

Chapter 16

The ω_a Measurement

The anomalous spin precession frequency ω_a is one of the two observables required to obtain the muon anomalous magnetic moment, a_μ . In order to ensure that the experiment's proposed goal of 140 ppb precision in a_μ is achieved, the error budget allows for a 100 ppb statistical uncertainty combined with equal 70 ppb systematic uncertainties from each of the ω_a and ω_p analyses. This chapter summarizes the procedure for the ω_a measurement, with Chapters 17-22 elaborating upon the design of each subsystem. First the decay kinematics are reviewed and the encoding of the muon spin information into the data set is explained (Section 16.1). From that basis several complementary analysis methods are described (16.1.2). Then a review of uncertainties is presented, first the statistical (Section 16.2) and then the systematic uncertainties intrinsic to the detectors (Section 16.3). Then the detector system organization is outlined and broken down into subsystems (Section 16.4). These subsystems include stored muon monitoring, decay positron tracking, electromagnetic calorimeter, signal digitization, data acquisition, and slow control systems.

16.1 Measurement Overview

In this experiment the polarized positive muons are stored in a magnetic ring. Their spins precess at a different rate than their momenta. The anomalous precession frequency, ω_a , is the difference between the ensemble-averaged muon spin precession and cyclotron frequencies. Direct measurement of the muon spin is not practicable so an indirect measurement is made via the decay positrons. Muon decay proceeds through the weak force and therefore is parity violating. The consequence of this behavior is that the emitted positron momentum is correlated with the muon spin direction. Therefore by measuring the decay positrons and analyzing their energies, a measurement of the muon spin is possible. We have elected to use calorimeters to measure the positron energy and time. The following sections describe the details of the muon decay pertinent to understanding the design of the calorimeters as well as the possible methods of analysis.




Figure 16.1: The decay positron energy spectrum as a function of time, modulo a complete $4362\text{-ns}(g-2)$ period. The muon spin and momentum are aligned at $\pi/2$ and anti-aligned at $3\pi/2$ in this figure, corresponding to about 1090 and 3271 ns, respectively.

16.1.1 Muon Decay and Boost Kinematics

In this section we summarize the most important aspects of muon decay with respect to the detector design. For a comprehensive discussion of the kinematics of muon decay, see Section 3.5. Starting in the muon rest frame, the angular distribution of emitted positrons from an ensemble of polarized muons is

$$dn/d\Omega = 1 + a(E) (\hat{S}_\mu \cdot \hat{P}_e), \quad (16.1)$$

where \hat{S}_μ is the muon spin direction and \hat{P}_e is the positron momentum direction. The asymmetry a depends on positron energy (E) and is such that the higher-energy positrons are emitted parallel to the muon spin. To boost to the laboratory frame we define θ^* as the angle between the positron momentum and the Lorentz boost,

$$E_{e,lab} = \gamma(E_e^* + \beta P_e^* \cos \theta^*) \approx \gamma E_e^* (1 + \cos \theta^*). \quad (16.2)$$

The starred quantities indicate the CM frame. The magic momentum requirement fixes γ at 29.3. Due to the correlation between the muon spin and the positron momentum direction the angle between the positron momentum and the boost direction from the muon center-of-mass frame (CM) to the lab frame acts as an analyzer of the muon spin. The maximum positron energy in the lab frame occurs when the positron decay energy (E_e^*) is the maximum and the positron momentum is aligned with the boost direction ($\cos \theta^* = 1$). Figure 16.1 shows the decay positron energy spectrum as a function of time for one $(g-2)$ period.

The modulation of the decay energy spectrum occurs at the frequency ω_a . Therefore by measuring this modulation a precision measurement of ω_a is possible.

16.1.2 Analysis Methods Summary

The standard analysis procedure is to identify individual decay positrons and plot the rate of their arrival versus time using only events having a measured energy above a threshold. The top panel of Figure 16.2 shows the result of this analysis method applied to simulation data. This method is named the T (time) method; it was the dominant analysis technique used in the Brookhaven experiment and it is well tested against systematic errors.

The rate of detected positrons above a single energy threshold E_{th} is

$$\frac{dN(t; E_{th})}{dt} = N_0 e^{-t/\gamma\tau_\mu} [1 + A \cos(\omega_a t + \phi)]. \quad (16.3)$$

Here the normalization, N_0 , average asymmetry, A , and initial phase, ϕ , are all dependent on the threshold energy. A parameterization of this function is used to fit the results from the T Method analysis and extract ω_a . The T method is sufficient to reach the experimental goal. The method is well understood and proven to work due to its implementation by the E821 collaboration.

However it is possible to extract more statistical precision from the data set. The information of the muon spin is encoded via the positron momentum. By weighting events in proportion to their energy, or the asymmetry associated with their energy, the statistical precision is improved. As in the T method, the data stream from the calorimeters must be first deconstructed into individual events and then binned in time before being fit to extract ω_a .

An alternative approach is to not identify individual decay positrons but rather to digitize the detector current vs. time, which is proportional to the energy deposited in the calorimeter vs. time from the decay positrons. From the relationship depicted in Figure 16.1, the energy in the detector will oscillate at the $g - 2$ frequency. That oscillation can then be fit to determine ω_a . This approach is called the Q (charge) method and it is inherently energy weighted, with a near-zero threshold. It was not used in E821 due to memory and readout limitations, but will be implemented in E989. This method has the attractive feature of being immune to the systematic effect of pileup (described in Section 16.3.2).

An example of this the spectrum that would result from a T or Q method analysis using the same simulated data set is shown in Figure 16.2. We summarize and compare the important features of the three T -based and the Q -based methods:

- **T Method:** Events in the calorimeter are individually identified, sorted and fit to obtain time and energy. The events vs. time-in-fill histogram is built from all events with reconstructed energy above a threshold. All events in the histogram are given equal weight. The figure-of-merit (FOM) is maximized for a positron energy threshold of 1.86 GeV, as discussed below. The quantity ω_a is obtained from a fit to a pileup-subtracted histogram. This is the standard method used in E821 and the benchmark for determining the statistical and systematic requirements for the E989 experiment.
- **E -Weighted Method:** Identical to the threshold T method except that the histogram is built by incrementing a time bin with a weight equal to the energy of an event, therefore producing the Energy vs. Time-in-Fill histogram.




Figure 16.2: Top: Monte Carlo data analyzed using the T method with a threshold cut at $y = 0.6$. Bottom: Same data analyzed using the Q method. Detector acceptance is included. The asymmetry A is much higher for the T method; however, the Q method has many more events (N). The ω_a Monte-Carlo truth is $R = 0$ and the uncertainty in R is a measure of the precision, in ppm. Both methods give a similar statistical uncertainty and acceptable fit central values.

- **A-Weighted Method:** Identical to the threshold T method except that the histogram is built by incrementing a time bin with the average value of the asymmetry corresponding to that positron energy. This technique yields the maximum possible statistical power for a given threshold using the event identification technique.
- **Q Method:** An alternative approach whereby the detector current—a proxy for the deposited energy—is digitized and plotted as a function of time. Individual positron events *are not* identified. This procedure leads to a histogram of energy vs. time-in-fill. No attempt to correct for pileup is necessary here and a very low threshold is desired. The statistical power of this approach almost reaches that of the threshold T method.

All of the mentioned analysis techniques will be applied to the data set by different research groups throughout the collaboration. The analysis will be conducted in a blind fashion. Next we provide an overview of the statistical and systematic uncertainties. We conclude the chapter by giving an overview of the detector systems (full details are given in subsequent chapters).

16.2 Statistical Uncertainty

The T and Q methods lead to similar histograms (see Figure 16.2) with different bin weights and asymmetries. A fit is performed using Equation 16.3 and the relevant parameter, ω_a , is obtained. The optimization of the experimental system follows from minimizing the uncertainty on that parameter, namely $\delta\omega_a$. A detailed study [1] of the statistical methods gives guidance to the statistical power of any data set built using various weighting methods. The uncertainty on ω_a can be parameterized as

$$\delta\omega_a = \sqrt{\frac{2}{N(\gamma\tau_\mu)^2} \cdot \frac{\langle p^2 \rangle_y}{\langle pA \rangle_y^2}}, \quad (16.4)$$

where N is the integrated number of decay positrons in the analysis, p is the weight function and therefore is method dependent, A is the asymmetry, and $\langle f \rangle_y$ is the value of f averaged over all detected positron energies above threshold. The parameter y is the fractional decay positron energy with respect to a maximum value; therefore y ranges from 0 to 1, with $y = 1$ corresponding to approximately 3.1 GeV.

Figure 16.3 shows *differential* plots of N , A , and NA^2 vs. energy for a uniform acceptance detector. This plot illustrates the importance of the higher-energy positrons (those with the greatest asymmetry). The asymmetry is negative for lower-energy positrons; thus, a single low threshold can be expected to dilute the average asymmetry. The modification of the ideal curves due to the detector acceptance is significant, as the detector placement has been designed to greatly favor the higher-energy events. low energy positrons are more likely to curl between detectors and be missed. The acceptance impacts the values of N and A , which are functions of the energy-dependent detector acceptance. In the T method, each event carries the same weight ($p = 1$) and the uncertainty $\delta\omega_a$ (Eq. 16.4) reduces to

$$\delta\omega_a = \frac{1}{\gamma\tau_\mu} \sqrt{\frac{2}{NA^2}}. \quad (16.5)$$




Figure 16.3: The differential distributions: normalized number of events (N), asymmetry (A), and the figure of merit (NA^2). Note, this plot assumes uniform detector acceptance across the full energy spectrum.

The figure of merit (FOM) that should be maximized to minimize $(\delta\omega_a)^2$ is NA^2 . The value of the threshold that maximizes the FOM corresponds to $A \approx 0.4$ and an energy of 1.86 GeV. Therefore the relative uncertainty in ω_a is

$$\frac{\delta\omega_a}{\omega_a} = \frac{1}{\omega_a} \cdot \frac{\sqrt{2}}{\gamma\tau_\mu AP} \cdot \frac{1}{\sqrt{N}} \approx \frac{0.0398}{\sqrt{N}}. \quad (16.6)$$

The variable P represents the average polarization of the muons. The end-to-end beamline transport simulations project a value of 0.95 for the stored muon polarization. For a statistical uncertainty on $\delta\omega_a/\omega_a$ of 100 ppb, $N = 1.6 \cdot 10^{11}$ fitted events will be required. The T Method is sufficient for reaching the goal of E989 and all benchmarks and estimates are based solely on this method. However there is an opportunity for additional precision by incorporating other analysis techniques.

The Q method is an energy-weighted ($p = y$) analysis with a single very low threshold, since events do not need to be individually identified. The computation of $\langle p^2 \rangle_y / \langle pA \rangle_y^2$ in Eq. 16.4 is non-trivial. We conducted a simulation that included the detector acceptance to determine a realistic FOM for the Q - and the A - and E -weighted methods. Figure 16.4 shows the results vs. threshold energy. The meaningful FOM for the Q method corresponds to the near-zero threshold end of the curve, which is circled. It should be compared to the T method at its peak.

The Q method was not employed in E821 owing to the high energy threshold, lack of sufficient memory in the digitizers, and limitations on the transfer speed to the DAQ. These technical limitations are easily overcome with today's large memories in such devices and faster bus speeds. Note that the data sets in the T and Q method are not identical, but substantial overlap exists. For example, in the T method, all events below ~ 1.86 GeV do not contribute and all events above are weighted with $p = 1$. The Q method includes all events that strike the detector and weights each by its energy, $p = y$. Therefore, a combination of the results of the two methods will enable a small reduction in the final uncertainty of ω_a ; more importantly, the two methods will also serve as an important cross checks on the final result because they have very different sensitivities to various systematic uncertainties.




Figure 16.4: The figure-of-merit for the T ($p(y) = 1$), energy-weighted ($p(y) = y$), and asymmetry weighted ($p(y) = A(y)$) methods are plotted versus threshold energy. They were calculated from the output of a simulated data set including detector acceptance. Two points are of particular interest. The first is the maximum of the unweighted distribution, which occurs at $y=0.6$ (1.86 GeV), and is therefore used for the T method threshold. The second is the low-threshold limit for the energy-weighted distribution, which is circled in the figure. It corresponds to the Q method analysis.

Table 16.1: Detector-specific systematic uncertainties in E821 and proposed upgrade actions and projected future uncertainties for E989.

E821 Error	Size [ppb]	Plan for the New $g-2$ Experiment	Goal [ppb]
Gain changes	120	Better laser calibration; low energy threshold; temperature stability; segmentation to lower rates	20
Pileup	80	low energy samples recorded; calorimeter segmentation; Fast Cherenkov light; improved analysis techniques	40

16.3 Detector-related Systematic Uncertainties

In this section we discuss the primary systematic error issues related to the Detectors, Electronics, DAQ, and the Offline Analysis. Table 16.1 lists the Gain and Pileup uncertainties and projections for improvements in the new $g-2$ experiment. The traditional T method analysis is assumed because uncertainties can be reliably projected based on our considerable experience in these analysis efforts. Since the Q method is new, we have not included its positive and partially independent impact on the final statistical result, nor are we presently able to fully project associated systematics. This topic is an active study in the collaboration. One key attractive feature of the Q method is pileup immunity; there is no correction necessary, so that systematic uncertainty is absent. Our preliminary studies also indicate that the sensitivity to a systematic gain variation vs. time is also reduced. However, there are unexplored issues with acquiring a large body of fully integrated waveforms that remain to be studied.

16.3.1 Gain Changes and Energy-Scale Stability

The error budget for E989 assigns a 20 parts per billion limit to the gain systematic error. To connect this limit to hardware specifications, we conducted a simulation and analysis where gain perturbations could be applied. In general, the real data analysis plan is to correct any systematic hardware gain drifts over the short term time scale of a fill, event by event for each calorimeter station and crystal. The gain correction function must be prepared by detector response using both the known laser calibration pulses and evaluating the time stability of the pileup-corrected overall energy spectrum. Since the event rate changes by more than four orders of magnitude over a 700 μs fill, one might expect a hardware gain instability to vary with the rate. Accordingly, our simulation assumed a fill-scale gain perturbation of the form $G(t) = 1 + \epsilon \exp(-t/\tau_\mu\gamma)$, where $\tau_\mu\gamma$ is the 64.4 μs time-dilated muon lifetime, and ϵ is the magnitude of the unknown perturbation. In the above description, $G(t)$ represents the difference between the true gain vs. time behavior of the detector and electronics systems and the corrected one; that is, where $G(t)$ is the error in the gain correction, not its actual magnitude.

Monte Carlo techniques were used to build the standard $g-2$ decaying oscillation spectra with the expected E989 statistics, realistic detector effects, and built-in gain variations according to the above function. For each oscillation spectrum, a complementary laser-calibration data set was generated; the calibration system will provide an independent gain

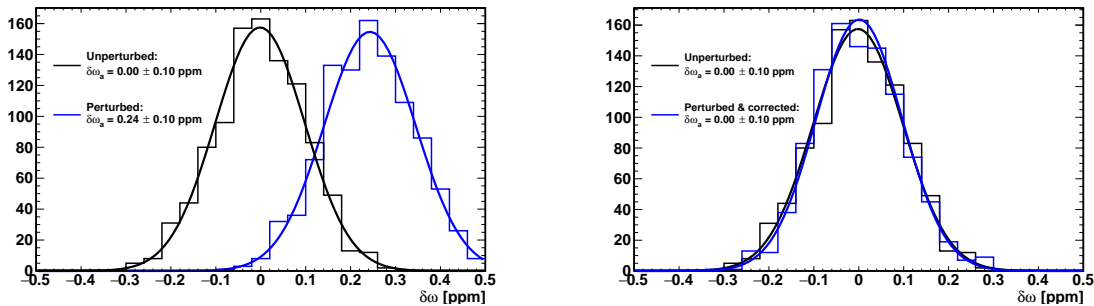


Figure 16.5: Histogram of extracted precession frequencies from 1000 simulated T -method histograms with and without an exponential gain perturbation of the form $5 \times 10^{-3} \exp(-t/\tau_\mu \gamma)$. The 100 ppb statistical uncertainty matches that expected in E989. The left plot does not include a correction from the simulated laser calibration data while the right plot does. After including the correction, the mean and width of the extracted ω_a distribution is restored to that of the unperturbed spectrum.

measurement every $5 \mu s$ with relative uncertainty no worse than 4×10^{-4} . For the purpose of this study, a statistical uncertainty of 4×10^{-4} was assumed. These laser data are fit and the resulting function used to correct the $g-2$ spectrum. The results for $\epsilon = 5 \times 10^{-3}$ are shown in Figure 16.5 for the T -based analysis method. The figure shows that correcting the gain with laser calibration data brings the ω_a perturbation well below the limit of 20 ppb. The target gain stability for E989 is $\delta G/G < 10^{-3}$ over a $700 \mu s$ fill, which this study shows is sufficient to achieve the gain systematic error goal while leaving breathing room for additional effects introduced by gain oscillations at ω_a . Such effects are the subject of ongoing studies.

The hardware gains of the E821 detectors [2] were determined to be stable to $\approx 0.15\%$ from early-to-late times within a storage ring fill. This limit was established by plotting the average energy for each $(g-2)$ period versus time after the PMTs were switched on. The gating circuitry in the tube base that allowed the PMTs to be turned off to avoid the initial burst of pions entering the ring also resulted in a small variation in the gain. For gain variations like this one, where the time constant is long compared to the $(g-2)$ oscillation period, the coupling to the ω_a frequency is small and, after correction, the residual systematic error was less than 20 ppb.

Several aspects of the current plan will be different. The first is that we will use silicon photo-multipliers (SiPMs), which can be saturated from a light burst and then recover with the same time constant as a low-light pulse. Each pixel recovers with a common time constant. Importantly, we do not intend to switch off these devices during injection because the anticipated hadronic-based flash will be (largely) absent. The initial pion flux at the target location will be reduced by the factor of 10^5 owing to the long beamline. Beam protons will be eliminated in the Delivery Ring using a kicker system timed to fire when the muon and proton bunches are well separated in space.

If the gain oscillates at a frequency ω_a , with an amplitude that varies in time, and with a phase that differs from that of the actual ω_a oscillation, then a direct error on the measured value of the anomalous precession frequency is produced. The average rate at

which energy is deposited into the calorimeters oscillates with frequency ω_a , and therefore any rate dependence in the gain of the detectors produces gain oscillations. In E821, we were able to demonstrate that the gain dependence on rate was small enough that its effect on ω_a was typically less than 20 ppb. In the new experiment, the increased beam rates will be partially offset by increased detector segmentation and our proposed monitoring system will be greatly improved compared to that employed in the past.

In E821, a UV-laser system was used to periodically pulse the scintillator in the detectors and thus monitor the complete gain and reconstruction chain during data collection against an out-of-beam reference counter. Unfortunately, the light distribution system included too many branches and not enough sub-branch reference detectors. Additionally, the laser intensity varied significantly on a shot to shot basis, making the per shot corrections large. Small fluctuations cascaded so that gain stability could be monitored to no better than a few tenths of a percent, which was not quite good enough to build a sensitive gain correction function. The system being designed for E989 will use cascaded distribution systems having multiple monitors at each stage. This is described in Chapter 17, Section 17.4.3.

The largest contribution to the gain systematic error in E821 came from analysis reconstruction induced gain oscillations at the ω_a frequency. The interpretation of the energy of a pulse from the fit to the waveform had a small bias. When a hardware signal rose above the waveform digitizer (WFD) trigger threshold, a pre-set minimum number of sequential samples was recorded. These data were fit offline to determine the peak height, time and the linear background under the pulse. However, if a trigger pulse was followed or preceded closely by another pulse, both pulses were fit together with a common background term, and the fitting region became longer compared to what is routinely used for a single pulse. In these pulses, the fitted energy was found to depend on the length of the fitting region, which was varying because of the hardware limitation. Because the data rate oscillates at frequency ω_a , and is higher at early than at late decay times, it follows that the fitting region length oscillated at frequency equal to ω_a and was, on average, longer at early times compared to late times. This produced a small, effective gain oscillation having an amplitude that decreases with time. A systematic error on ω_a results.

Given the current capabilities in data throughput, the new electronics will record all samples rather than triggered, fixed-length isolated islands. This avoids the intrinsic bias in the recorded data and allows reconstruction routines to compensate for the waveform islands that have more than one pulse. In addition, we will have one other new tool that will provide powerful information related to energy scale and gain. As discussed in Chapter 19, a large-acceptance tracker system will be built just upstream of at least two calorimeter stations. This system will reside inside a modified vacuum chamber. It will be capable of providing high-precision tracking with good momentum definition these calorimeter stations, which will provide an absolute energy scale. The position information obtained will also inform the calorimeter cluster algorithm development. The energy scale obtained from the directly calibrated stations can be bootstrapped to other calorimeters by comparing the average energy distributions from decay positrons, which are expected to be similar. In summary, we expect that the largest of the gain systematic errors from E821 will be eliminated by the design of the electronics and data acquisition systems, combined with the verification from the tracker. The smaller contribution will be reduced by a more precise hardware gain monitoring system.

16.3.2 Pileup

The term “Pileup” refers to the overlap of events in the calorimeter that originate from separate muon decays, too close to each other in time and space to be resolved into individual pulses. When two pulses overlap, the result is that the two individual events are lost, and one event with the sum of their energies is gained. In general, a finite time offset between the two pulses exists and the recorded pulse shape is widened so that the combined amplitude is somewhat less than the sum of the individual amplitudes. Because the fraction of pileup events increases with rate, a component is introduced into the time spectrum that decays with half the nominal lifetime, or $32 \mu\text{s}$, preventing a precise fit to the five-parameter function in Equation 16.3.

A more serious issue is that the muon spin precession phase varies with the energy of the pulse. A high energy positron has a larger radius of curvature and therefore a longer time-of-flight to the calorimeter, so it carries the phase of a muon that decayed earlier than one that produced a low energy positron. When two low energy pulses are lost and an apparent high energy pulse is gained, the high energy pulse still has the phase of the low energy pulses. The varying fraction of pileup over the fill causes an average early-to-late phase shift that directly distorts the fitted ω_a . Consequently, the (E, t) distribution of pileup pulses must be constructed and subtracted from the spectrum before it is fit.

The pileup distribution can be constructed based on the assumption that the probability of a pulse at time t is, to a good approximation, the average of the probabilities that it is found at times $t + \delta t$ and $t - \delta t$, provided that δt is small compared to the precession period. Consequently, secondary pulses from the “shadow” just before or after a primary pulse can be added to the primary pulse to form a constructed pileup event. To the extent that all fills have equal initial intensities of stored muons, the probability of a pulse at time t in one fill is also nearly the same as the probability of a pulse at the same time t in a different fill; this can provide another independent source of “shadow” events.

In E821, the construction of these distributions was complicated by the fact that only short islands around each pulse were stored, with a threshold of nearly 1 GeV required to store an island. The pileup distribution could only be fully reconstructed in a straightforward way at energies greater than twice this threshold. As described in the following chapters on electronics and data acquisition, in E989, the full waveform for the entire fill will be available. It will therefore be possible to construct pileup down to a very low energy threshold.

The unresolved pileup fraction scales quadratically with rate in each segment of the detectors. The effective size of the segment depends on the geometric extent of the shower. Our simulations demonstrate that an array of PbF_2 crystals, having 54 independent segments (see Chapter 17), and a smaller Molière radius compared to the Pb/SciFi used in E821, will provide an effective three-fold reduction in the intrinsic pileup based on the implementation of a very simple and robust shower separation routine and a 9-element cluster algorithm. The simulation includes a representative stored muon ensemble in the ring and correct spin physics in precession and decay.

Further improvements will also be associated with the correction for unresolved pileup. After following the pileup subtraction procedure described above, we will be able to check the result using an applied artificial deadtime (ADT). The ADT is the time established in the analysis software below which two pulses are not resolved (even if they can be). The

analysis proceeds by sorting data using a series of ADT values beginning with the intrinsic, device-specific constraints, and artificially extending to much larger values that exaggerate the pileup. The extraction of ω_a is then done for each data set, and ω_a will be plotted as a function of ADT. In principle, ω_a should not depend on ADT, but a small correction could be included by taking the deadtime-free value that occurs at the zero-ADT extrapolation point. We have spent considerable laboratory bench time and offline pulse-reconstruction efforts to determine and optimize the minimum hardware ADT that our detectors will permit. Our laboratory tests demonstrate that pulses separated by 5 ns or more can be resolved easily for most pulse-amplitude ratios expected.

In addition to the work that was done for E821, we have also carried out a precision muon lifetime analysis with a pileup correction algorithm based on this pileup construction and ADT extrapolation concept. The work is well documented [4, 3]. The algorithms will be tested using the one station that has a high-resolution tracker (see Chapter 19) that can resolve pileup events at the few mm level and provide the corresponding momentum of each. Comparing identified pileup events from the tracker to the interpretation of the same events by the calorimeter will give a great degree of confidence in the methods.

For reference we comment on what had been achieved in the past. The pileup systematic error of 80 ppb in the E821 experiment was obtained from three components listed below. The first two were correlated and add linearly. The third is not correlated so it was added in quadrature to the other two.

1. Pileup efficiency, 36 ppb. This is due to an estimated 8% uncertainty in the amplitude of the constructed pileup spectrum.
2. Pileup phase, 38 ppb. This is the error due to the uncertainty in the phase of the constructed pileup spectrum.
3. Unseen pileup, 26 ppb. This is the error due to pulses so small that they cannot be reconstructed and therefore they are not included in the pileup construction. In general, the energy from these pulses cancels out, because they occur as often in the pedestal, where they lower the fitted pulse energy, as under the pulse, where they raise it. This error accounts for the potentially incomplete cancellation.

We expect that the segmented detectors, better laser calibration, more complete waveform record storage, verification of methods by using the tracker, and the use of our more modern extrapolation algorithms will lead to a comprehensive pileup correction with minimal uncertainty. We assign up to 40 ppb here to account for any difficulties in the anticipated analysis. As mentioned earlier, the Q method is complementary to the traditional T method and has different sources of systematic errors. The most significant difference is the effect of pileup—it will be greatly reduced for the Q method.

16.4 Detector System Overview

Figure 16.6 shows the locations of the calorimeters and trackers with respect to one of the 12 vacuum chamber segments. Two decay positron trajectories are indicated in the figure




Figure 16.6: Scalloped vacuum chamber with positions of calorimeters indicated. A high- (low-) energy decay electron (or positron for the preferred positive muon storage) trajectory is shown by the thick (thin) red line, which impinges on the front face of the calorimeter array.

corresponding to high- and low energy events. The decay positrons have momenta below the muon storage momentum and therefore curl to the inside of the ring through the opening in the C -shaped magnet. Electromagnetic calorimeters are used to intercept the positrons and provide a measurement of energy and time of detection.

To operate the calorimeters and record the data several subsystems are required and shown schematically in Figure 16.7. First a segmented array of PbF_2 crystals absorbs the energy of the positrons and converts it to Cherenkov light. There are 24 of these arrays spaced evenly around the ring. Each station contains 54 crystals for a total of 1,296 individual PbF_2 crystals. Then the light is detected by SiPMs. These devices produce a pulse that is digitized separately for each block by custom digitizers running at 800 MSPS with 12 bit depth. The digitized signals are passed to a farm of graphics processing units (GPUs) where they are reduced to a form suitable for storage. In the case of the event reconstruction style analysis (T Method) regions of interest or “islands” are selected for recording. In the case of the current readout analysis (Q Method) the waveforms are summed across several digitizer samples as well as within individual stations to reduce the size of the recorded data.

The laser calibration subsystem provides the means to monitor the detector gain stability. The proposed SiPM readout devices are particularly sensitive to bias and temperature stability. We have dedicated subsystems to provide a stable bias supply to each device and to monitor temperature inside the calorimeter enclosures. Two tracker stations will gather data that contains a large number of resolved two-track events that might appear as unresolved pileup in the calorimeters. These data will be crucial to develop our calorimeter cluster and pileup-subtraction routines. Single track events will be used to determine the absolute energy calibration. The fiber harp system and the trackers are both needed to determine the stored muon beam distribution, which must be known to make the electric field and pitch corrections. The slow control data from the entire experiment will be gathered by a dedicated subsystem, which will monitor the performance and health of all the subsystems described in this TDR. The chapters which follow separately describe the requirements and design of these subsystems.

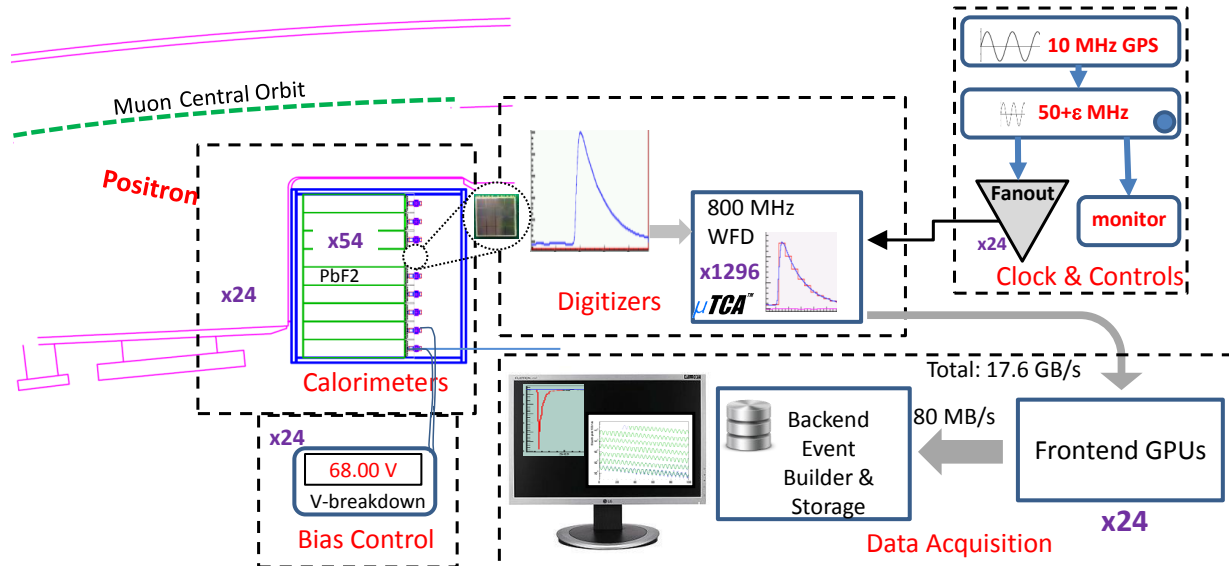


Figure 16.7: Schematic of the relationship of a number of the ω_a subsystems. The dashed boxes represent distinct responsibilities of different groups within the collaboration.

16.4.1 Calorimeter Subsystem Considerations

The calorimeters will be placed adjacent to the storage ring vacuum chambers, and located at 15 degree intervals around the ring. The 24 stations and their locations are constrained by the reuse of the E821 vacuum chambers, see Figure 16.6. These parameters were optimized in a study preceding E821 construction and the conclusions remain valid for E989. The number of emitted decay positrons vs. decay energy is shown in Figure 16.8. The geometry is designed to favor the high energy positrons because they are most correlated with the muon spin. low energy positrons have a higher chance to curl in between calorimeter stations and be lost.

The design of the new calorimeters is constrained by the unusual experimental demands. It is important to emphasize that the relevant time scale for most systematic uncertainties is one 700 μs long measuring period. The initial instantaneous event rate of several MHz drops by almost five orders of magnitude during the 700 μs measuring period; thus, any rate-dependent detector or readout response changes must be accurately known. The overall measurement system must be extraordinarily stable for each short-term storage ring fill; however, long time scale drifts can generally be tolerated. As discussed in Section 16.3.1, if the gain vs. time-in-fill, $G(t)$, is not constant, then ω_a might be incorrectly determined. Similarly, a time shift Δt owing to the clocking system or other influence can also change the fitted frequency. From our simulations the stability conditions that ensure less than a 50 ppb shift to ω_a require that $\Delta G < 0.1\%$ and $\Delta t < 10$ ps over a 200 μs interval.

As described in Section 16.3.2, incorrectly treated pileup can lead to a large systematic uncertainty. As illustrated in Figure 16.6 low energy positrons have a shorter flight path to the detector compared to higher-energy positrons; thus they correspond to muons having a different average spin at the time of the decay. The rate of fake high energy positrons




Figure 16.8: Number of decay events vs. decay energy (black), number of events that deposit detectable energy over 100 MeV (blue), and number of events that deposit detectable energy over 1860 MeV (red). This plot is generated from a simulation using full geometry, including pre-showering effects, and tracking the particles from muon decay through showering and energy deposition in the calorimeter.

coming from coincident low energy positrons has a $\sim e^{-2t/\gamma\tau}$ time dependence. This means the pileup rate falls twice as fast as the muon population decays. To minimize pileup, the calorimeter response must be fast (few ns) and the readout system must record information to enable the distinction between closely occurring pulse pairs, which strike the same detector elements. This information should also provide a mechanism to correct the data, on average, by removing the pileup events. Furthermore, if the detector segmentation is optimized, many simultaneous lower-energy positrons will be recorded in independent area of a calorimeter station and thus will not be interpreted as a pileup event. The goal in the detector design is to reliably resolve same-element pulses separated by 5 ns or more, to segment the detector to minimize pileup, and to accurately subtract unresolved pileup.

The calorimeter resolution must be moderately good near 1.8 GeV (better than 10%) to provide adequate energy discrimination. However, it also must be compact to avoid a preponderance of positrons that strike the inside face of the detector. Higher density materials enable a more compact detector but yield inferior energy resolution. For E989, we are aiming for a factor of 2 improvement in resolution compared to E821. The material PbF_2 meets both criteria of compactness and resolution. The energy resolution of the PbF_2 calorimeter system was tested at SLAC and shown to be 2.8% at 3.5 GeV.

16.4.2 Electronics and Data Acquisition Subsystems Considerations

To guarantee deadtime-free calorimetry readout, the signal from each of the 1,296 active calorimeter channels are continuously digitized for every 700 μs muon fill. Those waveforms are then transferred to the DAQ system for data reduction – isolation of the time windows containing electromagnetic showers – and storage. The DAQ must apply an energy threshold to identify showers within a station, so the 54 waveforms from that station must be summed to keep the threshold independent of the incident positron position in a crystal. All 54 WFD waveforms must therefore be transferred to the same frontend DAQ system, which will use the waveform sum to perform data reduction on the digitization stream: identification of time islands with activity (T method) and time rebinning of the waveform (Q method).

The energy range of interest for an individual calorimeter element is 25 to 3100 MeV for single events. Allowing pileup, suggests pushing the upper limit close to 5000 MeV. A digitizer with 12-bit depth (4096 channels) is ideal. It will allow good pulse definition, important for the energy resolution requirements, and it will not saturate for the highest energy events. As discussed in Chapter 17, the pulse shape rise time is approximately 2-4 ns. The laboratory tests have led to selection of a sampling rate of 800 MSPS. Therefore, each digitized waveform corresponds to 560K 12-bit words for each muon fill.

A precision oscillator (“clock”) will provide the time base from which the ω_a frequency is measured. It must be controlled to provide negligible error compared to the anticipated 100 ppb uncertainty on ω_a . In order to achieve this, the clock must have jitter that is significantly less than the 1.25 ns sampling period of the waveform digitizers. It must also have very low (< 10 ps) systematic shift across the time of a single fill. This latter requirement is important because of the large variation in event rate within a fill. A systematic time-slew that is correlated with muon or positron intensity would bias the result. The clock system

must also enable a convenient blinding scheme such that the actual precise clock frequency cannot be known to the data analyzers.

The data acquisition system must provide a deadtime-free readout of calorimeter segments using the waveform digitizers. Onboard memories in the digitizers will buffer the raw data and allow its asynchronous readout, thus decoupling the data acquisition cycles from storage ring fills. A frontend layer of multicore CPUs/GPUs will process the digitized records of each fill from every calorimeter segment into T -method, Q -method and other derived datasets. A backend layer of multicore CPUs/GPUs will handle the assembly of event fragments from the frontend layer and transfer of assembled events to the mass storage. Each stored event will represent a complete deadtime-free history of the entire activity in the detector system for every fill cycle.

References

- [1] G. W. Bennett *et al.* [Muon G-2 Collaboration], Nucl. Instrum. Meth. A **579**, 1096 (2007).
- [2] Sedykh S., et al. *Nucl. Instrum. Meth. A*455:346 (2000).
- [3] V. Tishchenko *et al.* [MuLan Collaboration], Phys. Rev. D **87**, 052003 (2013), arXiv:1211.0960 [hep-ex].
- [4] D. Webber, *A part-per-million measurement of the positive muon lifetime and a determination of the Fermi constant*, doctoral thesis, University of Illinois at Urbana-Champaign (2010).