrams in this appendix are included as a means of record keeping, as they might otherwise be lost to the sands of time. They include schematic diagrams of various components of the calorimeter electronics chain: the SiPM and on-board circuitry, the breakout boards used to fan out communication and bias voltage, and the pin allocations on the custom-made cape for the BeagleBone microcomputer. Each was designed by the Tim Van Wechel and David Peterson in the electronics lab at CENPA.

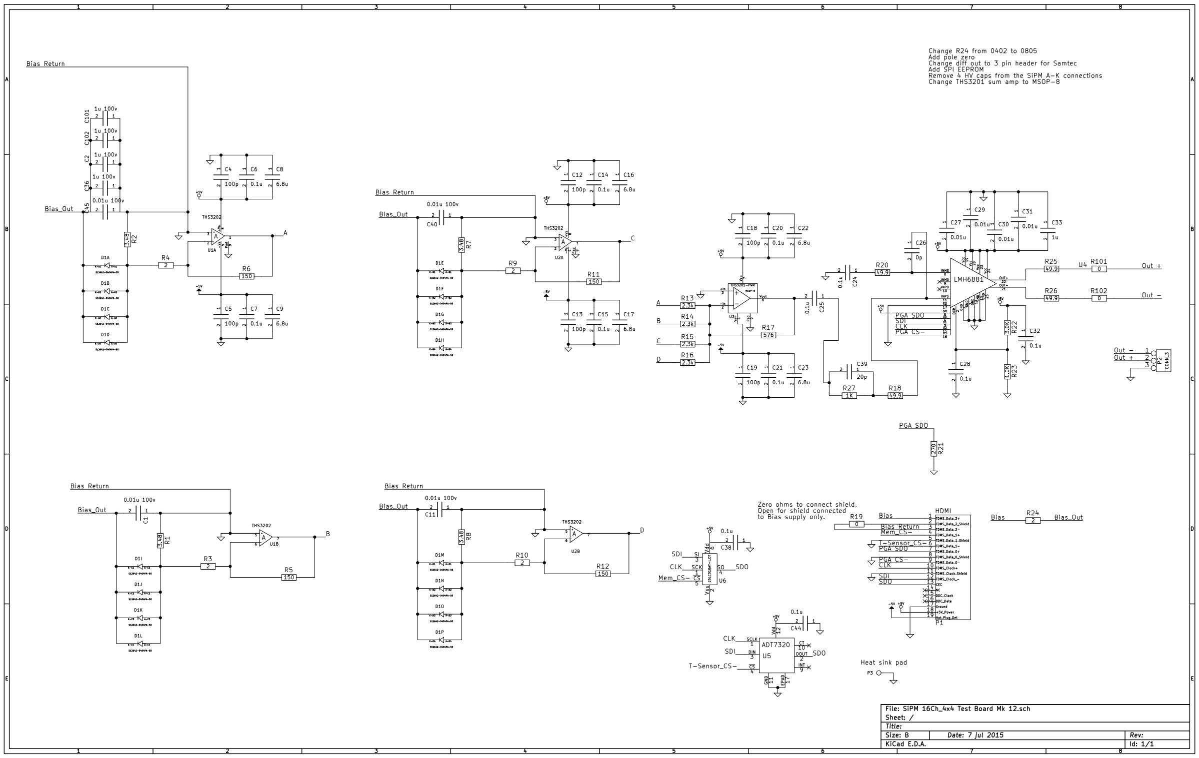

Figure A.1: The signal summing logic can be gleaned as well as the allocation of the output signals and HDMI communication.

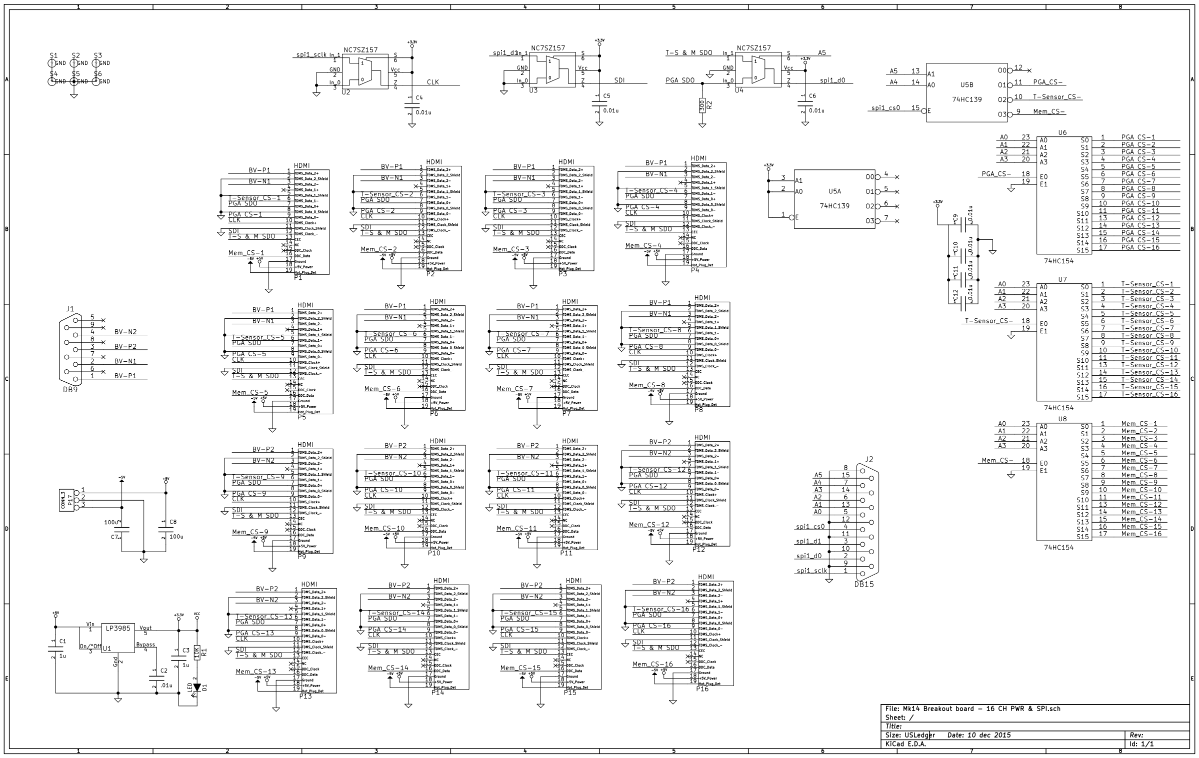

Figure A.2: Shown is the diagram of how communication and power is distributed to 16 SiPMs via the breakout boards in each calorimeter.

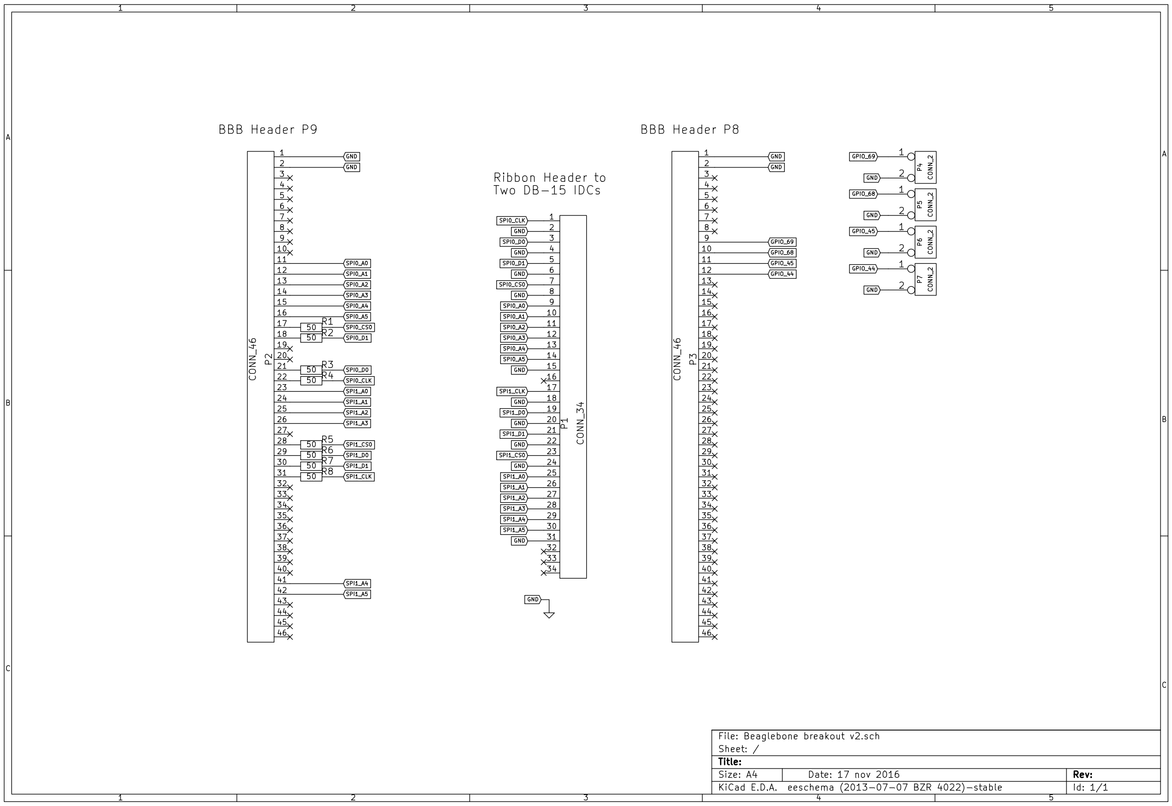

Figure A.3: SPI pin allocations on the top of the BeagleBone Black microcomputer that provides communication with the SiPMs. The connections are important when writing code that provides communication with the SiPMs and associated electronics.

Appendix B

PHOTOSTATISTICS

Start by assuming that the number of photons in a particular laser shot is governed by a Gaussian with mean $\bar{\gamma}$ and variance σ_γ^2, the number of photons fired by the laser, γ, follows this probability distribution function:

$$P(\gamma) = \frac{1}{\sigma_\gamma \sqrt{2\pi}} \exp\left(-\frac{(\gamma - \bar{\gamma})^2}{2\sigma_\gamma^2}\right) \tag{B.1}$$

This is a continuous approximation assuming a large number of photons. Now, assume that the number of photons transmitted through a filter of fractional transmittance f is governed by a binomial distribution. Once again using a continuous approximation, the probability of n, the number of photons transmitted, given γ and f is:

$$P\left(n|\gamma, f\right) = \frac{1}{\sqrt{2\pi f(1-f)\gamma}} \exp\left(-\frac{(n - f\gamma)^2}{2f(1-f)\gamma}\right) \tag{B.2}$$

Combining these two together, the probability that n photons are transmitted to a SiPM is given by:

$$P(n) = \int \frac{d\gamma}{2\pi\sigma_\gamma \sqrt{f(1-f)\gamma}} \exp\left(-\frac{(n - f\gamma)^2}{2f(1-f)\gamma} - \frac{(\gamma - \bar{\gamma})^2}{2\sigma_\gamma^2}\right) \tag{B.3}$$

To simplify the math, the γ appearing in the variance term $f(1-f)\gamma$ will be replaced by $\bar{\gamma}$. Given the stability of the lasers, this should be a good approximation.

$$\begin{aligned}
P(n) = & \int \frac{d\gamma}{2\pi\sigma_\gamma\sqrt{f(1-f)\bar{\gamma}}} \exp\left(-\frac{(n-f\gamma)^2}{2f(1-f)\bar{\gamma}} - \frac{(\gamma-\bar{\gamma})^2}{2\sigma_\gamma^2}\right) && \text{(B.4)} \\
= & \frac{1}{2\pi\sigma_\gamma\sqrt{f(1-f)\bar{\gamma}}} \int d\gamma \exp\Bigg(-\gamma^2 \frac{f^2\sigma_\gamma^2 + f(1-f)\bar{\gamma}}{2f(1-f)\bar{\gamma}\sigma_\gamma^2} + && \text{(B.5)} \\
& \gamma \frac{nf\sigma_\gamma^2 + \bar{\gamma}^2 f(1-f)}{f(1-f)\bar{\gamma}\sigma_\gamma^2} - \frac{n^2\sigma_\gamma^2 + \bar{\gamma}^2 f(1-f)\bar{\gamma}}{2f(1-f)\bar{\gamma}\sigma_\gamma^2}\Bigg) && \text{(B.6)} \\
= & \frac{1}{2\pi\sigma_\gamma\sqrt{f(1-f)\bar{\gamma}}} \exp\left(-\frac{n^2\sigma_\gamma^2 + \bar{\gamma}^2 f(1-f)\bar{\gamma}}{2f(1-f)\bar{\gamma}\sigma_\gamma^2}\right) \sqrt{\frac{2\pi f(1-f)\bar{\gamma}\sigma_\gamma^2}{f^2\sigma_\gamma^2 + f(1-f)\bar{\gamma}}} \times && \text{(B.7)} \\
& \exp\left(\left(\frac{nf\sigma_\gamma^2 + \bar{\gamma}^2 f(1-f)}{f(1-f)\bar{\gamma}\sigma_\gamma^2}\right)^2 \frac{2f(1-f)\bar{\gamma}\sigma_\gamma^2}{4f^2\sigma_\gamma^2 + f(1-f)\bar{\gamma}}\right) && \text{(B.8)} \\
= & \frac{1}{\sqrt{2\pi\left(f^2\sigma_\gamma^2 + f(1-f)\bar{\gamma}\right)}} \exp\left(\frac{-(n-f\bar{\gamma})^2}{2\left(f^2\sigma_\gamma^2 + f(1-f)\bar{\gamma}\right)}\right) && \text{(B.9)}
\end{aligned}$$

This last line indicates that the distribution of photons arriving at the SiPM will be Gaussian with mean $f\bar{\gamma}$ and variance $f^2\sigma_\gamma^2 + f(1-f)\bar{\gamma}$.

Next, assume that the efficiency of detecting a photon is e. This will once again be taken as the continuous approximation of the binomial distribution. The probability that a pixels fired (assuming the number of photons arriving at the SiPM is much less than the total number of pixels) for a given number of photons n is then:

$$P(a|n,e) = \frac{1}{\sqrt{2\pi e(1-e)n}} \exp\left(-\frac{(a-en)^2}{2e(1-e)n}\right) \tag{B.10}$$

Following the same math as above, the number of fired pixels will be Gaussian with mean given by $\bar{a} = e\bar{n} = ef\bar{\gamma}$ and variance given by

$$\sigma_a^2 = e^2\sigma_n^2 + e(1-e)\bar{n} = e^2(f^2\sigma_\gamma^2 + f(1-f)\bar{\gamma}) + e(1-e)f\bar{\gamma} = (ef)^2\sigma_\gamma^2 + ef(1-ef)\bar{\gamma}. \tag{B.11}$$

If the gain is linear, the output signal will be linearly proportional to the number of pixels fired, $S(a) = ga$. The distribution of output signals will follow the distribution of fired pixels:

$$1 = \int_{-\infty}^{\infty} da P(a) \tag{B.12}$$

$$= \int_{-\infty}^{\infty} da \frac{1}{\sigma_a \sqrt{2\pi}} \exp\left(-\frac{(a-\bar{a})^2}{2\sigma_a^2}\right) \tag{B.13}$$

$$= \int_{-\infty}^{\infty} \frac{dS}{g} \frac{1}{\sigma_a \sqrt{2\pi}} \exp\left(-\frac{(ga-g\bar{a})^2}{2g^2\sigma_a^2}\right) \tag{B.14}$$

$$= \int_{-\infty}^{\infty} dS \frac{1}{g\sigma_a \sqrt{2\pi}} \exp\left(-\frac{(S-g\bar{a})^2}{2g^2\sigma_a^2}\right). \tag{B.15}$$

So, the output signal will also be Gaussian with mean $\bar{S} = gef\bar{\gamma}$ and variance $\sigma_S^2 = g^2\left((ef)^2\sigma_\gamma^2 + ef(1-ef)\bar{\gamma}\right) = g\bar{S} + \bar{S}^2\left(\frac{\sigma_\gamma^2}{\bar{\gamma}^2} - \frac{1}{\bar{\gamma}}\right)$. The linear term dominates in the region where there are a large number of photons reaching the SiPM.

The upshot of all this math is that the coefficient of the linear term of the signal variance vs. mean plot is the gain of a particular SiPM. For this reason, it is not vital to know exactly the detection efficiency of the SiPM or fractional transmittance of the filters. Given several filter wheel positions, the variance and mean of the SiPM response to laser pulses can be extracted at each position. Once plotted, the coefficient of the first-order term will be the desired gain value.

Appendix C

CALCULATING THE χ^2 FOR KERNELIZED DATA

To avoid cyclotron aliasing effects entering into the calculation of ω_a, particularly at the individual calorimeter level, each positron hit is smeared uniformly across a full cyclotron period in the histogram that is ultimately fit. Thus, each hist still contributions a total weight of 1 to the T-method fit, but any coherent effects from the decoherence of the cyclotron motion are diminished. The downside is that this will introduce correlations between time bins that must be accounted for in the formulation of the χ^2 used when fitting.

Before getting into the calculation of the bin correlations, some notation for how the kernel method works will be presented. Even though the final analysis histogram is binned in cyclotron period (T_c)-sized bins, it will be convenient throughout this discussion to use $T_c/2$-sized bins. For the nominal data (no smearing performed), the full cyclotron bins $N_i = A_{i,1} + A_{i,2}$ where the $A_{i,j}$ s denote half-cyclotron bins. The contribution of each hit to the kernelized half-cyclotron bins $B_{i,j}$ is as follows:

$$w_{i,1}(t) = \begin{cases} 1/2 + \frac{t - iT_c}{T_c} & \text{if } (i-1)T_c + \frac{T_c}{2} \leq t < iT_c \\ 1/2 & \text{if } iT_c \leq t \leq iT_c + \frac{T_c}{2} \\ 1 - \frac{t - iT_c}{T_c} & \text{if } iT_c + \frac{T_c}{2} < t \leq (i+1)T_c \\ 0 & \text{otherwise.} \end{cases} \tag{C.1}$$

This can be shifted to get a similar formulation for the weights contributing to $B_{i,2}$. If the rate across each bin is assumed to be constant, this leads to a bin content of:

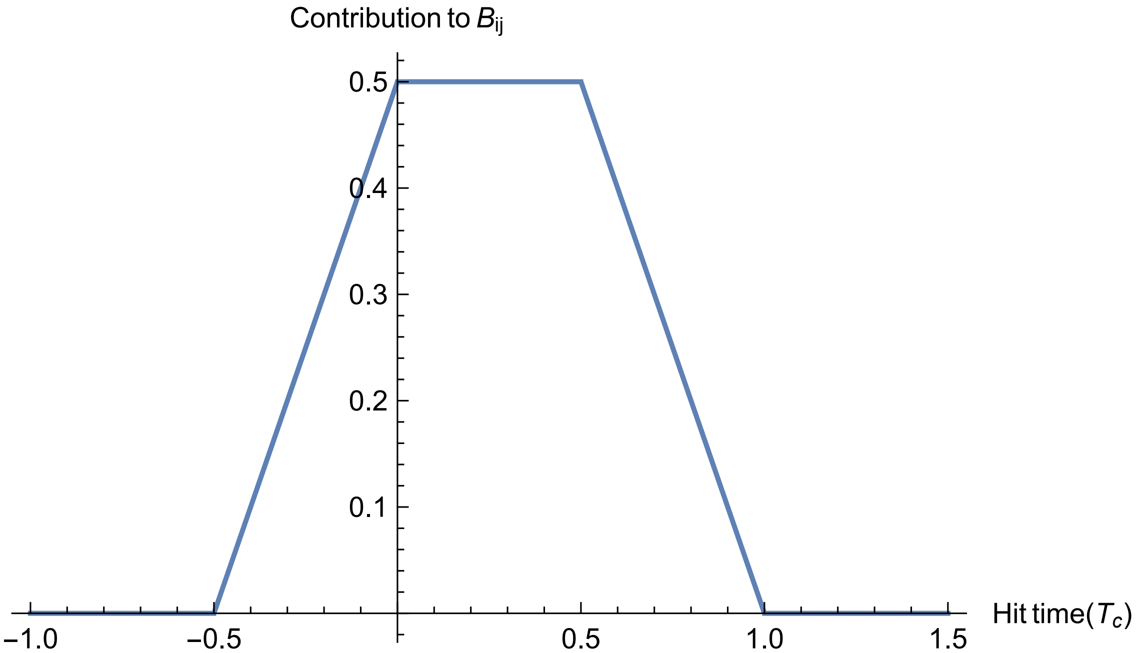

Figure C.1: The contribution from a single hit at time t, given in units of the cyclotron period, is shown.

$$
\begin{aligned}
B_{i,1} &= \int w_{i,1}(t)R(t)dt && \text{(C.2)}\\
&= \int_{(i-1)T_c+\frac{T_c}{2}}^{iT_c} \left(\frac{1}{2}+\frac{t-iT_c}{T_c}\right)\frac{A_{i-1,2}}{T_c}dt \qquad + \int_{iT_c}^{iT_c+\frac{T_c}{2}} \left(\frac{1}{2}\right)\frac{A_{i,1}}{T_c}dt && \text{(C.3)}\\
&\qquad + \int_{iT_c+\frac{T_c}{2}}^{(i+1)T_c} \left(1-\frac{t-iT_c}{T_c}\right)\frac{A_{i,2}}{T_c}dt && \text{(C.4)}\\
&= \frac{1}{4}A_{i-1,2}+\frac{1}{2}A_{i,1}+\frac{1}{4}A_{i,2} && \text{(C.5)}
\end{aligned}
$$

This can once again be generalized to the other half-cyclotron bin $B_{i,2}$ very easily. When combined into the bins used for fitting, this becomes:

$$
\begin{aligned}
K_i &= B_{i,1} + B_{i,2} && \text{(C.6)}\\
&= \int w_{i,1}(t)R(t)dt + \int w_{i,2}(t)R(t)dt && \text{(C.7)}\\
&= \int \left(w_{i,1}(t) + w_{i,2}(t)\right) R(t)d && \text{(C.8)}\\
&= \int_{(i-1)T_c+\frac{T_c}{2}}^{iT_c} \left(\frac{1}{2} + \frac{t-iT_c}{T_c}\right) \frac{A_{i-1,2}}{T_c}dt + \int_{iT_c}^{iT_c+\frac{T_c}{2}} \left(\frac{1}{2} + \frac{t-(iT_c+T_c/2)}{T_c}\right) \frac{A_{i,1}}{T_c}dt && \text{(C.9)}\\
&+ \int_{iT_c}^{(i+1)T_c} \left(1 - \frac{t-iT_c}{T_c} + \frac{1}{2}\right) \frac{A_{i,2}}{T_c}dt + \int_{(i+1)T_c}^{(i+1)T_c+\frac{T_c}{2}} \left(-t/T\right) \frac{A_{i+1,1}}{T_c}dt && \text{(C.10)}\\
&= \frac{1}{4}A_{i-1,2} + \left(\frac{1}{2} + \frac{1}{4}\right) A_{i,1} + \left(\frac{1}{4} + \frac{1}{2}\right) A_{i,2} + \frac{1}{4}A_{i+1,1} && \text{(C.11)}\\
&= \frac{1}{4}A_{i-1,2} + \frac{3}{4}\left(A_{i,1} + A_{i,2}\right) + \frac{1}{4}A_{i+1,1} && \text{(C.12)}\\
&= \frac{1}{8}N_{i-1} + \frac{3}{4}N_i + \frac{1}{8}N_{i+1}. && \text{(C.13)}
\end{aligned}
$$

The last step was made by assuming that the rate in each half of the cyclotron bin is equal ($A_{i,1} = A_{i,2} = N_i/2$). This is to show that, as stated above, the kernel has introduced correlations between time bins that must be accounted for when fitting.

To include these correlations, the χ^2 used for fitting is formulated as

$$
\chi^2 = \sum_i \sum_j \left(K_i - F(t_i)\right) \mathrm{Cov}^{-1}(K)_{ij} \left(K_j - F(t_j)\right), \tag{C.14}
$$

where K_i is the content of the kernelized bin, $F\left(t_i\right)$ is the fit function evaluated at the center of that bin, and $\mathrm{Cov}^{-1}(K)_{ij}$ is the inverse of the covariance matrix at position (i, j). Constructing the covariance matrix and taking the inverse are the most complicated part of this process.

If the number of counts in each time slice dt is taken to be a Poisson process, the variance in each slice will be the square of the weight times the number of counts from that slice. Thus,

$$\mathrm{Var}(K_i) = \int \left(w_{i,1}(t) + w_{i,2}(t)\right)^2 R(t)d \tag{C.15}$$

$$= \int_{(i-1)T_c+\frac{T_c}{2}}^{iT_c} \left(\frac{1}{2} + \frac{t - iT_c}{T_c}\right)^2 \frac{A_{i-1,2}}{T_c}dt + \int_{iT_c}^{iT_c+\frac{T_c}{2}} \left(\frac{1}{2} + \frac{t - (iT_c + T_c/2)}{T_c}\right)^2 \frac{A_{i,1}}{T_c}dt \tag{C.16}$$

$$+ \int_{iT_c}^{(i+1)T_c} \left(1 - \frac{t - iT_c}{T_c} + \frac{1}{2}\right)^2 \frac{A_{i,2}}{T_c}dt + \int_{(i+1)T_c}^{(i+1)T_c+\frac{T_c}{2}} \left(-t/T\right)^2 \frac{A_{i+1,1}}{T_c}dt \tag{C.17}$$

$$= \frac{1}{24}N_{i-1} + \frac{7}{12}N_i + \frac{1}{24}N_{i+1}. \tag{C.18}$$

The covariance between two adjacent bins K_i and K_{i+1} will involve the product of the weights that contribute to each:

$$\mathrm{Cov}(K_i, K_{i+1}) = \mathrm{Cov}(B_{i,2}, B_{i+1,1}) \tag{C.19}$$

$$= \int w_{i,2}(t)w_{i+1,1}(t)R(t)dt \tag{C.20}$$

$$= \int_{(i+1)T_c+\frac{T_c}{2}}^{(i+1)T_c+\frac{T_c}{2}} \frac{1}{2}\left(\frac{1}{2} + \frac{t - (i+1)T_c}{T_c}\right) \frac{A_{i,2}}{T_c}dt \tag{C.21}$$

$$+ \int_{iT_c+\frac{T_c}{2}}^{(i+1)T_c} \frac{1}{2}\left(\frac{1}{2} - \frac{t - (i+1)T_c}{T_c}\right) \frac{A_{i+1,1}}{T_c}dt \tag{C.22}$$

$$= \frac{1}{6}\left(A_{i,2} + A_{i+1,1}\right) \tag{C.23}$$

$$= \frac{1}{12}\left(N_i + N_{i+1}\right). \tag{C.24}$$

Thus, the covariance matrix:

$$\mathrm{Cov}(K)_{ij} = \begin{cases} \frac{1}{24}N_{i-1} + \frac{7}{12}N_i + \frac{1}{24}N_{i+1} & \text{if } i = j \\ \frac{1}{12}\left(N_i + N_j\right) & \text{if } |i-j| = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{C.25}$$

Notice that this matrix is sparse – the only non-zero elements are on the diagonal and the subdiagonals immediately above and below. This *tridiagonal* structure lends itself well to taking the inverse. There is a readily available algorithm for doing so that will be summarized below.

The tridiagonal algorithm, as developed by the same Thomas of Thomas precession, is a modified form of Gaussian elimination. The generic form of the equation to be solved is

$$\begin{bmatrix} b_1 & c_1 & 0 & & & \\ a_2 & b_2 & c_2 & 0 & & \\ 0 & a_3 & b_3 & c_3 & 0 & \\ & & & \ddots & & \\ & & 0 & a_{n-1} & b_{n-1} & c_{n-1} \\ & & & 0 & a_n & b_n \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_{n-1} \\ x_n \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \\ d_{n-1} \\ d_n \end{bmatrix} \tag{C.26}$$

where the empty entries are all zeroes. To use this equation to solve for the inverse of the covariance matrix laid out above, identify the following:

$$a_i = c_{i-1} = \frac{1}{12}\left(N_i + N_{i-1}\right) \tag{C.27}$$

$$b_i = \frac{1}{24}N_{i-1} + \frac{7}{12}N_i + \frac{1}{24}N_{i+1}. \tag{C.28}$$

The inverse matrix will be constructed by solving the above equation n times, each time identifying $\vec{x}$ as the j-th column of $\mathrm{Cov}^{-1}(K)$ and $\vec{d}$ as the j-th column of $\mathbb{1}_n$.

The tridiagonal algorithm is as follows. First, iterating for $k = 2, 3, ..., n$:

$$w = \frac{a_k}{b_{k-1}} \tag{C.29}$$

$$b_k \to b_k - wc_{k-1} \tag{C.30}$$

$$d_k \to d_k - wd_{k-1}. \tag{C.31}$$

Then:

$$x_n = \frac{d_n}{b_n} \tag{C.32}$$

$$x_k = \frac{d_k - c_k x_{k+1}}{b_k} \text{ for } k = n-1, n-2..., 1. \tag{C.33}$$

To avoid problems with the precision of doubles, any entries $|d_k| < 10^{-50} \to d_k = 0$ and similarly for values of x_k. This algorithm is remarkably fast while remaining accurate. For a benchmark comparison, the time to take the inverse of a 4000×4000 element matrix is approximately 7 seconds while the linear algebra library Eigen takes about 60 seconds. This speed-up is important when attempting to do studies involving fitting different calorimeters, energy ranges, or time intervals.