

The development, testing and characterisation of a straw tracking detector and readout system for the Fermilab muon g-2 experiment

Thomas Simon Stuttard
University College London

Submitted to University College London in fulfilment
of the requirements for the award of the
degree of **Doctor of Philosophy**

December 30, 2017

Declaration

I, Thomas Simon Stuttard confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

A handwritten signature in black ink, appearing to read 'Tom Stuttard', with a stylized flourish at the end.

Thomas Simon Stuttard

Abstract

The anomalous magnetic moments of leptons can be both measured and theoretically predicted with high precision, and as such provide stringent tests of the Standard Model (SM) of particle physics. The anomalous magnetic moment of the muon, a_μ , in particular offers sensitivity to new physics within the reach of current experiments. The current world's best measurement of a_μ made at the Brookhaven National Laboratory (BNL) deviates from the SM prediction by over 3σ , providing a tantalising but inconclusive hint that contributions from new physics may be present [1].

A new a_μ experiment is currently under construction at the Fermi National Accelerator Laboratory (Fermilab), known as the Fermilab muon g-2 experiment. Due to commence in 2017, this new experiment is expected to improve the experimental measurement precision of a_μ by a factor of 4, and is seeking to confirm or reject the current discrepancy [2]. The experiment, like its BNL predecessor, will measure the precession of the μ^+ spin in a magnetic storage-ring.

This new experiment has three straw tracking detectors used to measure the beam profile of the stored μ^+ , as well as other dynamic properties of the beam. This thesis describes the design of these detectors, and the data acquisition system and subsequent data processing systems that have been developed to read them out and make their data available for physics analysis. The performance of the straw trackers and the readout systems at a number of beam tests and test stands is presented, in addition to Monte-Carlo simulations of the trackers in the final experiment.

Acknowledgements

I have received immense support during my time working on the Fermilab muon g-2 experiment from people too numerous to name. In particular, I would like to thank my supervisor Mark Lancaster, and the University College London muon g-2 team Rebecca Chislett, Erdem Motuk, Matt Warren, Samer Al-Kilani and Gleb Lukicov. Thanks is also due to other teams involved in the tracker development at the University of Liverpool, Boston University, Fermilab, Oxford University and Northern Illinois University with whom I have worked closely and travelled extensively. These include but are not limited to Joe Price, Saskia Charity, Will Turner, Themis Bowcock, Steve Maxwell, Barry King, James Mott, Nick Kinnaird, Rob Carey, Lee Roberts, Brendan Casey, Tammy Walton and Farrukh Azfar. Thanks also to those across the g-2 experiment software, DAQ and DQM teams with whom I have collaborated and interfaced with, including Wes Gohn, Tim Gorringer, Sudeshna Ganguly, Aaron Fieberg, Kim Siang Khaw, Adam Lyon and David Sweigart. I can only apologise to all these people for the sheer number of TODO statements left strewn throughout the g-2 software.

I am particularly grateful to the support of Mandy Rominsky and the rest of the Fermilab Test Beam Facility team for all their efforts and support during the tracker test beams.

Thanks must also go to Mario at Best Western hotel in Naperville, Illinois for his constant beyond the call of duty hospitality, and to World of Beer for providing a light at the end of the test beam tunnel.

Finally, I would like to thank my fiancé Signe for all her support during my PhD and putting up with my long absences away at Fermilab, and my baby daughter and boss Ivy for granting me just enough sleep to write a thesis.

Contents

List of figures	8
List of tables	23
1 Introduction	24
2 The theory of lepton anomalous magnetic moments	28
2.1 Lepton magnetic moments	28
2.2 The anomalous magnetic moment of the muon	31
2.3 Electric dipole moments	33
2.4 Standard model value of a_μ	34
2.4.1 QED contributions to a_μ	35
2.4.2 Electroweak contributions to a_μ	36
2.4.3 Hadronic contributions to a_μ	36
2.4.4 Value and uncertainty of a_μ^{SM}	38
2.5 Possible new physics contributions to a_μ	39
3 The Fermilab muon g-2 experiment	41
3.1 Experimental goals	41
3.2 Measuring g-2	42
3.3 The Fermilab muon g-2 experiment	48
3.3.1 Muon production and delivery to the storage-ring	48
3.3.2 Muon storage in the ring	52
3.3.3 Full treatment of precession in the storage ring	54
3.3.4 μ^+ decay in the storage-ring	55
3.3.5 Detector systems	56
3.3.6 Storage ring magnetic field	60
3.3.7 Combining ω_a and ω_p	65
3.3.8 Measuring an EDM	66

4	The straw trackers	68
4.1	Overview	68
4.2	Tracker goals	69
4.2.1	Observation of stored beam	69
4.2.2	Supporting the calorimeter	69
4.2.3	Observation of vertical angle oscillation	70
4.3	Tracker high-level design	71
4.4	Detection principle	73
4.5	Tracker detailed design	76
4.5.1	Straw design	76
4.5.2	Choice of gas	77
4.5.3	Tracker module design	79
4.6	Track formation	82
4.7	Ageing	84
5	The tracker data path	85
5.1	Readout electronics	85
5.1.1	Frontend electronics	85
5.1.2	Backend electronics	89
5.2	Data acquisition software	91
5.2.1	DAQ software description	92
5.2.2	Emulators	96
5.2.3	MIDAS	96
5.2.4	Data rates	98
5.3	Offline processing	102
5.3.1	MIDAS-to-art	102
5.3.2	Tracker offline data processing	104
5.4	Data monitoring systems	105
6	Testing the tracker	108
6.1	Test beams	108
6.1.1	The MTest beam	111
6.1.2	Event reconstruction	113
6.2	Test stands	114
6.3	GEANT4 simulation	114
6.4	GARFIELD simulation	116

6.5	Detector testing results	118
6.5.1	Straw drift times	118
6.5.2	Straw drift velocity	121
6.5.3	Single module hit reconstruction	123
6.5.4	Straw resolution	128
6.5.5	Reconstruction resolution	132
6.5.6	Straw efficiency	133
6.5.7	Optimising wire voltage	137
6.5.8	Optimising ASDQ discriminator threshold	142
6.6	Conclusions	145
7	Simulations of tracker performance	146
7.1	GEANT4 simulation of the g-2 experiment	146
7.2	Straw and tracker occupancy	147
7.3	Tracker acceptance	152
7.4	Tracker and calorimeter e^+ matching	157
8	Outlook	159
	Bibliography	161

List of figures

2.1	Feynman diagrams at (a) tree-level (Dirac) and (b) one-loop level (Schwinger) of the electromagnetic interactions determining the value of g_e . In both cases the electron couples to an external magnetic field.	30
2.2	(a) LO and (b) NLO QED Feynman diagrams contributing to a_μ [2].	35
2.3	LO Feynman diagrams of the EW processes contributing to a_μ [3]. There are contributions from the W (a) and Z (b) massive vector bosons and (c) the Higgs boson.	36
2.4	LO Feynman diagram of the hadronic contribution to a_μ , where H represents any hadron [2].	37
2.5	Example higher (beyond leading) order Feynman diagrams for hadronic contributions to a_μ , where H represents any hadron [2].	37
2.6	Feynman diagram of an example light-by-light hadronic contribution to a_μ , where H represents any hadron [2].	38
2.7	SM contributions to the (a) value and (b) uncertainty of a_μ	39
	(a)	39
	(b)	39

- 2.8 Maximum possible value of a_μ in the flavour-aligned H2DM, shown as a function of the mass M_A of one of the additional Higgs bosons in this model, A , and for three values of the lepton Yukawa coupling, ζ_l , where these parameters are found to be those that most strongly drive large predictions for a_μ [39]. The masses of the remaining three Higgs boson states in this model (excluding A and the SM model Higgs) are set to 250 GeV in this plot. The yellow band indicates the value of a_μ measured at BNL and its uncertainty. The H2DM contribution to a_μ is computed using a two-loop calculation. 40
- 3.1 μ^+ decay kinematics, with the post-decay state shown being the preferred state due to parity violation in the weak decay and spin and momentum conservation. LH and RH refer to left- and right-handed particles respectively. Red arrows represent the momentum vectors and blue arrows represent the spin vectors. 44
- 3.2 The e^+ emission angle in μ^+ decays, expressed relative to the μ^+ spin direction in the μ^+ rest-frame. 45
- 3.3 The orientation of the μ^+ momentum, p , and spin, s , vectors as they orbit in the storage-ring, as seen from above the ring. The magnetic field, B , vector is also shown. The angle between the spin and momentum vectors, θ_a , is known as the anomaly angle. The rate of change of θ_a is known as the anomaly frequency, ω_a , and is one of the fundamental quantities measured in the Fermilab muon g-2 experiment. The circular motion of the spin vector is shown by the dashed blue line. 46
- 3.4 The magnetic storage-ring during installation at Fermilab. The ring diameter is 14 m. 49

-
- 3.5 Histograms of the times that the decay e^+ hit the calorimeters in a Monte-Carlo simulation. Hits from all e^+ are shown in red, whilst hits from e^+ with energies greater than 1.8 GeV are shown in blue. Both distributions have a falling exponential component resulting from the μ^+ decays, where the time dilated lifetime of the μ^+ is 64 μs . The energy cut introduces a sinusoidal component which varies with the anomaly frequency, 4.4 μs . The simulation includes effects such as momentum spread and time width of the injected beam. 50
- 3.6 Diagram of the locations of beam-control instrumentation and detectors around the storage-ring. The ideal orbit path and direction is shown in green. 53
- 3.7 Example trajectory of a decay e^+ following the decay of an orbiting μ^+ . The e^+ is stopped and detected in a calorimeter, three of which are pictured. Image courtesy of J. Kaspar. 56
- 3.8 Photograph of a calorimeter during construction. The perspective is from the top and back of the detector. Individual PbF_2 crystals can be seen connected to readout SiPMs. Image courtesy of J. Kaspar. . . 57
- 3.9 Photograph of a fiber harp plane. The harp is in the deployed position, with vertical scintillating fibers used to measured the horizontal beam profile. 59
- 3.10 The cross-section of the g-2 storage ring. Indicated are the coils, pole pieces and yoke of the superconducting magnet, the μ^+ storage region, the location of fixed NMR probes, and various shimming devices: top-hats, edge-shims, surface-coils, and wedges. 61
- 3.11 The storage-ring magnetic field relative to the overall average field strength, plotted versus azimuth, at BNL (blue) and Fermilab (red). The Fermilab measurement was taken following the first shimming phase, and will be improved further following the installation of ring instrumentation. 62

3.12	The azimuthally-averaged storage-ring magnetic field relative to the overall average field strength at Fermilab following the first shimming phase, plotted versus vertical and radial position within the μ^+ storage region. The field uniformity will be further improved during future shimming before physics runs begin.	62
3.13	Angular momentum vectors of the μ^+ spin precession resulting from magnetic and electric dipole moments, and the resulting tilt of the overall spin precession plane [2].	67
4.1	A view from above of the trajectory of a decay e^+ and the region available for the trackers between two calorimeters.	72
4.2	Top-down view of a charged particle traversing a Fermilab g-2 straw in a GARFIELD simulation. The charged particle track is shown in red, with black circles representing the primary ionisations. Blue lines represent the drift trajectories of the primary electrons until they terminate on the central sense-wire. Green lines represent the drift of the avalanche ions. The electron paths are visibly curved due to the external magnetic field.	75
4.3	Diagram of the relationship between the drifting of primary ionisation electrons to the straw wire and the DCA reconstructed from the time taken by this drift. The drift cylinder specified by the DCA is also shown.	76
4.4	One straw tracker module. The perspective is downstream, i.e. looking in the direction of e^+ flight.	79
4.5	The straw layers in a straw tracker module. The straws are arranged in two pairs of two layers. In each pair of layers, the straws have the same stereo angle. The two pairs of layers have equal and opposite stereo angles.	80

-
- 4.6 A view into the manifold of a straw tracker module. The ASDQ boards (see section 5.1.1) are mounted directly on the ends of the wires, and are connected to the off-detector readout electronics via the flexibles. Copper cooling bars make a thermal connection between the electronics boards and a water-cooling channel drilled into the manifold. The white end caps of the straws can be seen in the lower left, which hold the wires under tension and have openings to allow gas to flow through. 81
- 4.7 Diagram showing how 8 straw modules are arranged to form a single straw tracker, and how this tracker is aligned relative to the storage region and the downstream calorimeter. The direction of the orbiting μ^+ is shown by the blue line. 82
- 4.8 Staircase design of the vacuum chamber to which tracker modules are mounted. The left-hand side will be mounted closest to the downstream calorimeter. 83
- 5.1 The path of clock, control signals and straw-hit data through the various layers of frontend and backend boards in the straw tracker readout system. The physical location of the boards is also indicated. 86
- 5.2 The hierarchy of frontend and backend boards and the numbers of each used in the straw tracker readout system. 86
- 5.3 Simplified schematic of the signal path in the ASDQ. (a) The raw signal includes a number of fast spikes corresponding to the avalanches from individual primary ionisations. (b) The shaped signal integrates the raw signal into a single pulse. (c) The discriminator records where the shaped pulse crosses a threshold (the red line represents the threshold, the blue lines represent the two crossing points). The time of the first crossing is considered to be the hit time. (d) The output ASDQ signal records only the leading and trailing edges of the shaped pulse following discrimination. Note that the ion tail is not included in this diagram. 87

-
- 5.4 Four FLOBBERS mounted on the outside of the vacuum chamber, along the inside radius of the storage-ring. Each is attached via a snout through the flange to a tracker module within the vacuum chamber. 88
- 5.5 The tracker DAQ software FSM. Red boxes represent states. Blue ellipses represent commands, with the arrows indicating the states a command may be executed on and what the resulting state of the command is. 94
- 5.6 GUI for board register read and write operations using IPbus. The GUI periodically reads and displays all IPbus registers for a given board. Clicking on a given register allows a new value to be written to the register. 95
- 5.7 The top-level status page of the MIDAS web interface. 97
- 5.8 The MIDAS ODB, as viewed using the MIDAS web interface. The page shown contains the run-time settings for the tracker AMC13. 98
- 5.9 Proposed fill structure for the Fermilab muon g-2 experiment [2]. . . 100
- 5.10 Maximum event rate achievable by the DAQ software versus event data size (blue line), measured at a test stand. Also indicated are the Fermilab muon g-2 experiment fill rate (red), expected event size during physics runs (green) and maximum possible event size (magenta). The time taken to readout one event scales linearly with the event size, and as such the number of events that can be readout in 1 s scales as $1/(\text{event size})$ 101
- 5.11 Communication between MIDAS and MIDAS-to-art when running in online mode. 103
- 5.12 Offline data processing stages for the straw tracker data. The green box represents MIDAS data, the red boxes are art records, and the blue ellipses represent art modules performing the data processing. . 104
- 5.13 Online tracker data monitoring web display. Readout hardware monitoring is shown. 106

5.14	Online tracker data monitoring web display. Counts of hits in ASDQ-TDC channels (each corresponding to a single straw) for single tracker module are shown.	107
5.15	Tracker slow-control data monitoring web display. This example shows the temperatures of the boards in the manifold and FLOBBER for a single tracker module.	107
6.1	Detector arrangement at the 2015 test beam. The beam path is indicated by a green line. The location of the tracker module within the vacuum chamber is indicated by a red box.	109
6.2	Prototype tracker module used for the 2015 test beam.	110
6.3	Auxiliary detectors used in the 2015 test beam	111
	(a) MWPC	111
	(b) Silicon telescope	111
6.4	Beam profile as measured by the MWPC immediately upstream of the tracker module. The top left and right plots show the fitted vertical (y) and horizontal (x) beam profiles respectively. The bottom left plot shows the spread of the measured hits. The bottom right plot shows a close-up of the central region of the bottom left plot, with a red box indicating the size of the silicon telescope's active area.	112
6.5	A single tracker module mounted in a cosmic-ray test stand. Scintillator paddles clad in black tape can be seen above and below.	115
6.6	116
	(a) GEANT4 simulation of the test beam. The silicon planes are coloured orange, whilst the straws are coloured yellow and red. A particle trajectory is shown in white. The vacuum chamber is not shown.	116
	(b) GEANT4 simulation of the cosmic ray test stand. The red lines are particle tracks, with the straw modules shown in grey.	116

6.7	Event display showing a 120 GeV proton traversing a straw as simulated using GARFIELD. The perspective is top-down into the straw. The proton track is shown in red, the locations of primary ionisation events with small black circles, and the path of the drifting primary ionisation electrons is shown in blue. The large black circle represents the straw wall.	117
6.8	118
(a)	Raw signal on the wire from GARFIELD simulations	118
(b)	Signal following the readout electronics model from GARFIELD simulations. The red line represents the threshold to trigger a hit, and indicates the leading and trailing edges of the hit recorded. The leading edge threshold crossing is taken as the hit time.	118
6.9	Drift times in the straws for the 2015 test beam	119
(a)	Test beam data with 1800 V wire voltage	119
(b)	GARFIELD simulations with 1500 V (blue) and 1800 V (red) wire voltage	119
6.10	Drift times in the straws for the cosmic-ray test stand	120
(a)	Test stand data with 1650 V wire voltage	120
(b)	GARFIELD simulation with 1650 V wire voltage	120
6.11	Scatter plot of the drift times measured in the straws (y-axis) versus the track-to-wire DCA reconstructed from the silicon track (x-axis). The plot shows test beam data taken with argon-ethane 50:50, a wire voltage of 1800 V, and ASDQ discriminator threshold voltage of 300 mV. 121	
6.12	Scatter plot of drift times in the straws versus the track-to-wire DCA from GARFIELD simulations of the straw response to the MTest proton beam.	122
(a)	1500 V wire voltage	122
(b)	1800 V wire voltage	122

6.13 Top-down view of a straw cross-section, showing the DCA to the wire for a primary ionisation cluster produced at the track-to-wire DCA, and of a cluster offset from this position. Two cases are shown: in (a) the track is far from the wire, and the difference between the DCA of the two clusters to the wire is small, whilst in (b) the track is close to the wire, and the difference between the DCA of the two clusters to the wire is large. In both cases the separation of the two clusters along the track is the same. 123

6.14 The mean DCA of primary ionisation clusters to the wire in the straws versus the track DCA from GARFIELD simulations of the response on the straws to the MTest proton beam. The simulated wire voltage is 1800 V. The coloured lines represent the closest (black), second closest (red) and third closest (green) cluster to the straw wire liberated by a single traversing particle. 124

6.15 125

(a) Diagram of a charged particle traversing the four straw layers in the tracker module at the test beam. The perspective is top-down. The blue line represents the charged particle, and the red lines represent the DCA of the particle to the straw wires. 125

(b) Diagram showing the anti-correlation between the DCA to the two wires in a straw doublet for particles normally incident on the straw layers. The particle trajectory is shown in blue. The DCA in each straw is shown by solid red lines, whilst the corresponding drift cylinder (seen as a circle in this top-down perspective) is shown in the red dashed line. 125

6.16 The 2D hit reconstruction process in a single tracker module. 127

6.17 The event selection process for hit reconstruction in a single tracker module. 127

6.18 The reconstructed 2D hit positions in the straw module. 128

6.19 130

(a)	Scatter plot of the track-to-wire DCA reconstructed in the straws (y-axis) versus that reconstructed by the silicon (x-axis). This data was taken with argon-ethane 50:50 with a wire voltage of 1800 V. This plot consists of approximately 4000 proton tracks.	130
(b)	Residuals of the track-to-wire DCA measurement made with the straws with respect to the silicon track-to-wire DCA measurement (blue). A Gaussian function (red) is fitted to the data, where the fit parameter $\sigma = 200 \pm 6 \mu\text{m}$ is defined as the straw resolution. This plot consists of approximately 4000 proton tracks.	130
6.20	Residuals of the reconstructed versus true values of the track-to-wire DCA from GARFIELD simulations of the response of the straws to the MTest proton beam. Wire voltages of 1500 V (blue) and 1800 V (red) are shown.	131
6.21		133
(a)	Histogram of the radial component of the residuals from the silicon tracks to the reconstructed straw hit positions. A Gaussian fit to the data is shown in red, and a Gaussian curve with $\sigma_x = 100 \mu\text{m}$ corresponding to the expectation from equation 6.3 is shown in blue.	133
(b)	Histogram of the vertical component of the residuals from the silicon tracks to the reconstructed straw hit positions. A Gaussian fit to the data is shown in red, and a Gaussian curve with $\sigma_x = 750 \mu\text{m}$ corresponding to the expectation from equation 6.3 is shown in blue.	133
6.22		134
(a)	Straw efficiency versus the track-to-wire DCA measured by the silicon at the test beam. The discriminator threshold voltage used is 300 mV. Each coloured curve represents a single straw plane. The straw planes are identified by their “view” number, referring to a pair of adjacent straw planes with the same stereo angle in a straw module, and “layer” number, which discriminates between the two straw planes within a given view.	134

(b)	GARFIELD simulation of the straw efficiency versus track-to-wire DCA at the test beam, shown for the wire voltages 1500 V (blue) and 1800 V (red)	134
6.23	Number of straw layers hit for a particle traversing a single straw module in Monte-Carlo simulation, where 100% straw efficiency is assumed.	135
6.24	Number of straw layers hit for a particle traversing a single straw module	136
(a)	Test beam	136
(b)	Test stand (one module)	136
6.25	The ratio of the number of tracks leaving hits in 3 vs 4 layers of straws from a Monte-Carlo simulation, plotted as a function of the simulated straw efficiency. The value of this ratio is also indicated for the test beam and test stand data.	136
6.26	The straw hit rate when exposed to an ^{55}Fe source, measured across a range of discriminator threshold voltages. The charge corresponding to these discriminator threshold voltages, as calibrated by lab measurements, is shown on the x-axis as the top of the plot in red.	138
6.27	Histogram of the number of primary electrons liberated in a straw by a 5.9 keV photon (characteristic of ^{55}Fe radioactive decay) from a GARFIELD simulation. The larger peak is due to the photon depositing all energy in the gas. The smaller peak results from the case where a secondary photon is emitted by an argon ion and escapes the straw, reducing the total amount of energy available to liberate electrons. This is known as an argon escape peak.	139
6.28	Gain in the g-2 straws versus wire voltage, determined by extrapolation from gain measurements using an ^{55}Fe source at wire voltages of 1600 V, 1650 V and 1700 V.	140

- 6.29 The hit rate from a ^{90}Sr source in a single straw versus wire voltage, as measured at a test stand using an ASDQ discriminator threshold of 300 mV (black). Also shown are the expected straw hit rates from GARFIELD simulations using ASDQ discriminator thresholds of 200 mV (red), 300 mV (green) and 400 mV (blue). 142
- 6.30 The straw hit rate versus the wire voltage from GARFIELD simulations. The results for simulations of cosmic-ray muons (black) and a ^{90}Sr source (green) at the test-stand, 120 GeV protons at the test beam (red), and 1.5 GeV e^+ at the Fermilab muon g-2 experiment (blue) are shown. Note that the simulated results for the Fermilab muon g-2 experiment use the lower discriminator threshold of 200 mV, compared to 300 mV for the other cases. This is because the low-voltage supplies to be used in the final experiment have improved noise characteristics than those used during testing, meaning a lower discriminator threshold can be used. 143
- 6.31 The number of background hits in each ASDQ channel in the absence of a particle source, measured as a function of the ASDQ discriminator threshold voltage. The curves for each channel are offset vertically for clarity. 144
- 7.1 Top-down view of the storage-ring in the g-2 GEANT4 simulation. The design μ^+ orbit is clockwise and shown in green, with e^+ trajectories shown in red. White lines indicate the outline of geometric elements. 24 calorimeters can be seen equally spaced around the ring, as well as three tracker stations. The vacuum chambers are not shown for clarity except at the three tracker stations. All other geometry has been hidden for clarity. The diameter of the ring is 14 m 148
- 7.2 Top-down view of a detector station in the g-2 GEANT4 simulation, using the same colouring scheme described for figure 7.1. Calorimeters can be seen in the far left and right, with eight tracker modules immediately upstream of the right-hand calorimeter. To the left of the tracker modules is an inlet to the vacuum chamber. 149

7.3	Straw occupancy heat map showing the mean number of hits per fill in each straw in a single tracker station, where the colour scale is the number of hits. The storage region is just beyond straws at the top of this plot, with μ^+ orbiting from left to right. The downstream calorimeter is just beyond the straws in the right-most layer on the plot. Each of the 8 tracker modules has 4 layers, and the last module before the downstream calorimeter is layers 28-31 in this figure. Each tracker layer has 32 straws.	150
7.4	150
	(a) Mean straw occupancy per fill from decay e^+ hits, where the colour scale is the number of hits.	150
	(b) Mean straw occupancy per fill from secondary particle hits, where the colour scale is the number of hits.	150
7.5	151
	(a) Histogram of the number of hits in an individual straw during a single μ^+ fill of the storage-ring. The plot includes data from all straws in all three tracker stations, and so averages across the exposure experienced by different straws at different locations.	151
	(b) Histogram of the number of hits in an individual TDC during a single μ^+ fill of the storage-ring. The plot includes data from all TDCs in all three tracker stations, and so averages across the exposure experienced by different TDCs at different locations. .	151
7.6	Straw heat map showing the mean number of pile-up hits per fill in each straw in a single tracker station, where the colour scale is the number of hits.	151
7.7	Detector azimuthal acceptance. In the top panel of each plot, the red curve shows the true μ^+ distribution whilst the blue curve shows the distribution accepted by the detector. The bottom panels show the ratio of accepted to total μ^+ for the detector.	152
	(a) Azimuthal acceptance of the 24 calorimeters	152

(b)	Azimuthal acceptance of the 3 trackers	152
7.8	Heat map showing the decay vertex positions of those μ^+ whose decay e^+ are tracked in a single tracker station located at $x \sim 6800$ mm, $z \sim 100$ mm. The perspective is top-down on the ring.	153
7.9	Detector e^+ energy acceptance. In the top panel of each plot, the red curve shows the true μ^+ distribution whilst the blue curve shows the distribution accepted by the detector. The bottom panels show the ratio of accepted to total μ^+ for the detector. The distributions are normalised to a single detector station.	154
(a)	Calorimeter e^+ energy acceptance	154
(b)	Tracker e^+ energy acceptance	154
7.10	The stored μ^+ beam profile in the simulation, as determined from the decay vertices of the μ^+ . The colour scale is the number of μ^+ in each $[x, y]$ bin, normalised to 1.	155
7.11	Detector radial decay vertex acceptance for e^+ above the energy cut. In the top panel of each plot, the red curve shows the true μ^+ distribution whilst the blue curve shows the distribution accepted by the detector. The bottom panels show the ratio of accepted to total μ^+ for the detector. The distributions are normalised to a single detector station.	155
(a)	Calorimeter radial decay vertex acceptance	155
(b)	Tracker radial decay vertex acceptance	155
7.12	Detector vertical decay vertex acceptance for e^+ above the energy cut. In the top panel of each plot, the red curve shows the true μ^+ distribution whilst the blue curve shows the distribution accepted by the detector. The bottom panels show the ratio of accepted to total μ^+ for the detector. The distributions are normalised to a single detector station.	156
(a)	Calorimeter vertical decay vertex acceptance	156
(b)	Tracker vertical decay vertex acceptance	156

7.13	156
(a)	Calorimeter versus tracker relative radial (red) and vertical (blue) decay vertex acceptance. The error bars show the statistical uncertainty in these distributions, which increases away from the centre of the storage region due to the fall off in numbers of μ^+ orbiting in these fringe regions.	156
(b)	Tracker-to-calorimeter beam profile translation map, where the colour scale shows the magnitude of the translation to be applied to the μ^+ counts measured by the tracker at a given $[x, y]$ coordinate.	156
7.14	Probability of a calorimeter e^+ hit having a matching track in the upstream tracker, plotted as a function of hit position on the calorimeter front-face. The colour scale is the probability of a hit at that position. The view is downstream into the calorimeter, with the storage region on the left.	158

List of tables

5.1	Data rates in interfaces between components in the DAQ system. The transmission (Tx) protocols and corresponding data transmission rates are shown. Payload fraction refers to the fraction of each data packet containing payload data, where the remaining packet capacity is occupied by information required by the protocol.	99
-----	---	----

Chapter 1

Introduction

In the last century, precision measurements of the magnetic moments of subatomic particles have played a fundamental role in the development of Quantum Field Theories (QFT) and the Standard Model (SM) of particle physics [3]. Early measurements of the electron magnetic moment provided key evidence supporting Dirac's relativistic theory of spin- $\frac{1}{2}$ particles, and subsequent higher precision measurements led to the discovery of vacuum polarisation and the understanding of quantum loop processes in the SM. Additionally, proton and neutron magnetic moment measurements indicated the composite nature of these particles decades before experimental evidence for quarks. To this day, measurements of the contribution of quantum loop processes to magnetic moments, known as anomalous magnetic moments, a , are amongst the most precise performed in the field of particle physics. Combined with detailed theoretical calculations, to the five-loop level for leptons, these provide stringent tests of the SM.

Lepton anomalous magnetic moments are of particular interest, as the absence of hadronic contributions in the dominant lower order loop processes means that they are minimally affected by the theoretical difficulties in Quantum ChromoDynamics (QCD) calculations at low energies. This makes precise theoretical calculations possible, and the existence of stable or relatively long lived leptons such as the electron and muon allow precision measurements to be made.

The sensitivity of charged lepton anomalous magnetic moments to new physics is $\propto m^2$, resulting in the heavier muon being 40000 times more sensitive to new physics than the electron [4]. This has fueled a decades-long campaign of increasingly precise muon anomalous magnet moment, a_μ , measurements at Columbia, CERN

and the Brookhaven National Laboratory (BNL). The most recent and precise of these measurements, from BNL, achieved a measurement precision of 540 ppb. This measured a_μ value deviates from the SM prediction by over 3σ . Enticingly, one possible explanation for this discrepancy is a contribution to a_μ from new physics, and this deviation is presently one of the most significant hints of Beyond the SM (BSM) physics known today.

Although this deviation is appealing, it falls short of the 5σ discovery threshold normally required in the field of particle physics. A new a_μ experiment is currently being constructed at the Fermi National Accelerator Laboratory (Fermilab) with the aim of increasing the precision of the a_μ measurement by a factor of 4 [2]. This is expected to result in conclusively disfavouring a purely SM value of a_μ , if the observed BNL deviation is indeed the result of new physics.

The Fermilab muon g-2 experiment can be considered a direct evolution of the BNL concept, using the same experimental principle of observing the Larmor precession of the spins of μ^+ orbiting in a 1.45 T, 14 m diameter magnetic storage-ring. The BNL storage-ring is being re-used for the Fermilab experiment following a complex move from New York to Illinois. A new source of μ^+ that is more intense and with a lower pion contamination will be provided by the Fermilab accelerator complex, and new or upgraded detector systems will be used. The magnetic field in the storage-ring will be more uniform owing to an extensive magnet shimming campaign, and the field measurement precision will be significantly improved.

Amongst these new detector systems are three straw tracking detectors consisting of 2.5 mm radius mylar straws containing a central sense wire held at high voltage and filled with an argon-ethane gas mixture. These trackers will reconstruct the trajectories of the e^+ resulting from μ^+ decays in the storage-ring, and by extrapolation build up a spatial profile of the stored μ^+ beam. The trackers will also be able to observe time-dependent dynamics in the beam including betatron oscillations, and quantify the impact these effects have on the measured a_μ value and allow corrections to be made. The trackers also complement other detector systems in the experiment such as the calorimeters, for example by providing independent measurements of particle momentum, and by identifying pile-up events.

The use of trackers in the experiment also allows the possibility of a separate physics measurement in addition to a_μ . Unlike the magnetic dipole moment, the electric

dipole moment (EDM) of charged leptons is predicted to be effectively zero in the SM, as a non-zero EDM would violate CP-symmetry. A non-zero EDM would influence the precession of μ^+ spins in the g-2 storage-ring, tilting the spin precession plane away from the muon orbit plane. This tilt would produce a vertical modulation in the average emission angle of the e^+ from μ^+ decay which could be observed by the trackers. The trackers are expected to allow the world's most sensitive measurement of the muon EDM to be made.

The design and development of the tracking detectors is detailed in this thesis, and results from detector testing and characterisation are presented. Prototype and production tracker modules have been tested in a number of environments, including three test beams at the Fermilab Test Beam Facility (FTBF), with cosmic-rays in a vertical test stand with scintillator triggers, and using radioactive sources. These tests have provided measurements of the straw-hit resolution and efficiency, determined the optimum operating parameters for the detectors, and have probed the relationship between the drifting of charge in the straw gas and the distance of closest approach of the particle track to the straw wire, which is fundamental to the accurate reconstruction of particle tracks. Detector data is compared to detailed simulations of the straw gas and electronics properties using the **GARFIELD** gas-detector simulation software.

The tracker readout system and online and offline data processing systems are also described. These systems have been developed and tested in parallel to the detector itself, and have been shown to meet the speed and reliability demands of the experiment.

Finally, Monte-Carlo simulations of the full Fermilab muon g-2 experiment using the **GEANT4** detector simulation toolkit are presented, along with an analysis of the performance of the trackers in these simulations to demonstrate that their performance meets the physics requirements. Measurements of straw occupancy and hit rates, and the tracker acceptance with respect to a number of correlated μ^+ orbit and decay parameters are presented. The relative acceptance with respect to the experiment's calorimetry systems are also presented, quantifying the translation that must be applied to interpret tracker beam profile measurements from the perspective of other detector systems.

Original contributions to the Fermilab muon g-2 experiment by this author cover a range of areas. These include the development of online software, including the

tracker data acquisition and online-monitoring software, and contributions to the experiment-wide MIDAS run control and data acquisition system, as well as the development of offline software, including the tracker low-level offline software, and experiment-wide MIDAS-to-art online-to-offline data conversion software. Another major area of work has been testing the tracker at a number of test beams and test stands, where contributions have included setting up the hardware and software, development of the auxiliary detector DAQ systems, conducting the testing and operating the detectors, and the analysis of the resulting data and the development of supporting Monte-Carlo simulations. Another area of significant work was in the development of a GARFIELD simulation of the response of the straw gas, wire and electronics to charged particles traversing the straw. An active role has also been played in the development of the experiment's GEANT4 simulation, with contributions including the straw readout model, the mechanism for recording particle trajectories, and the handling of coordinate-systems and transformations. Physics studies have also been performed using the simulated data.

Chapter 2

The theory of lepton anomalous magnetic moments

2.1 Lepton magnetic moments

Magnetic moments characterise the torque experienced by charges in an external magnetic field, such as the well known classical example of a compass needle aligning with the Earth's magnetic field. Magnetisation is a result of the motions of electric charges, known as currents. Classically, this results in a magnetic field in a solenoid, where current flowing through loops of wire produces a magnetic field in the direction perpendicular to these loops. At the scale of individual particles, a charged particle undergoing orbital motion acts like the current loop in a solenoid, and the induced magnetisation results in a magnetic moment in the orbiting particle. For a charged particle, there is a fundamental relation between angular momentum and magnetic moments, characterised by the gyromagnetic ratio, g , also referred to as the g -factor.

In quantum mechanics, particles have an intrinsic form of angular momentum known as spin. Charged-particles with non-zero spin will thus have an intrinsic magnetic moment, μ , even at rest [5]. For a spin- $\frac{1}{2}$ charged-particle such as a lepton, quark, or proton, this is expressed by:

$$\boldsymbol{\mu} = g \frac{q}{2m} \mathbf{s}, \quad (2.1)$$

where g is the electric charge of the particle in units of the electron charge, m is the mass of the particle and \mathbf{s} is the particle's spin.

Following the then recent development of the quantum description of spin, Dirac predicted that $g = 2$ for fundamental spin- $\frac{1}{2}$ charged-particles, such as the electron [3]. Measurements of the magnetic moment of the proton, at the time considered to be a point-like Dirac particle, by Frisch, Stern and Estermann in 1933 yielded a surprising result 2.5 times larger than expected for $g_p = 2$, with 10 % measurement uncertainty [6, 7]. Combined with the subsequent inference of an unexpected non-zero magnetic moment for the neutral neutron from deuteron magnetic moment measurements in 1934 by Estermann and Stern, and Rabi, Kellogg and Zacharias, this served as the first experimental indication of the composite nature of these nucleons [8, 9]. For the fundamental electron however, experiments at the time appeared to confirm Dirac's prediction of $g_e = 2$ within the limitations of their measurement uncertainty. Subsequently however, more precise measurements during the 1940s began to indicate a deviation for the electron magnetic moment, culminating in the 1947 Kusch and Foley experiment finding a $(0.119 \pm 0.0005)\%$ discrepancy from $g_e = 2$ [10].

A theoretical explanation for the electron g_e discrepancy soon followed as the theory of Quantum Electrodynamics (QED) was developed. In QED, the electromagnetic interaction includes contributions from vacuum polarisation effects of short-lived virtual-loop particles. Dirac's prediction is exactly correct before these vacuum polarisation effects are taken into account, or equivalently for the tree-level (and dominant) Feynman diagram of the electromagnetic vertex. Schwinger calculated the leading-order (LO), one-loop correction to g_e to be $\frac{\alpha}{\pi}$, where α is the electromagnetic coupling constant, which was in agreement with the experimentally observed deviation [4]. Figure 2.1 (a) shows the Feynman diagram describing the Dirac term, whilst 2.1 (b) shows the Schwinger term. This work pioneered the theoretical understanding of vacuum polarisation effects that now underpins QFTs, and to this day successive generations of ultra-precise measurements of the electron's magnetic dipole moment provide one of the most stringent tests of the SM [11].

The contributions of beyond tree-level processes to g for a particle are now commonly referred to as the anomalous magnetic moment, a , which is defined for a charged lepton as:




Figure 2.1: Feynman diagrams at (a) tree-level (Dirac) and (b) one-loop level (Schwinger) of the electromagnetic interactions determining the value of g_e . In both cases the electron couples to an external magnetic field.

$$a_l = \frac{g_l - 2}{2}, \quad (2.2)$$

where l represents the lepton flavour and thus takes the values $l = e, \mu, \tau$.

Equivalently, the so-called anomalous contribution is also often referred to simply as $(g - 2)_l$.

Measurements of a_l directly probe the large variety of loop-diagrams contributing to the fundamental electromagnetic vertex. Beyond $\mathcal{O}(\alpha)$, the Schwinger term, a broad spectrum of particles and forces are present in the loop-diagrams, and so a broad range of SM physics is tested by a_l measurements. Strong constraints are also imposed on new physics models, many of which would result in additional loop processes contributing to a_l .

Lepton anomalous magnetic moment measurements are particularly attractive since QCD contributions only appear at $\mathcal{O}(\alpha^2)$ and thus are strongly suppressed, allowing sub-ppm theoretical predictions to be made. This has motivated a series of increasingly precise a_l measurements over the past 50 years. All these measurements have exploited the Larmor precession experienced by charged leptons in an external magnetic field (see section 3.2), with the resulting a_l measurements being the most precise performed in particle physics. For the electron, a_e has been measured to a precision of 0.28 ppt, with this measurement also providing the present most precise value of α [11]. The measured a_e value is in agreement with the SM value, although the precision of this comparison is limited by the requirement to use an independent

α measurement to calculate the SM value, which at present is 20 times less precise than the value obtained from the a_e measurement and dominates the uncertainty. Future independent α measurements with higher precision [12] will allow a more precise comparison of the measured a_e with the SM value, hence improving the sensitivity to new physics.

There is also a long history of a_μ measurements which persist to this today, and which are the subject of the next section.

2.2 The anomalous magnetic moment of the muon

As one of the key motivations for precision measurements of a_l is to search for and constrain new physics, it should be considered what sensitivity a_l has to new physics models. Helicity considerations imply that the impact on a_l due to additional loop contributions from new physics, δa_l , can be expressed as [4]:

$$\delta a_l \propto \frac{m_l^2}{\Lambda^2}, \quad (2.3)$$

where Λ is the energy scale of new physics and m_l is the lepton mass, and it is assumed that $\Lambda \gg m_l$.

The energy scale Λ could represent the mass of a new particle, an energy scale where the SM is no longer valid, or the scale at which the lepton is actually a composite, and not point-like, particle [4].

Equation 2.3 indicates that the sensitivity of a_l increases rapidly with increasing lepton mass. As the muon mass is approximately 200 times that of the electron, a_μ is thus approximately 40000 times more sensitive to new physics than a_e when both are measured to the same precision. This makes a_μ an appealing candidate measurement for new physics searches, and a number of precision a_μ measurements have been made. The tau lepton would in principle be more sensitive than the muon to new physics due to its larger mass. However, the short lifetime of the tau and difficulties producing them in large numbers has prohibited dedicated precise a_τ measurements to date [13], although limits have been imposed at $e^+ e^-$ colliders such as LEP [14].

These a_τ measurements are at present not yet even sensitive enough to probe the $\mathcal{O}(\alpha)$ Schwinger contribution however. With current technology and the limited precision of non- a_e α measurements, the muon is currently the favoured lepton for searching for new physics through anomalous magnetic moment measurements.

The most precise a_μ measurement to date was achieved by the Brookhaven National Laboratory (BNL) muon g-2 experiment, which ran between 1997 and 2001 and measured a_μ to a precision of 540 ppb [1]. The experiment utilised a large magnetic storage-ring in which orbiting muons were stored and the precession of their spin vectors in the magnetic field due to the muon magnetic moment was observed. Strikingly, this measurement deviates from the SM prediction by $3.3 - 3.6\sigma$, where the range here results from differing approaches to determining hadronic contributions to the theoretical value of a_μ [2]. One possible cause of this deviation from the SM expectation is contributions to a_μ from new physics, but the significance of the deviation falls short of the 5σ discovery threshold normally required in the field of particle physics.

Two new experimental efforts are planned to further probe and hopefully unambiguously confirm or reject this discrepancy. Both utilise the same fundamental measurement principle as the BNL experiment, where the correlation between the spin vector of a muon and the momentum of its decay products due to parity violation in the weak decay is exploited to observe the precession of the spin vector in a magnetic field that results from the muon's intrinsic magnetic moment. The first is based at the Fermi National Accelerator Laboratory (Fermilab), and is a direct evolution of the BNL experiment utilising the same magnetic storage-ring but a new muon source and new detector systems [2]. Construction of the experiment is now almost complete, and data taking will commence in June 2017 with a target of reaching an a_μ measurement precision of 140 ppb. With two years of data this is sufficient to confirm the observed a_μ discrepancy with respect to theory beyond the 5σ discovery threshold, assuming no change in the central value or improvement in the precision of the SM prediction. This Fermilab experiment is the subject of this thesis and is described further in chapter 3.

The second experiment is based at the Japan Proton Accelerator Research Complex (J-PARC), and differs significantly from the Fermilab and BNL experiments despite the common fundamental measurement principle. This experiment uses an ultra-cold muon beam to eliminate the need for the focusing electric fields employed at

Fermilab and BNL, and the resulting lower energy muons allow a smaller storage-ring to be used [15]. The injection of the muons into the storage-ring also follows a radically different scheme compared to BNL and Fermilab, utilising a vertical rather than a horizontal kicker field. These two experiments will have very different systematic uncertainties, and having both measurements provides a valuable cross-check, particularly if the measured values continue to deviate from the SM value. The J-PARC experiment is currently in a phase of intensive research and development to demonstrate that an ultra-cold muon beam of the required intensity and emittance can be injected into the experiment's MRI magnet, and data taking is not expected to begin until the end of Fermilab g-2 data taking [16]. The experiment aims to reach a statistical precision matching Fermilab, but the systematic uncertainty estimate is not yet finalised.

It is also expected that the uncertainty in the SM theory value of a_μ will be reduced by approximately a factor of 2 on a timescale commensurate with the experiment efforts, allowing an even more precise comparison between theory and experiment [2].

2.3 Electric dipole moments

Dirac's theory of relativistic particles allows the possibility of an intrinsic electric dipole moment (EDM) in spin- $\frac{1}{2}$ particles as well as a magnetic dipole moment. This would result in particles experiencing a torque in the presence of an electric field [2]. An EDM, d , can be expressed in a very similar way to μ , and characterised by a dimensionless constant η , that is analogous to g :

$$\mathbf{d} = \eta \frac{q}{2mc} \mathbf{s}, \quad (2.4)$$

where c is the speed of light in a vacuum.

Although Dirac's theory makes provision for an EDM, this does not imply that one exists in nature.

The Hamiltonian, \mathcal{H} , for a spin- $\frac{1}{2}$ particle under the influence of electric and magnetic fields can be expressed as:

$$\mathcal{H} = -\boldsymbol{\mu} \cdot \mathbf{B} - \mathbf{d} \cdot \mathbf{E}, \quad (2.5)$$

where \mathbf{E} is the electric field vector.

The dipole moments $\boldsymbol{\mu}$ and \mathbf{d} flip sign under time reversal (T) transformations. The magnetic field vector also changes sign under T, whilst the electric field vector does not. This means that the magnetic dipole moment term, $\boldsymbol{\mu} \cdot \mathbf{B}$, is invariant under T whilst the EDM term, $\mathbf{d} \cdot \mathbf{E}$, is not [4]. Assuming CPT invariance, T-symmetry violation implies CP-symmetry violation, and hence EDMs cannot exist in any theory that is invariant under CP, implying that $d = 0$ in such theories.

At present, CP violation has only been conclusively observed in weak interactions in the quark sector, although searches for CP violation in neutrino oscillations are also underway [17]. EDMs in charged leptons can only be generated by highly suppressed higher-order virtual-loop contributions to the charged lepton electromagnetic vertex, such as quark loops. Charged lepton EDMs are thus predicted to be unmeasurably small in the SM. The observation of an EDM larger than the SM prediction in the charged lepton sector can only arise from BSM physics, and would be an additional source of CP violation and so potentially could help explain the observed matter-antimatter asymmetry of the Universe: one of the major outstanding problems facing the SM and particle physics today.

Storage-ring anomalous magnetic moment experiments such as the Fermilab muon g-2 experiment are also sensitive to non-zero EDMs, which would be manifest as additional, orthogonal components in the spin vector precession that these experiments observe. The current world's best limit on the muon EDM was set by the BNL muon g-2 experiment, and this will be significantly improved upon by the Fermilab muon g-2 experiment [18]. This is discussed further in section 3.3.8.

2.4 Standard model value of a_μ

The high precision to which a_μ has been and will be measured demands corresponding precision in the theoretical prediction. This demands a challenging calculation of loop-

diagrams to at least $\mathcal{O}(\alpha^4)$. The SM value of a_μ can be subdivided into contributions from loop processes involving QED, electroweak (EW) and hadronic interactions:

$$a_\mu^{\text{SM}} = a_\mu^{\text{QED}} + a_\mu^{\text{EW}} + a_\mu^{\text{Hadronic}}. \quad (2.6)$$

The QED contributions dominate the value of a_μ , but are very precisely predicted; conversely the challenging nature of calculating low-energy hadronic interactions means that the hadronic contribution dominates the uncertainty. The details of each contribution are outlined in the following sections.

2.4.1 QED contributions to a_μ

Pure QED contributions to a_μ involve virtual-loops consisting of photons and charged leptons, where lepton flavour is conserved and no coloured particles are present [19]. The sole LO contribution, first calculated by Schwinger, is shown in figure 2.2 (a). Nine next-to-leading-order (NLO) two-loop processes exist, where one example is given in figure 2.2 (b) [3].




Figure 2.2: (a) LO and (b) NLO QED Feynman diagrams contributing to a_μ [2].

It was required that a_μ^{QED} be determined to $\mathcal{O}(\alpha^4)$ to reach theoretical precision commensurate with the BNL experimental precision [20], and this has since been surpassed by an $\mathcal{O}(\alpha^5)$ calculation of more than 10000 Feynman diagrams [21]. Thanks to these efforts, despite the QED contribution being by far the largest term in equation 2.6 it contributes the smallest uncertainty [22].

2.4.2 Electroweak contributions to a_μ

Electroweak contributions to a_μ result from loops containing the weak massive W and Z bosons, or the Higgs boson. The LO processes are at the one-loop level and are shown in figure 2.3. The large mass of the W and Z bosons and the small Yukawa coupling of the Higgs boson to the muon result in a_μ^{EW} being a small contribution to a_μ , and at the LO only the W and Z boson loops contribute significantly [2]. The one and two-loop processes contributing to a_μ^{EW} have been calculated, yielding a theoretical uncertainty comparable with the QED contribution [23].




Figure 2.3: LO Feynman diagrams of the EW processes contributing to a_μ [3]. There are contributions from the W (a) and Z (b) massive vector bosons and (c) the Higgs boson.

2.4.3 Hadronic contributions to a_μ

Hadronic contributions to a_μ arise when coloured particles such as quarks are produced in virtual-loops, with the LO contribution at the two-loop level. The LO Feynman diagram is shown in figure 2.4, and a number of possible higher-order (HO) diagrams are shown in figure 2.5.

Calculating these contributions poses a theoretical challenge, as the theory governing the behaviour of coloured particles, QCD, is non-perturbative at the energy scales at which the dominant hadronic contributions to a_μ occur [3]. Perturbation theory can therefore not be applied to calculate the loop contributions as is the case for QED and EW loops, and other methods have instead been pursued to determine a_μ^{Hadronic} , including data-driven methods using experimental data.

The first of these techniques exploits commonalities between the LO hadronic contribution to a_μ and the LO contribution to the process $e^+e^- \rightarrow \text{hadrons}$ that is

measured in e^+e^- experiments [3]. These two processes are related via a so-called dispersion-relation such that experimental e^+e^- annihilation data can be used to infer the LO a_μ^{Hadronic} contribution [2]. This technique can be extended to determine the contributions from some, but not all, HO loop processes. Data from the KLOE, BaBar, BELLE, VEPP and BES experiments have been used to compute a_μ^{Hadronic} [24]. An alternative data-driven technique utilising measurements from tau lepton decays to hadrons, also from e^+e^- colliders, has also been used to estimate the LO a_μ^{Hadronic} contribution, and the latest results are found to be in good agreement with $e^+e^- \rightarrow \text{hadrons}$ results [25].




Figure 2.4: LO Feynman diagram of the hadronic contribution to a_μ , where H represents any hadron [2].

The LO hadronic contribution has also been calculated using lattice QCD techniques, and is found to be in good agreement with the value from data-driven methods [26].




Figure 2.5: Example higher (beyond leading) order Feynman diagrams for hadronic contributions to a_μ , where H represents any hadron [2].

The LO processes dominate the value of a_μ^{Hadronic} , and contribute significant uncertainty. However, a particular class of HO processes known as hadronic Light-by-Light (LbL) scattering upon which the dispersion-relation techniques cannot be applied also contribute a significant uncertainty [5]. The form of the hadronic LbL processes is shown in figure 2.6. The current best estimate of these contributions is an internationally agreed consensus on combining the results of various model-dependent

approaches [27]. The resulting uncertainty remains relatively large: although LbL scattering contributions correspond to less than 2 % of the total value of a_μ^{Hadronic} , they contribute nearly 40 % of the uncertainty.




Figure 2.6: Feynman diagram of an example light-by-light hadronic contribution to a_μ , where H represents any hadron [2].

2.4.4 Value and uncertainty of a_μ^{SM}

Collecting the QED, EW and hadronic contributions to a_μ yields a SM theory value of [28]:

$$a_\mu^{\text{SM}} = (116591763.7 \pm 51.8) \times 10^{-11}. \quad (2.7)$$

This represents a theoretical precision of 440 ppb. The relative weightings of these contributions to the value of a_μ and their uncertainty are shown in figure 2.7, clearly showing the dominance of the QED term to the value and the hadronic terms to the uncertainty.

Improvement in the uncertainty of a_μ^{Hadronic} is expected on a time scale comparable with the Fermilab muon g-2 experiment a_μ measurement. In particular, lattice QCD techniques are expected to become increasingly competitive for hadronic LbL and hadronic LO contributions [29], and new $e^+ e^-$ collider data from the VEPP-2000 [30], BESIII [31] and KLOE-2 [32] experiments will reduce the uncertainty in the data-driven techniques.




Figure 2.7: SM contributions to the (a) value and (b) uncertainty of a_μ .

2.5 Possible new physics contributions to a_μ

As previously stated, one possible explanation for the observed discrepancy between the BNL a_μ measurement and the SM prediction could be contributions to a_μ from BSM physics. Any BSM interaction contributing to a_μ must conserve CP and flavour, and flip the chirality of the muon. In general the contribution to a_μ decreases as the mass scale of the BSM physics increases. In the majority of BSM scenarios, the BNL deviation suggests an energy scale of the order of a TeV or less for the new physics [33].

A large range of BSM models have been considered to explain the observed a_μ discrepancy, including supersymmetry (SUSY) [34, 35], new EW bosons such as W' or Z' [36, 37], alternative Higgs models [38, 39], and dark sector models [40, 41]. In many cases the phase-space of these models probed by a_μ measurements is complementary to other experiments, such as those at the LHC.

One example from the Higgs sector is the two-Higgs doublet model (H2DM), where electroweak symmetry breaking is minimally extended to include two SU(2) doublets [39]. This model predicts five Higgs bosons, one of which is compatible with the SM Higgs boson. These additional bosons can contribute to a_μ via loop-diagrams in a similar manner to other electroweak bosons, as indicated in figure 2.3. Figure 2.8 indicates regions of H2DM parameter space that are consistent with the measured a_μ deviation with respect to the SM. These regions are not ruled out by collider

results, and as such constraints from a_μ are complementary to direct searches for the extended Higgs boson sector at the LHC.




Figure 2.8: Maximum possible value of a_μ in the flavour-aligned H2DM, shown as a function of the mass M_A of one of the additional Higgs bosons in this model, A , and for three values of the lepton Yukawa coupling, ζ_l , where these parameters are found to be those that most strongly drive large predictions for a_μ [39]. The masses of the remaining three Higgs boson states in this model (excluding A and the SM model Higgs) are set to 250 GeV in this plot. The yellow band indicates the value of a_μ measured at BNL and its uncertainty. The H2DM contribution to a_μ is computed using a two-loop calculation.

The improved precision of the Fermilab a_μ measurement will further constrain the parameters of the BSM theories, and in the event of the deviation with respect to the SM being confirmed beyond 5σ it would be a landmark observation in particle physics that would change the landscape just as the Kusch and Foley measurement transformed QED.

Chapter 3

The Fermilab muon g-2 experiment

3.1 Experimental goals

The E989 Fermilab muon g-2 experiment [2], due to commence data taking in 2017, aims to measure the anomalous magnetic moment of the muon, a_μ , to a world's best precision of 140 ppb. This represents a factor 4 improvement compared to the current world's best precision of 540 ppb from the E821 g-2 experiment at BNL [1]. This measurement is made by observing the precession of the spins of μ^+ stored in a magnetic storage-ring due to the interaction between their intrinsic magnetic moments and the external magnetic field. The goal of this improved measurement is to reduce the uncertainty such that the present BNL deviation of 3.6σ in a_μ relative to the SM value increases to 5σ , should the measured value remain the same, or else rule out a deviation at this magnitude and set a new limit on the magnitude of any new physics contribution to a_μ .

The improvement in measurement precision for the Fermilab muon g-2 experiment will be achieved in four main ways:

- An increase in statistics brought about by storing a larger integrated number of muons in the storage-ring. This will be achieved using the higher intensity μ^+ beam with a far lower pion contamination provided by the Fermilab accelerator complex, and by increasing the fraction of muons injected into the ring that are stored via improvements in the beam control systems within the ring [2].

- The uniformity of the magnetic field in the storage-ring will be improved and the uncertainty in the measurements of this field reduced.
- The systematic uncertainty in the measurement of the precession of μ^+ spin vectors via the detection of the μ^+ decay products will be reduced using new detector systems with improved spatial and timing resolution, and with enhanced calibration systems.
- A simulation campaign of significantly higher fidelity than was available at the BNL experiment shall be pursued that will allow systematic effects resulting from the beam dynamics and the detector performance to be identified and quantified.

The following sections describe the measurement principle that will be employed at Fermilab muon g-2 experiment and the various experimental systems required to make the measurement.

3.2 Measuring g-2

Section 2.2 outlined the theoretical interest in precision measurements of a_μ , or equivalently $(g - 2)_\mu$. This section will describe a technique for making these precision measurements. The Fermilab muon g-2 experiment uses the same measurement principle as the BNL experiment, which in turn can be considered an evolution of previous measurements made at CERN [42], and it is this technique that is described below.

A particle with a non-zero magnetic moment experiences a torque in the presence of a magnetic field, resulting in a precession of the particle's spin axis about the magnetic field vector, known as Larmor precession [43]. In the classical limit, such as for non-relativistic particles, the angular frequency of this precession, ω_L , is given by:

$$\omega_L = g \left(\frac{e\mathbf{B}}{2m} \right), \quad (3.1)$$

where g is the gyromagnetic ratio, e is the electric charge of the particle in units of the electron charge, \mathbf{B} is the magnetic field, and m is the mass of the particle.

In principle, a measurement of ω_L could be used to determine g . However, the spin vector direction and hence spin precession are not directly observable quantities, and a measurable signal must be sought. For the case of the muon, the parity violating nature of the muon's weak decay can be exploited to determine information about the muon spin direction at the time of decay from the momentum direction of its decay products. The mechanics of this relation are discussed now.

The Fermilab muon g-2 experiment will predominantly utilise μ^+ rather than μ^- for its a_μ measurement. The μ^+ are produced at Fermilab from a pion beam which itself is produced via the interaction of a proton beam with a fixed target. These pions then decay to μ^-/μ^+ . The interaction of the positively charged proton with the target yields a greater number of π^+ than π^- , and the π^+ predominantly decays to μ^+ [19]. The Fermilab accelerator complex thus produces μ^+ at a higher rate than μ^- . Only a single type of muon is used in the storage-ring at a given time, and as the μ^+ production rate is greater, the Fermilab muon g-2 experiment will initially run with stored μ^+ . In the subsequent sections μ^+ are thus considered.

The μ^+ decays almost exclusively via a single decay channel [19]:

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu. \quad (3.2)$$

This decay occurs via the weak interaction, which is observed to be parity violating. This means that the weak force couples only to left-handed particles and right-handed anti-particles. Handedness here strictly refers to chirality, but in the massless limit this is equivalent to helicity, the projection of a particle's spin onto its momentum. For massive particles, a given chiral state contains both left- and right-handed helicity components whose relative weighting depends on the mass. In terms of helicity, a left-handed particle is one whose spin and momentum are anti-parallel, and for right-handed particles they are parallel.

In the μ^+ decay the ν_e , a particle, and the $\bar{\nu}_\mu$, an antiparticle, can be treated as effectively massless, meaning that they are solely produced in left-handed and right-

handed helicity states respectively. The e^+ is an antiparticle and so the right-handed helicity state is preferentially produced by the weak interaction, although as it is not massless there is still a small left-handed component. There is thus a preferred handedness for all decay products, and so a preferred orientation of the spin vectors relative to the momentum vectors in this decay.

As well as helicity considerations, spin must be conserved in this decay. All particles in the interaction are spin- $\frac{1}{2}$ fermions. The kinematic configuration that conserves spin whilst maximally satisfying the preferred handedness is where the e^+ is emitted in the opposite direction to the two ν . This case is shown in figure 3.1. These decay kinematics are thus preferred in the decay. The e^+ energy is also maximal in this configuration, as can be seen by considering momentum conservation.

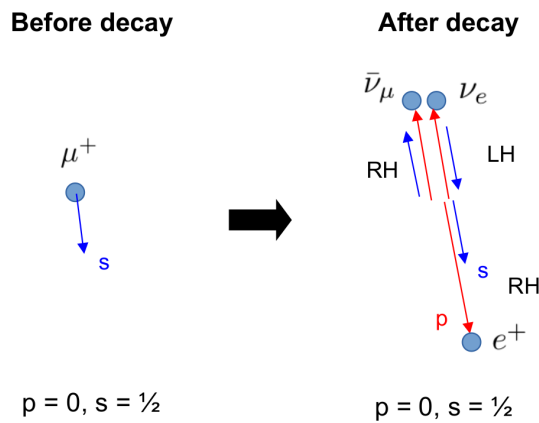


Figure 3.1: μ^+ decay kinematics, with the post-decay state shown being the preferred state due to parity violation in the weak decay and spin and momentum conservation. LH and RH refer to left- and right-handed particles respectively. Red arrows represent the momentum vectors and blue arrows represent the spin vectors.

Returning to the question of observing the precession of the μ^+ spin vector, it can be also be seen from figure 3.1 that in this preferred decay, the e^+ spin direction is parallel to that of the parent μ^+ . That is, the spins of the two ν cancel, leaving the e^+ with the same spin as the μ^+ via spin angular momentum conservation. As the e^+ is right-handed, its momentum vector is aligned with its spin vector, and thus the e^+ momentum vector is parallel to the direction of the μ^+ spin vector at the time of the decay. The decay e^+ are thus preferentially emitted along the μ^+ spin vector axes, as can be seen in figure 3.2, and hence the average direction of the emitted

e^+ will precess as the μ^+ spin precesses. This is an observable quantity that can be used to probe the μ^+ spin precession.




Figure 3.2: The e^+ emission angle in μ^+ decays, expressed relative to the μ^+ spin direction in the μ^+ rest-frame.

In order to observe Larmor precession, the μ^+ must be in an external magnetic field. Charged particles moving in a magnetic field will undergo cyclotron motion, and for a constant magnetic field the particle will follow a circular path of constant radius. The cyclotron angular frequency, ω_c , for a non-relativistic particle is given by:

$$\omega_c = \frac{e\mathbf{B}}{m}. \quad (3.3)$$

A μ^+ in a magnetic field is thus subject to two simultaneous precession effects: its spin vector precesses with the Larmor frequency, whilst its momentum vector precesses at the cyclotron frequency. The relative orientation of these vectors in the storage-ring is shown in figure 3.3. The difference between these two precession frequencies, i.e. the frequency at which the spin vector precesses relative to the momentum vector, is referred to as the anomaly frequency, ω_a :

$$\omega_a = \omega_L - \omega_c = \left(\frac{g}{2} - 1\right) \frac{e\mathbf{B}}{m} = a_\mu \frac{e\mathbf{B}}{m}. \quad (3.4)$$




Figure 3.3: The orientation of the μ^+ momentum, p , and spin, s , vectors as they orbit in the storage-ring, as seen from above the ring. The magnetic field, B , vector is also shown. The angle between the spin and momentum vectors, θ_a , is known as the anomaly angle. The rate of change of θ_a is known as the anomaly frequency, ω_a , and is one of the fundamental quantities measured in the Fermilab muon g-2 experiment. The circular motion of the spin vector is shown by the dashed blue line.

An important property of ω_a is that it depends only on the anomalous magnetic moment of the muon, a_μ , rather than g_μ . If a signal varying with ω_a can be measured, then a_μ can be directly determined and thus only the quantum-loop effects will be measured.

Such a signal is found if the μ^+ are moving with relativistic velocities. In this case, the decay e^+ will experience a Lorentz boost in the laboratory frame in the μ^+ momentum direction. e^+ emitted in the same direction as the μ^+ momentum will thus have the maximal average energy, and those emitted in the opposite direction will have the minimal average energy. As the e^+ emission angle is correlated with the μ^+ spin direction, this means that the average e^+ energy depends on the angle between the μ^+ spin and momentum vectors, known as the anomaly angle, which varies with the anomaly frequency. This means that for a population of μ^+ , the number of e^+ above some energy cut seen by an observer will vary periodically with time accordingly to the anomaly frequency, with the maxima when the μ^+ spin vector is parallel to the momentum vector and the minima when they are anti-parallel. This counting of e^+ above an energy threshold is the fundamental measurement of the modern a_μ experiments, including the CERN-III, BNL and Fermilab measurements.

It should be noted that in using this technique, there are only two unknowns that must be measured to determine a_μ : ω_a and B . The Fermilab muon g-2 experiment is divided into systems for the precise measurement of these two parameters, with the ω_a measurement being the subject of this thesis. It is also important to note that ω_a has a linear dependence on B . If this were not the case, e.g. if the dependence were quadratic, then non-uniformities in the magnetic field would have a significant effect.

In the relativistic domain the observed spin vector precession is also subject to an effect known as Thomas precession [44], whereby the particle rest-frame rotates relative to the observer frame. This gives an overall apparent spin precession frequency, ω_s of:

$$\omega_s = g \frac{e\mathbf{B}}{2m} + (1 - \gamma) \frac{e\mathbf{B}}{\gamma m}, \quad (3.5)$$

where γ is the Lorentz factor.

The cyclotron frequency is also modified by relativistic effects:

$$\omega_c = \frac{e\mathbf{B}}{\gamma m}. \quad (3.6)$$

These relativistic modifications to the spin and cyclotron precession frequencies cancel, resulting in an unchanged anomaly precession frequency compared to the non-relativistic case shown in equation 3.4.

When measuring a periodic effect, it is desirable to observe the effect over many periods to make an accurate measurement. The μ^+ is an unstable particle, and so it is required that ω_a is large enough that on average the μ^+ spin vector completes multiple rotations relative to its momentum vector before the μ^+ decays. An advantage of using relativistic μ^+ is that their lifetimes are extended by time dilation, increasing the number of periods that can be observed. The experiment uses 3.094 GeV μ^+ which have a lifetime of 64 μs . Equation 3.4 shows that using a strong magnetic field is also desirable to increase ω_a .

3.3 The Fermilab muon g-2 experiment

Section 3.2 described the fundamental principle of the g-2 measurement to be undertaken by the Fermilab muon g-2 experiment. This chapter describes the design of the experiment and how it will achieve this measurement.

The central component of the experiment is a 14 m diameter magnetic storage-ring, which is re-used from the previous BNL experiment. The ring can be seen in figure 3.4. Bunches of polarised μ^+ from the Fermilab accelerator complex are injected into the ring periodically, with each injection known as a “fill”. They enter a ring-shaped storage region under vacuum where a highly uniform 1.45 T vertical dipole magnetic field is present. Under the influence of this magnetic field, the μ^+ orbit within the ring whilst their spins precess until they decay. The e^+ produced in the decays have lower energies than their parent μ^+ and thus follow more highly curved trajectories in the magnetic field, and over some fraction of an orbital period exit the storage region on the inner side of the ring. This inner side is instrumented with detector systems, chiefly 24 calorimeters which detect the e^+ and measure their energies. The observed number of e^+ detected in these calorimeters above some energy threshold versus time modulates periodically with the frequency ω_a , as well as falling exponentially due to the decreasing population of μ^+ in the ring as they decay following each injection. An example of this behaviour can be seen in figure 3.5, and it is this signal that is used for the Fermilab a_μ measurement.

In reality there are of course a large number of details, subtleties and complications to this simple premise, which are the subject of the following sections.

3.3.1 Muon production and delivery to the storage-ring

μ^+ production at Fermilab will be achieved via the chain $p \rightarrow \pi^+ \rightarrow \mu^+$ [2]. 8 GeV protons from the Fermilab booster are injected into the former Tevatron recycler ring where each booster batch is separated into four distinct bunches. At the time of injection into the storage-ring, the temporal width of the μ^+ bunch is ~ 120 ns, which is less than the cyclotron period of 149 ns. Each of these bunches will form one fill of approximately 16000 μ^+ in the storage-ring. This reduced number relative to the original booster batch reduces the instantaneous hit rate in the experiment’s




Figure 3.4: The magnetic storage-ring during installation at Fermilab. The ring diameter is 14 m.

detectors, mitigating pile-up effects. Four booster batches per 1.4 s Main Injector (MI) accelerator cycle are available for muon g-2 use with the rest used by the Fermilab neutrino programme. The fills are typically concentrated at the start of the MI cycle with a separation of 10 ms, and the average fill rate is ~ 12 Hz.

Following this re-bunching, the protons are transported to a pion production target located at the former Tevatron antiproton target station [45]. Protons interacting with this target produce secondary particles, including charged and neutral pions. The target is a cylinder of Inconel that is resistant to the effects of heating produced by the incident beam [46]. The target includes an air cooling channel to mitigate heating effects, and an outer layer of beryllium to protect against oxidation whilst minimising scattering of secondaries. The positively charged secondaries produced by the target then pass through a lithium lens which provides transverse focusing, before momentum selection is performed using a pulsed magnet to select the π^+ that will yield μ^+ with the desired momentum downstream following their decay. These selected particles, comprising the desired π^+ and other secondaries (notably protons), then pass through a long beamline such that the majority of the π^+ decay. The resulting μ^+ dominated beam enters a small-ring beam-line, known as the delivery




Figure 3.5: Histograms of the times that the decay e^+ hit the calorimeters in a Monte-Carlo simulation. Hits from all e^+ are shown in red, whilst hits from e^+ with energies greater than 1.8 GeV are shown in blue. Both distributions have a falling exponential component resulting from the μ^+ decays, where the time dilated lifetime of the μ^+ is 64 μs . The energy cut introduces a sinusoidal component which varies with the anomaly frequency, 4.4 μs . The simulation includes effects such as momentum spread and time width of the injected beam.

ring, which is re-purposed from the Tevatron antiproton debuncher. The μ^+ orbit in this ring long enough for any remaining π^+ to decay, and also to force any protons remaining in the beam bunch to separate from the μ^+ due to their larger mass, allowing them to be removed from the beam using an abort kicker magnet. This process results in a pure population of μ^+ ready for injection into the g-2 storage-ring.

A requirement of the ω_a measurement is that the μ^+ delivered to the storage-ring must be initially polarised, as otherwise the anomaly precession for each μ^+ would be out of phase with the rest of the population, and the observed modulation of e^+ counts above an energy threshold versus time would be washed-out. To achieve this polarisation, the parity violating nature of the weak decay is exploited again here, now to produce a polarised μ^+ source directly from the π^+ decays [47, 48]. Positive pions resulting from the proton collisions with the target decay predominantly via the following channel [19]:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu. \quad (3.7)$$

Using similar arguments from parity violation as those made in section 3.2, the ν_μ must be in a left-handed helicity state, and spin conservation from the scalar pion implies a left-handed μ^+ in the pion rest-frame. Note that this is not the preferred handedness of an antiparticle in the weak interaction but is the only allowed possibility via the decay kinematics, and hence this decay is helicity suppressed. It is however less helicity suppressed than the corresponding decay to a e^+ due to the larger mass of the μ^+ . In the pion rest-frame, the left-handedness of the μ^+ means that its spin vector is anti-parallel to its momentum vector. There is no preference for any emission angle in the pion rest-frame. The pions emitted from the target decay in-flight, meaning that in the lab-frame the pion momentum is either added to or subtracted from the μ^+ momentum depending on whether the μ^+ is emitted forwards or backwards compared to the pion direction. The most forward μ^+ receive the greatest additional momentum contribution, and hence have the largest lab-frame momenta. As their spins are anti-parallel with their momenta, selecting the highest momenta μ^+ in the lab-frame selects those whose spin vectors are highly backwards pointing: this high-momenta population is thus polarised in the backwards direction. The same argument yields a forward polarised beam by selecting the backwards,

e.g. lowest energy, μ^+ . The forward μ^+ will be selected for the Fermilab muon g-2 experiment, resulting in a beam polarisation of 90% or greater prior to injection into the storage-ring [2].

At the end of this production sequence, the μ^+ bunches are polarised, have a narrow momentum spread and a time distribution of less than one cyclotron period in the storage-ring.

3.3.2 Muon storage in the ring

Following their production, the polarised μ^+ are injected into the storage-ring. The μ^+ enter through a channel in the magnet that is near tangential to the ideal orbit at the point of entry. Just before entering the storage region they enter an inflector magnet which acts to cancel the magnetic field produced by the storage-ring magnet and allow the μ^+ to pass straight through relatively unperturbed into the storage region. Were there no inflector, then the field at the edge of the storage region would cause the μ^+ to bend inwards as they enter and the resulting trajectory would quickly result in the μ^+ exiting the storage region. A diagram showing the locations of instrumentation around the storage-ring can be found in figure 3.6.

Once injected, the μ^+ begin to orbit within the ring. As the inflector aperture is at the edge of the storage region, the orbit that the μ^+ initially follow is radially offset relative to the ideal case. A kicker magnet, located one quarter of a turn around the ring, pulses during this first orbit to provide a force to move the μ^+ on to the ideal orbit, i.e. central within the storage region.

The average orbit radius of the μ^+ population is now close to the ideal radius. However, the μ^+ have a small momentum spread which results in a spread of orbit radii. Collimators truncate the stored beam to remove μ^+ with momenta near the extremes of the distribution, resulting in a final momentum spread of 0.15 % in the stored beam [2].

The stored beam diverges vertically as it circulates, causing particles to move away from their stable orbits and eventually to exit the storage region. Focusing is required to mitigate this issue. Four electrostatic quadrupoles spaced evenly and symmetrically around the ring provide discrete regions of vertical focusing for the stored μ^+ as they orbit. Electrostatic rather than magnetic quadrupoles are chosen to minimise




Figure 3.6: Diagram of the locations of beam-control instrumentation and detectors around the storage-ring. The ideal orbit path and direction is shown in green.

perturbations to the dipole magnetic field. This quadrupole field acts to defocus the stored beam in the radial direction, but the combined effect of the radial electric field and the vertical storage-ring magnetic field provides radial focusing, with the storage-ring acting as a weak focusing betatron [2]. The stored μ^+ undergo periodic oscillations about their average orbit path, known as betatron oscillations, in both the radial and vertical directions due to this focusing. The frequency of the betatron oscillations is an order of magnitude larger than ω_a , and hence these effects do not directly bias the g-2 measurement. A given detector however observes the beam once per orbit from a fixed location. This spectrometer-like discrete sampling of a periodic function leads to an apparent net periodic motion in the beam with a frequency more comparable to ω_a , an effect known as coherent betatron oscillation (CBO). Care must be taken in setting the operating parameters of the focusing quadrupoles to avoid CBO periods or harmonics close to ω_a .

3.3.3 Full treatment of precession in the storage ring

The presence of an electric field from the vertical focusing quadrupoles modifies both the relativistic ω_c and ω_s expressions previously introduced in equations 3.6 and 3.5. Additionally, accounting for the case where the stored μ^+ momentum vectors are not orthogonal to the magnetic field vector results in an additional modification to ω_s . The resulting ω_a expression given these effects is [42]:

$$\omega_a = \frac{e}{m} \left[a_\mu \mathbf{B} - \left(a_\mu - \frac{1}{\gamma^2 - 1} \right) \frac{\boldsymbol{\beta} \times \mathbf{E}}{c} - a_\mu \left(\frac{\gamma}{\gamma + 1} \right) (\boldsymbol{\beta} \cdot \mathbf{B}) \boldsymbol{\beta} \right], \quad (3.8)$$

where β is the particle velocity as a fraction of the speed of light in vacuum, c .

Equation 3.8 is significantly more complicated than the expression presented in equation 3.4, and includes the additional free parameters $\boldsymbol{\beta}$, γ and \mathbf{E} . However, two simplifying assumptions can be made to reduce the ω_a expression to the form in equation 3.4. This reduces the number of parameters that need to be measured in the experiment and simplifies the expression that is ultimately fitted when measuring a_μ , but results in two small correction factors being required to reach the correct a_μ value.

The first simplification to equation 3.8 is achieved when the μ^+ momentum direction is orthogonal to the magnetic field, e.g. when $\boldsymbol{\beta} \cdot \mathbf{B} = 0$. This is true on average since the μ^+ are orbiting about the magnetic field. However, the presence of vertical betatron oscillations results in a periodic vertical component in the μ^+ motion as their momentum vectors “pitch” up and down. This requires a correction to be applied to ω_a when calculated using equation 3.4, which is commonly known as the pitch correction.

The second simplification to equation 3.8 that can be made is to remove the term that depends on \mathbf{E} . Fortunately, there exists a choice of γ that reduces this term in equation 3.8 to zero: this is known as the “magic” γ and is defined as:

$$\gamma_{\text{magic}} = \sqrt{1 + \frac{1}{a_\mu}}. \quad (3.9)$$

This particular γ_{magic} value corresponds to a μ^+ momentum of 3.094 GeV, and is known as the “magic” momentum. This is the target momentum for the μ^+ injected into the storage-ring, and the design orbit radius in the ring is for μ^+ of this momentum. The stored μ^+ will in reality have a momentum spread about this magic value, and a correction to ω_a will be applied to compensate for this effect: this correction is commonly referred to as the E-field correction. Both the pitch and E-field corrections can be computed either analytically or from simulations.

3.3.4 μ^+ decay in the storage-ring

At the magic momentum, μ^+ have a time dilated lifetime of 64 μs , a substantial increase from their lifetime at rest of 2.2 μs [19]. In the storage-ring magnetic field, the period corresponding to the anomaly frequency is 4.4 μs , meaning that most μ^+ will undergo many precessions before decaying.

As outlined in section 3.2, the resulting decay e^+ are Lorentz boosted, and are emitted within 30° of the μ^+ momentum direction, or less than 2° for the high energy e^+ counted by the calorimeters. The decay e^+ are thus initially moving approximately tangentially to the μ^+ orbit. The e^+ carry a fraction of the μ^+ energy, the rest being carried by the two neutrinos also produced in the decay, and thus the e^+ follow

more highly curved trajectories in the magnetic field than the μ^+ orbit path. The e^+ trajectories take them radially inwards, exiting the storage region on the inside of the ring where they can be measured by the experiment's detector systems. Figure 3.7 shows an example decay e^+ trajectory from the initial decay vertex at the μ^+ orbit radius until it is stopped in a calorimeter.




Figure 3.7: Example trajectory of a decay e^+ following the decay of an orbiting μ^+ . The e^+ is stopped and detected in a calorimeter, three of which are pictured. Image courtesy of J. Kaspar.

Data is accumulated for 700 μs : after this time virtually all the μ^+ will have decayed.

3.3.5 Detector systems

The inner edge of the storage-ring is instrumented with detector systems to measure the e^+ resulting from the μ^+ decays. 24 calorimeters are equally spaced around the ring with the purpose of measuring the arrival time and energy of the incident e^+ , and are used to count the numbers of high energy e^+ versus time during fills and measure ω_a .

The calorimeters are each composed of 54 PbF_2 crystals, arranged in grid of 9 crystals horizontally and 6 vertically, to form a segmented detector that is 225 mm wide, 150 mm high and 140 mm deep [49]. Relativistic charged particles traversing

a PbF_2 crystal produce Cerenkov light. The decay e^+ produce particle showers in the dense crystals, resulting in detectable Cerenkov light which travels along the crystal and is detected at the downstream end by a silicon photo-multiplier (SiPM). The crystals are clad in reflective wrapping to minimise light transmission between crystals. The showering e^+ lose energy rapidly in the crystals until all of their kinetic energy is lost, and hence the measured light yield can be used as a measure of the e^+ energy. A Cerenkov medium was selected due to the fast and short signal it provides, which minimises the pile-up of signals. The detector segmentation similarly serves to minimise pile-up in individual crystals. The electronics were selected to maximise the gain stability over time, and for the ability to operate in strong magnetic fields [2]. A photo of a calorimeter is shown in figure 3.8.




Figure 3.8: Photograph of a calorimeter during construction. The perspective is from the top and back of the detector. Individual PbF_2 crystals can be seen connected to readout SiPMs. Image courtesy of J. Kaspar.

A laser calibration system is employed to calibrate the gains of individual calorimeter crystals and SiPMs, and to monitor the gain variation over time [50]. Laser pulses are delivered into the face of every calorimeter crystal periodically between and during the μ^+ fills. This allows both fast changes in gain during the readout process and long term changes over hours and days to be corrected for.

Although the calorimeters provide the key observables in the ω_a measurement, a number of support detectors are also present. These are primarily used for particle

tracking to observe the motions of the stored μ^+ during the fill and measure the beam profile. The anomaly frequency, ω_a , for a given μ^+ depends on the magnetic field that the μ^+ traverses during its lifetime in the storage region, and this effective-field can be obtained by convoluting the measured beam profile with the magnetic-field measured as a function of position in the storage region. These stored beam properties are measured as a function of time and at multiple azimuthal positions around the ring in order to detect time dependent variations. These could be periodic, for example the aforementioned betatron oscillations, or other time dependent effects within the fill. Examples include the movement radially outwards of the mean stored beam position as the fill progresses due to the μ^+ at larger radii being those with larger momentum, which thus have longer lifetimes due to time dilation. μ^+ at small radii thus decay more quickly on average, shifting the mean radial position of the beam outwards during the fill.

Two distinct detector systems are used for measuring the stored muon profile. The first are three straw tracking detectors spaced azimuthally around the ring, which are located on the inside of the ring immediately upstream of 3 of the 24 calorimeters and within the storage-ring vacuum. Decay e^+ reach the straw trackers in the same manner as the calorimeters, leaving hits in individual straws. Tracks can then be fitted to these hits which are then extrapolated back to the point of decay, giving a measurement of the stored beam profile. More details on the straw trackers are given in section 4.

The second detector system for measuring the stored beam profile are scintillating so-called fiber-harp detectors. Each holds seven parallel scintillating fibers between two arms reminiscent of the strings of a harp, as can be seen in figure 3.9. The harps can be moved into the path of the stored beam, and are arranged in pairs orientated orthogonally to one another in order to measure both the radial and vertical beam profile. Their location in the stored beam itself has the advantage that the μ^+ profile can be directly measured, rather than extrapolating from the e^+ profile as for the straw trackers. The disadvantage is that since the μ^+ on average traverse many orbits and pass through the harps on each orbit, the integrated harp material the μ^+ pass through is high and the resulting scattering is large. The fiber-harps will thus not be deployed during nominal operations, but instead will be used in dedicated beam-profile measurement-runs. The straw trackers by contrast can run at all times.




Figure 3.9: Photograph of a fiber harp plane. The harp is in the deployed position, with vertical scintillating fibers used to measure the horizontal beam profile.

The final particle tracking detector system are the injection beam monitors. These detectors are also scintillating fiber detectors and are mounted up and downstream of the inflector. The beam profile delivered by the upstream accelerator infrastructure can thus be measured just prior to the g-2 storage-ring and immediately after the inflector. This information will be used to tune the beam to best match the storage-ring acceptance, and so maximise the number of stored μ^+ .

As well as spatial tracking of particles, scintillator counters are placed at the opening into the storage-ring in order to measure the injection time of a fill, allowing data from multiple fills to be phase aligned and so avoid smearing the ω_a modulation of the e^+ count signal. This is particularly important because the Fermilab accelerator runs on a different clock to the g-2 experiment.

All detectors and other devices that are in or near the magnetic field region have been constructed minimising the number of magnetic components to avoid perturbing the storage-ring magnetic field.


3.3.6 Storage ring magnetic field

As can be seen in equation 3.4, a determination of a_μ requires a precise knowledge of the storage-ring magnetic field in addition to the ω_a measurement. This is achieved by ensuring the magnetic field in the storage region is highly uniform and stable versus time, and by precisely mapping the field. The required measurement precision of the field-region sampled by the beam is 70 ppb when averaged over time [2], which can be compared to 170 ppb at BNL. This improvement will be achieved by improvements in both the field uniformity and the measurement precision.

The storage-ring is an iron dipole magnet with superconducting coils above and below the storage region providing the magnetic field [51], as shown in figure 3.10. Iron pole pieces are located between the superconducting coils to produce a uniform dipole field. An iron yoke forms a C-shape around the top, bottom and outer side of the ring and provides the flux return. This design leaves the inner side clear for the decay e^+ and the detectors. A cryogenic system is used to cool the coils to superconducting temperatures. The magnet design is driven by the requirement of high uniformity both azimuthally and radially within the storage region, to minimise discontinuities at the intersections of the separate pieces, and to minimise higher-order multipole fields.

The ring was originally designed and constructed for the E821 BNL muon g-2 experiment, and is reused for the new Fermilab measurement. The superconducting coils were shipped in one piece from BNL to Fermilab in the summer of 2013. The large iron components such as the pole and yoke pieces were removed and shipped separately before reassembly in Fermilab.

To achieve the required field uniformity, many small-scale degrees of freedom are built into the magnet to allow non-uniformities to be corrected for in a process known as “shimming”. This flexibility is provided by move-able or insert-able metal pieces that can be adjusted to make small local changes to the field. Although the overall magnet is the same as the BNL experiment, the field uniformity has been substantially improved by an extensive shimming process. The shimming components range in scale from tens of cm iron blocks to sheets of foil with sub-mm details. Examples include large “top-hat” blocks above and below the yoke, small edge-shims used at the edges of the pole pieces, and wedge-shaped shims between the poles and yoke that are used to compensate for the quadrupole effects caused by the C-shape of



g-2 Magnet in Cross Section

Figure 3.10: The cross-section of the g-2 storage ring. Indicated are the coils, pole pieces and yoke of the superconducting magnet, the μ^+ storage region, the location of fixed NMR probes, and various shimming devices: top-hats, edge-shims, surface-coils, and wedges.

the yoke. As well as these passive shimming techniques, wire coils on the surface of the poles can be used to apply variable corrections to the field, which is particularly useful for correcting slow temporal changes.

The goal for the Fermilab muon g-2 experiment is to achieve azimuthally-averaged magnetic field uniformity of 1 ppm, with local variability limited to 50 ppm [2]. The first shimming campaign was completed in September 2016, with the resulting magnetic field uniformity versus azimuth shown in figure 3.11, and the azimuthally-averaged field shown in figure 3.12. Figure 3.12 shows that the azimuthal variability achieved is within 25 ppm, a factor of 2 better than the requirement and a factor of 4 improvement relative to the BNL muon g-2 experiment. Further shimming will be performed following the installation of ring instrumentation.

The technique employed to measure the magnetic field is pulsed Nuclear Magnetic Resonance (NMR), and is a continuation and enhancement of the BNL E821 methodology and hardware. The principle of the measurement is as follows. A sample of material containing protons (typically petroleum jelly) whose spin directions have no




Figure 3.11: The storage-ring magnetic field relative to the overall average field strength, plotted versus azimuth, at BNL (blue) and Fermilab (red). The Fermilab measurement was taken following the first shimming phase, and will be improved further following the installation of ring instrumentation.

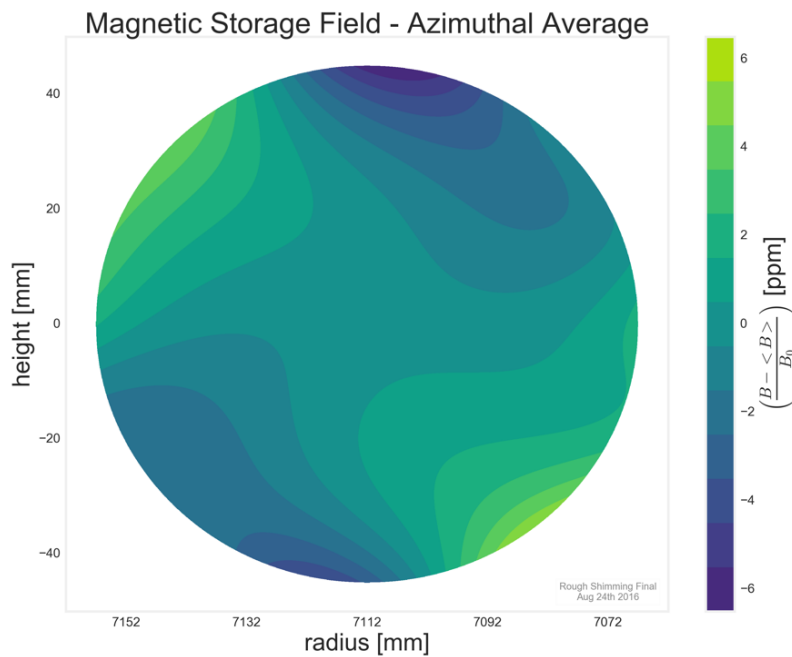


Figure 3.12: The azimuthally-averaged storage-ring magnetic field relative to the overall average field strength at Fermilab following the first shimming phase, plotted versus vertical and radial position within the μ^+ storage region. The field uniformity will be further improved during future shimming before physics runs begin.

net orientation is placed in the magnetic field. The proton spins have a tendency to align or anti-align with the field: preferring the aligned state, they produce a net magnetisation in the material. As protons have a magnetic moment, they undergo Larmor precession in the same manner as was discussed for muons in section 3.2, meaning their spin vectors precess about the field. A Radio-Frequency (RF) pulse can be used to align the spin vectors in the transverse direction to the magnetic field, producing a net transverse magnetisation component [52]. Following this pulse, the spin vectors will then precess in an initially coherent way, producing a sinusoidal rotation of the transverse component of the net magnetisation. However, over time decoherence in the precession phase takes place, and the net transverse magnetisation will decay exponentially in a process known as Free Induction Decay (FID). The precessing net magnetisation during this decay period induces a voltage in a pickup coil surrounding the sample which is digitised and Fourier transformed to measure the Larmor frequency. The magnetic field is then related to the measured proton Larmor frequency, ω_p , using the equivalent expression to equation 3.1 for protons [2]:

$$\omega_p = 2\mu_p \mathbf{B}, \quad (3.10)$$

where μ_p is the proton magnetic moment.

The pulsed NMR ω_p measurements are made using refurbished custom NMR magnetometers, often called probes, inherited from the BNL muon g-2 experiment [53]. The probes have a small cylinder containing petroleum jelly, which acts as the sample of protons whose spin precession is observed. The sample is surrounded by a pickup coil in which the voltage from the spin precession is induced, and the probe includes electronics for delivering the RF pulse and digitising the signal.

The probes are used in two separate ways in the experiment. A few hundred probes are located in fixed positions azimuthally above and below the ring as close to the storage region as possible. These probes make regular measurements of the fringe field during muon storage. Additionally, 17 probes are mounted in a trolley system which is pulled around the ring inside the storage region to directly measure the field experienced by the muons. The trolley data is used to produce a magnetic field map of the storage region versus azimuth, radius and height. These measurements take

place in designated trolley runs where no muons are present as the trolley would interfere with the beam. These runs will be performed every 2 hours: much more frequently than was the case at BNL. The fixed probe data of the fringe field is used to interpolate the field map between trolley runs.

There are two additional types of pulsed NMR probe in the experiment that are used for calibration. The protons in the petroleum jelly in the trolley probes are not truly free protons, as they are within the jelly and also within the rest of the material of the probe and the trolley itself. A correction must therefore be made to obtain the free proton Larmor precession from the measured Larmor precession frequency in each probe, which is required for the determination of the magnetic field. A NMR probe built for both the BNL experiment and the Los Alamos National Laboratory (LANL) E1054 experiment is used to provide an absolute reference measurement of a magnetic field relative to the free proton Larmor frequency. This probe has undergone extensive testing and calibration, and features a highly spherical pure water sample which acts to minimise diamagnetic shielding within the sample itself, which is the dominant factor in the deviation of the measured Larmor precession versus the free proton precession frequency [53]. Measurements of the storage-ring field with this absolute calibration probe can be compared to trolley measurements at the same position in order to accurately determine the Larmor precession correction factor for the trolley probes. This probe will also be cross-calibrated with a probe being used in a J-PARC muonium Hyper-Fine Splitting (HFS) experiment that is very precisely measuring the ratio of the muon and proton magnetic moments [12].

Access to the storage region is required to insert the relatively large absolute calibration probe, meaning that this measurement cannot be performed once the storage region is sealed and at vacuum. Smaller probes known as plunging probes can be inserted to cross-check the trolley measurements even during vacuum. These plunging probes are calibrated against the absolute calibration probe, and are thus used to effectively transfer the calibration between the trolley and the absolute calibration probe.

In order to cross-check the existing absolute calibration probe, two additional absolute calibration probes will be developed for the Fermilab muon g-2 experiment. The first features a similar concept and design to the E821 probe, but with a number of improvements. The second utilises ^3He instead of petroleum jelly and an optical pumping system to achieve net spin polarisation in the sample [2]. This second

system should exceed the the precision of the spherical pure water devices, and will be utilised once the efficiency of the optical pumping is demonstrated.

3.3.7 Combining ω_a and ω_p

The two major measurements made in the Fermilab muon g-2 experiment are of ω_a by the calorimeters, characterising the μ^+ spin precession, and ω_p by the various NMR probes, characterising the storage-ring magnetic field. Combining equations 3.4 and 3.10 gives an expression for a_μ with respect to these measured quantities:

$$a_\mu = \frac{2m_\mu\mu_p}{e} \frac{\omega_a}{\omega_p}. \quad (3.11)$$

The precision to which a_μ can be determined using equation 3.11 depends not only on measurement precision of ω_a and ω_p , but also on the precision to which the physical constants e , m_μ and μ_p have been measured by other experiments. It is possible to re-express equation 3.11 in terms of the ratios of a number of physical constants in such a way as to reduce the total uncertainty in a_μ compared to equation 3.11, exploiting the high precision to which these ratios have been measured.

The anomalous magnetic moment of the electron, μ_e , is given by:

$$\mu_e = \frac{g_e e}{4m_e}, \quad (3.12)$$

where m_e is the mass of the electron.

Substituting equation 3.12 into 3.11 yields the new expression for a_μ :

$$a_\mu = \frac{g_e m_\mu \mu_p}{2 m_e \mu_e} \frac{\omega_a}{\omega_p}, \quad (3.13)$$

The value of $\frac{g_e}{2}$, e.g. the anomalous magnetic moment of the electron, has been measured with 0.28 ppt precision as detailed in section 2.1 [11]. The ratios $\frac{\mu_p}{\mu_e}$ and $\frac{m_\mu}{m_e}$ are determined from the LAMPF muonium HFS experiment [54], and are known to 3 ppb and 22 ppb precision respectively [55]. Additionally, the MuSEUM experiment, due to start taking data within 2 years, will measure muonium HFS with a factor of 10 improvement in precision compared to LAMPF [12].

3.3.8 Measuring an EDM

The case for interest in a non-zero muon EDM was outlined in section 2.3. A non-zero muon EDM would result in the μ^+ in the $(g - 2)$ storage-ring experiencing a torque in the presence of an electric field, which exists for the orbiting μ^+ due to both the apparent electric field in the μ^+ rest-frame resulting from the Lorentz boost to the laboratory-frame magnetic field, and from the electrostatic quadrupole fields [18]. This torque results in the precession of the μ^+ spin vector as is the case for the magnetic dipole moment. This effect is thus in principle measurable at the Fermilab muon g-2 experiment.

The angular frequency of the precession resulting from a muon EDM, ω_η , is given by [18]:

$$\omega_\eta = \eta \frac{e}{2m} \left[\frac{\mathbf{E}}{c} + \boldsymbol{\beta} \times \mathbf{B} \right]. \quad (3.14)$$

The $\boldsymbol{\beta} \times \mathbf{B}$ term, which represents the contribution from the induced electric field from the Lorentz-boosted magnetic field, dominates over the \mathbf{E} term, which represents the quadrupole electric field contribution. The resulting ω_η vector is orientated radially inwards towards the centre of the storage-ring, orthogonal to the vertically orientated ω_a [5]. The spin precession plane for a non-zero EDM is thus tilted out from the orbit plane towards the ring centre, as shown in figure 3.13.

Since the decay e^+ are emitted preferentially along the μ^+ spin axis, an observable consequence of this tilt is that the e^+ will be on average downward going when the spin vectors are pointing towards the centre of the ring, and upward going when the spin points away from the centre. The average e^+ vertical angle away from the orbit




Figure 3.13: Angular momentum vectors of the μ^+ spin precession resulting from magnetic and electric dipole moments, and the resulting tilt of the overall spin precession plane [2].

plane will thus vary with ω_η . It is this signal that is used to search for an EDM, which can only be measured accurately at the Fermilab muon g-2 experiment using the trackers.

The addition of an EDM spin precession component changes the overall magnitude of the spin precession, and in the event of a non-zero EDM being measured, would require the measured precession frequency to be corrected for the ω_η contribution. The vertical angle oscillation amplitude and the small correction to ω_a increases linearly with η .

The ω_a periodic signal peaks when the μ^+ spin and momentum are aligned, whilst the vertical angle oscillation peaks when the spin is pointing radially outwards. The two signals are therefore out of phase by $\pi/2$. This can be used to distinguish an EDM signal from other effects that can cause an apparent vertical angle oscillation, such as detector misalignments resulting in a different detector acceptance for upwards versus downwards going e^+ .

Chapter 4

The straw trackers

4.1 Overview

The straw tracking detectors, or trackers, for the Fermilab muon g-2 experiment form the major subject of this thesis. Their primary purpose is to reconstruct the trajectories of e^+ from μ^+ decays and by extrapolation infer the position of the μ^+ at the time of decay. This can be then be used to determine the properties of the stored μ^+ beam and the implications for the a_μ measurement.

The experiment will use three straw trackers mounted on the inside of the storage-ring just outside of the storage region. If the azimuth, θ , of the injection point in the ring is defined to be $\theta = 0^\circ$, the three trackers are located at $\theta = 15^\circ$, 180° and 270° . These locations are chosen to spread the trackers around the ring and so measure the stored beam properties in different locations, whilst avoiding conflicts with other pieces of equipment mounted along the inside of the ring such as the NMR-trolley garage and the kickers. Having trackers at multiple locations increases the number of e^+ that are tracked and provides redundancy.

Each of the three trackers are located immediately upstream of a calorimeter, allowing the e^+ hitting these three calorimeters to be tracked and direct comparisons between the two detector systems to be made, such as measuring $\frac{E}{p}$. The trackers and calorimeters have a similar acceptance meaning the two detector systems are observing e^+ with comparable initial momenta and vertex positions.

4.2 Tracker goals

4.2.1 Observation of stored beam

The trackers can contribute in a number of ways to the Fermilab muon g-2 experiment. Their primary goal is to measure the spatial profile of the stored μ^+ beam and measure its properties over time. The position of the stored μ^+ is inferred using the tracker by forming tracks for the e^+ that pass through the tracker and extrapolating back to the storage region to determine the μ^+ decay vertices. Each vertex gives a measurement of the position of a μ^+ during its orbit, and hence the extrapolated tracks of many e^+ can be used to form a map of the μ^+ beam. The precession of the μ^+ spin depends on the integrated magnetic field it experiences during its life in the ring. Although the storage-ring magnetic field is extremely uniform, it is not perfect, and convoluting the measured beam profile with the measured magnetic field map is used to determine the integrated, average magnetic field experienced by the μ^+ as they orbit.

Dynamical properties of the stored beam can also be observed by considering the time-dependence of the stored muon profile. Time-dependent effects such as betatron oscillations in the stored beam affect the calorimeter acceptance, which directly affects the number of high energy e^+ observed in the calorimeter, and hence the determination of ω_a . Drifts in beam parameters over longer time scales can also be observed. The profile of the stored beam directly affects the E-field and pitch-corrections made to the ω_a measurement, as discussed in section ??, and tracker measurements can also be used to quantify these effects.

The properties of the beam measured by the tracker will also be used to optimise the simulation, which in turn will be used to quantify some of the systematic uncertainties in the measurement.

4.2.2 Supporting the calorimeter

As well as extrapolating e^+ tracks backwards to determine μ^+ decay vertices, it is also possible to extrapolate the tracks forwards to the downstream calorimeters. This then predicts where calorimeters are hit, which can be compared to the measured calorimeter hit positions. This serves as a cross-check between the two systems, and

can be used to assess the alignment of the detectors. Matching tracks and calorimeter hits in this way also allows measurements of the e^+ energy in the calorimeter to be compared to momentum measurements inferred from the tracker curvature, and can thus provide an independent monitor of the calorimeter gain calibration.

A significant contribution to the systematic uncertainty in the ω_a measurement is particle pile-up in the calorimeter, where multiple particles hit the same area of the calorimeter within a small time window of each other. The multiple Cerenkov light cones in the calorimeter and resulting electronics signals for the different particles are difficult to differentiate, leading multiple particles to be reconstructed as a single particle with an energy that is the sum of the pile-up particle energies. This will result in false counts of e^+ seemingly above the calorimeter energy cut, and would thus affect the ω_a measurement. This was a significant source of error in the E821 experiment, and has driven the design of the Fermilab muon g-2 experiment calorimeter to use a fast signal and segmented crystals to improve the time and spatial resolutions respectively. Despite this, there remain pile-up events that the calorimeter cannot resolve itself. The trackers can identify these pile-up events by observing the distinct particle tracks causing the single pile-up hit in the calorimeter. The tracker can thus be used to correct for pile-up in the three calorimeters directly downstream of the trackers, and estimate the systematic uncertainty resulting from this effect across all 24 calorimeters.

4.2.3 Observation of vertical angle oscillation

The nominal μ^+ orbit is in a horizontal plane perpendicular to the B-field: vertical deviations from this can be determined from the e^+ trajectories measured in the trackers, and the average vertical angle of the μ^+ spin vectors at a given time can then be inferred. An oscillation of this vertical angle is the signature of a non-zero muon EDM (see section 3.3.8). The tracker is the only detector system in the Fermilab muon g-2 experiment capable of precisely measuring this angle.

4.3 Tracker high-level design

The goals for the trackers outlined in section 4.2 drive the design of the detector. Fundamentally, the design objective is to maximise the number of e^+ tracked and the precision with which the track properties are determined, whilst minimising any detrimental effects on other systems. Tracking a particle requires measurements of the particle's three-dimensional position at a number of points along its trajectory. The detector is designed to make multiple position measurements per particle at a high rate. The position measurements must also be sufficiently spatially separated such that the curvature of the track in the fringe magnetic field of the storage-ring can be measured to allow accurate track extrapolation and momentum reconstruction.

The μ^+ can decay at any point within the storage region, requiring that the tracker have a large azimuthal acceptance to maximise the number of e^+ that are tracked. Additionally, the e^+ momentum range from μ^+ decay is large, meaning that the tracker must be capable of tracking particles with a wide range of track curvatures. This includes e^+ that exit the detector before traversing its full length. These two requirements can be achieved using a long and wide detector, with numerous points along the detector length where track points can be measured. In particular, the highest energy e^+ have radii of curvature close to the μ^+ orbit radius, requiring a detector length of order 1 m in order to measure the track curvature with a reasonable precision.

The tracker length is thus maximised within the region available between two calorimeters. The width covers the region radially inwards from the μ^+ storage region up until the "shadow" of the upstream calorimeter, beyond which the e^+ flux drops rapidly. The e^+ receive a large Lorentz boost in the μ^+ momentum direction tangential to the orbit, and hence are largely directed in the orbital plane. This means that the vertical coverage of the tracker need not be particularly large, and is commensurate with the storage region vertical aperture.

The calorimeter and tracker are orientated to be perpendicular to the highest energy e^+ , as it is these that are counted in the calorimeters to determine ω_a (see section 3.2). As the e^+ path curvature reduces with increasing energy, these high energy e^+ leave the storage region at shallow angles relative to the orbit circle, and so the detectors must be located as close to the storage region as possible.

The envelope available for a tracker matching these design requirements is shown in figure 4.1.

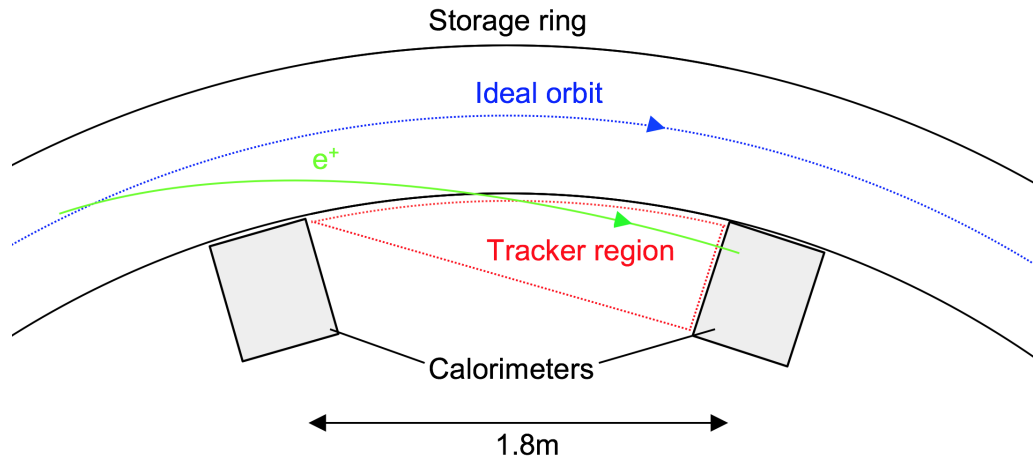


Figure 4.1: A view from above of the trajectory of a decay e^+ and the region available for the trackers between two calorimeters.

Using a large detector potentially introduces large amounts of material. Particles can scatter in the detector material, causing deviations to their trajectory which introduce systematic uncertainty in track fitting and hence the stored beam profile measurement: minimising the material in the detector is thus crucial. This leads to the decision to use a gas-based detector system, which is a lower density and cheaper active medium than alternative technologies such as solid state detectors, and hence the detector can be made large whilst keeping the material and cost budget relatively low. Additionally, the active detection area is divided into 8 evenly spaced tracking planes, meaning that position measurements can be made along the length of the detector without filling the entire volume with material. The wide range of μ^+ decay vertices and e^+ track curvatures mean that not all e^+ will traverse all planes. 8 planes is chosen as a trade-off between maximising the number of particles traversing enough planes for track fitting and minimising material.

To further reduce the material the e^+ must traverse, the entire tracker is placed in vacuum, meaning that between the tracker planes there is minimal material. The storage-ring vacuum chambers are extended in the regions where the three trackers are located to facilitate this. The tracking planes are subdivided into distinct cylindrical “straws”, whose circular geometry can resist the pressure difference between the gas inside and the vacuum outside with the minimum wall thickness, again reducing the

material in the detector. Each straw contains a single central wire along its long axis which is held at high voltage, and the induced signal on this wire is used to detect the passage of charged particles.

4.4 Detection principle

The straws in the tracker planes act as individual drift tube detectors. They detect the electrons and ions produced by ionisation when a charged particle such as a e^+ traverses the gas in the straw. The charged particles encounter gas molecules at random during their passage through the gas, and sufficient energy can be transferred in these collisions to liberate electrons from the gas. For gases commonly used in drift detectors, $\mathcal{O}(10)$ electrons are typically liberated in this manner per cm traversed by the charged particle [56]. These liberated electrons and the ions that result from the ionised gas molecules are mobile charge-carriers within the gas. Multiple electrons may be liberated in a single collision, with the resulting free electrons referred to as a cluster. The electrons liberated by ionisation from the traversing charged particle are known as primary electrons. Secondary electrons may also be liberated by the subsequent ionisation of the gas by the primary electrons.

Each straw has a central wire held at high voltage acting as an anode and grounded straw walls acting as a cathode, creating an electric field radially out from the wire. This electric field causes the liberated charge carriers to experience a radial force, towards the wire in the case of the electrons and away from it (towards to the straw walls) for the ions. The liberated charges undergo a net motion under this force, but also experience collisions with the gas molecules as they travel, slowing their progress. This overall motion is known as drift.

As well as the radial electric force, the vertical storage-ring magnetic field introduces an orthogonal force on the drifting charges. This curves the trajectory of the particles, in particular the light electrons [56].

As these liberated charges drift, they induce a current in the wire. This is the signal that is used to detect the passage of a charged particle, with the wire being referred to as a sense-wire. The induced current for individual electrons is small. However, if the wire is thin then the electric field close to the wire is strong enough that

electrons can accelerate enough in the time between collisions to ionise the next gas molecule they encounter and liberate further electrons [56]. These newly liberated electrons are also accelerated enough to further ionise the gas, creating a so-called avalanche effect where the number of electrons multiplies rapidly. This amplification in the number of moving charges creates a much larger induced signal, resulting in a signal detectable above noise levels from only a few primary ionisation electrons. Detectors operating in this manner are known as a proportional wire chambers, as the post-avalanche signal is approximately proportional to the number of primary ionisations [56].

In the avalanche, each new ionisation produces an ion as well as an electron. These ions drift back out to the straw walls and hence also induce a signal in the wire. As these ions travel a large distance (approximately one straw radius), they induce a much larger signal than for electrons. Due to their large mass however they travel much more slowly than the electrons and hence this signal is delayed, of order μs relative to the electron-induced signal, and is often referred to as the ion tail. The faster electron signal forms the rising edge of the pulse, and is used to trigger the electronics rather than the slower, but larger, ion signal, which is instead deliberately suppressed by the electronics to avoid interference with signals from subsequent charged particles, allowing the detector to cope with a high rate of particles.

An example ionisation event in a straw modelled using the GARFIELD [57] gaseous detector simulation software package is shown in figure 4.2.

The measured quantity in the straws is the time that the fast electron signal, following some shaping, crosses a threshold in the readout electronics and triggers a hit. The time of this hit is recorded. Registering a straw-hit indicates that a particle has traversed that straw, but it is desirable to know the particle's trajectory to a greater precision than this. The measured hit time, t_h , includes both the time that the charged particle traversed the straw, known as t_0 , and the time it took the primary ionisation electrons to drift to the wire, cause an avalanche and trigger the electronics, known as the drift time t_d . If t_0 is known, the drift time can be extracted from the hit time. The drift time can then be converted to a distance if the behaviour of the charge in the gas is well understood, yielding the Distance of Closest Approach (DCA) of the charged particle to the sense-wire, as can be seen in figure 4.3. This distance measurement is a magnitude with no directional information, and thus specifies a cylinder around the wire that the traversing particle intersected at a single point.




Figure 4.2: Top-down view of a charged particle traversing a Fermilab g-2 straw in a GARFIELD simulation. The charged particle track is shown in red, with black circles representing the primary ionisations. Blue lines represent the drift trajectories of the primary electrons until they terminate on the central sense-wire. Green lines represent the drift of the avalanche ions. The electron paths are visibly curved due to the external magnetic field.

Combining this drift cylinder with cylinders from straws at different orientations traversed by the same charged particle allows a full 3D particle trajectory to be formed. Reconstructing track positions by combining the drift cylinders in multiple straws is further discussed in section 6.5.3.

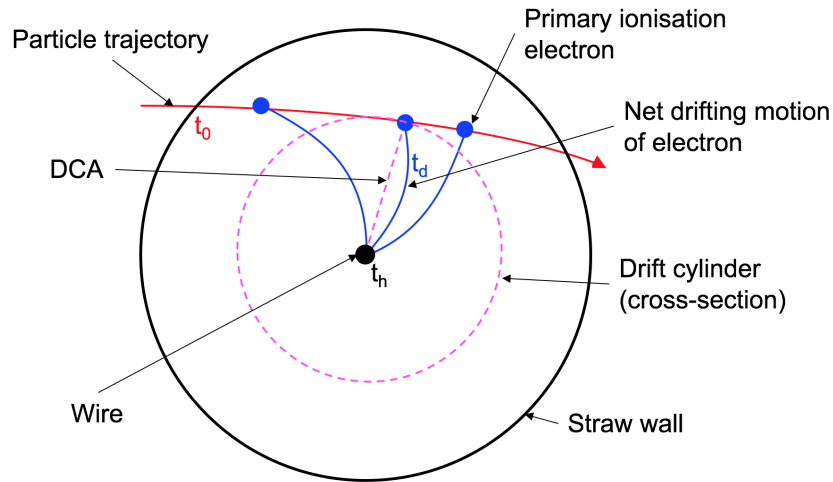


Figure 4.3: Diagram of the relationship between the drifting of primary ionisation electrons to the straw wire and the DCA reconstructed from the time taken by this drift. The drift cylinder specified by the DCA is also shown.

4.5 Tracker detailed design

4.5.1 Straw design

The design of the straws is driven by the desire to minimise the material whilst achieving good resolution, efficiency and reliability, and being able to operate within a vacuum. Each straw is 100 mm in length with a radius of 2.5 mm. The walls are 15 μm thick and constructed from spiral-wound aluminised mylar, which is strong enough to resist the pressure gradient between the inside and outside of the straws even at this low thickness. Thin aluminium and gold layers on the inside of the straw provide electrical conductivity to allow the straw walls to act as a cathode, and an outer layer of aluminium reduces gas permeation through the walls, reducing the leak rate of the gas into the vacuum. The straws are fixed in place at either end and tensioned against the pressure gradient.

The central sense-wire must be extremely thin to facilitate a large avalanche amplification. However, the wires must also be very strong so that they can be held at high tension to ensure they are extremely straight, as otherwise their positions will not be well known which will introduce errors in the track finding. Without this high tension, the wires can move under the force of the electric field or gravity. 25 μm tungsten wires are chosen to meet these requirements.

The use of a straw-based design has the advantage that if a particular wire snaps, it is contained within that straw and only a single straw is disabled. Detectors which have many wires within a single chamber can suffer from snapped wires becoming wrapped around other wires and thus disabling large parts of the detector.

End plugs are designed to cap the straws at each end and hold the wire in place. These plugs contain channels for gas to flow through, with the gas entering the straws at one end and leaving through the other, creating a slow flow through the straw.

4.5.2 Choice of gas

There are many factors to consider when choosing the gas and operating conditions for an ionisation detector. Of particular importance are the yields of primary ionisation electrons, the stability and amplification of the avalanche, the charge drift velocity, and the ageing of the detector, e.g. the gain deterioration over time.

Noble gases are common choices for ionising gas detectors as they are formed of atoms rather than molecules. Molecular gases have additional modes in which they may be excited compared to noble gases, such as rotational or vibrational energy states, which do not result in mobile charge carriers. In noble gases, ionisation is the dominant form of energy loss for particles in energy ranges typically used in high-energy physics, and hence the signal measured for a traversing particle is maximised. The energy required for ionisation decreases with increasing atomic number for noble gases as the outer shell electrons are increasingly shielded from the nucleus, while the multiple scattering increases with atomic number such that argon is a common choice [56].

Within the avalanche, photons are produced as well as electrons and ions via the excited states of gas atoms and molecules. These Ultraviolet (UV) photons may themselves have enough energy to liberate more electrons. The photons have a longer

mean free path than electrons however, and can leave the avalanche region and cause new avalanches. This runaway process can cause gas breakdown, and can be avoided by adding a second gas in addition to the argon which has excitation modes other than ionisation to absorb these photons before they trigger new avalanches. These gases are known as quenchers, and good candidate gases are those with many degrees of freedom for forming excited states, such as organic molecular gases.

As outlined in section 4.4, the time between the charged particle traversing the straw and a hit being registered is used to determine the position the charged particle crossed the straw. This time is dominated by the time it takes the primary electrons to drift to the wire, and thus the characteristics of this drift must be well understood for the chosen gas mixture. The drift velocity is a net result of the competing effects of the acceleration of the electrons in the electric field, which increases their radial velocity, versus the rate of collisions with gas atoms/molecules, which decreases their radial velocity either via energy loss or a change of direction due to scattering. The collision probability includes quantum effects such as the Ramsauer–Townsend effect [56] which result in a strong energy dependence, such that as the electrons accelerate in the electric field their collision probability changes. This results in a complicated relation between the drift velocity and the electric field that is particular to a given gas composition, and is sensitive to impurities in the gas. It is however possible to select gas compositions where these effects largely cancel such that the drift velocity is relatively constant with respect to the electric field strength. This is called a saturated drift velocity, and results in a drift velocity that is approximately constant with respect to the radial position. This constant drift velocity results in a linear relationship between the drift time and the radial distance from the wire to the primary ionisation position, which is desirable when using the drift time for position measurements: the drift velocity can be determined from data, then accurately modelled by simulation.

Taking these considerations into account, a mixture of argon and ethane is selected. Argon is a noble gas with a relatively high atomic number so is a good ionisation medium, and is also affordable compared to for example xenon which has a lower ionisation energy but is prohibitively expensive. Ethane is selected as a quencher gas as it provides good photo-absorption, and in a 50:50 mix with argon provides a saturated drift velocity. This mixture has been used successfully for a number of drift detectors in previous high-energy and nuclear physics experiments.

Argon-CO₂ in a 80:20 ratio mixture was also considered and tested with detector prototypes. This mixture has similar properties to the argon-ethane 50:50 mix, such as a saturated drift velocity, and was initially preferred to ethane because ethane is flammable. However, the straw leak rate with argon-CO₂ was found to be uncomfortably close to the pumping capabilities of the experiment's vacuum system. The permeation rate of argon-ethane 50:50 through aluminised mylar is known to be lower than for argon-CO₂ 80:20, and so the decision was made to use argon-ethane 50:50 for the tracker gas. Rate of rise tests in a vacuum chamber were performed to confirm the reduction in leak rate.

4.5.3 Tracker module design

As described in section 4.3, each of the three trackers is composed of 8 distinct tracking planes of straws. Each of these planes is constructed as a distinct unit with dedicated gas, power and cooling, and is hereafter referred to as a tracker module. Each module is 32 straws wide, and has four layers of straws, resulting in 128 straws in total per module. A whole tracker has 1024 straws. Each module is identical, and a single straw module is shown in figure 4.4.




Figure 4.4: One straw tracker module. The perspective is downstream, i.e. looking in the direction of e^+ flight.

As each individual straw provides a radial position measurement from the measured hit time, the straws are mounted at two distinct stereo angles such that the radial measurements of two straws at different angles can be combined to yield vertical information. Two of the straw layers in each module are mounted at $+7.5^\circ$ to the vertical, with the other two layers at -7.5° . Having two straw layers at each angle provides overlap to help resolve left-right ambiguities in the drift cylinders (see section 4.6). Offsetting the straws in adjacent layers allows one layer to cover the gaps in coverage of the other, since there are necessarily gaps between the straws to provide sufficient structural support and for the gaps to match the machining tolerances of the milled holes in the manifold housing the straws. This arrangement of straws can be seen in figure 4.5.




Figure 4.5: The straw layers in a straw tracker module. The straws are arranged in two pairs of two layers. In each pair of layers, the straws have the same stereo angle. The two pairs of layers have equal and opposite stereo angles.

The straws are mounted between two aluminium manifolds above and below, with these manifolds protruding from an aluminium flange which bolts into the vacuum chamber wall. Inside the manifolds, which can be seen in 4.6, the straws and wires are fixed under tension, and readout electronics and high voltage are directly connected to

the wire ends. Digital readout signals on flexicables and high voltage cables lead out of the manifolds via channels through the flange. Beyond the flange these openings are extended within aluminium “snouts” until they terminate in a feed-through electronics board to which the downstream electronics are connected. The snouts also have gas connections, with the gas being supplied to one snout, flowing through one manifold and into the straws, then exiting through the other manifold before eventually reaching an outlet on the opposite snout. The purpose of the snouts is to extend these connections radially inwards away from the storage-ring to clear the magnet pole pieces whilst maintaining the gas seal and providing an unbroken Faraday cage for the electronics to shield against RF noise.

A hollow carbon fiber support post at the open end of the module is used to withstand deformation of the manifolds under vacuum. Carbon fiber is chosen here to provide sufficient strength with a small material budget, as these posts are at the edge of the storage region in the line of sight of the calorimeters.




Figure 4.6: A view into the manifold of a straw tracker module. The ASDQ boards (see section 5.1.1) are mounted directly on the ends of the wires, and are connected to the off-detector readout electronics via the flexicables. Copper cooling bars make a thermal connection between the electronics boards and a water-cooling channel drilled into the manifold. The white end caps of the straws can be seen in the lower left, which hold the wires under tension and have openings to allow gas to flow through.

To cool the electronics boards within the manifolds, each manifold includes a cylindrical cooling channel with a concentric internal cylinder, allowing cooling water to be pumped in and extracted within the same channel. Copper bars form a thermal connection between the water-cooled manifold and the electronics.

The arrangement of 8 modules to form a single tracker is shown in figure 4.7. The modules are evenly spaced in the azimuthal direction, but are radially staggered to follow the curve of the storage-ring and ensure each module is as close as possible to the storage region. The vacuum chamber wall to which the module flanges attach are machined in a staircase design to achieve these module positions, as can be seen in figure 4.8.




Figure 4.7: Diagram showing how 8 straw modules are arranged to form a single straw tracker, and how this tracker is aligned relative to the storage region and the downstream calorimeter. The direction of the orbiting μ^+ is shown by the blue line.

4.6 Track formation

A charged particle trajectory in the straw trackers is determined from the DCA of the track to the straw wires using the straw-hit times combined with the known locations and geometry of the straws. It is a complex process combining many layers of algorithms and iteration, and multiple algorithms will be developed and employed in the experiment. Broadly speaking however, the general steps for fitting a track from straw-hits are:




Figure 4.8: Staircase design of the vacuum chamber to which tracker modules are mounted. The left-hand side will be mounted closest to the downstream calorimeter.

- Group straw-hits temporally and spatially to find the hit candidates that could form a track
- Determine the time the charged particle enters the detector to provide a t_0
- Use this t_0 to extract drift times from the hit times
- Convert from drift times to a track-to-wire DCA to define drift cylinders in each hit straw
- Fit a helical trajectory that intersects these drift cylinders whilst taking the magnetic field map into account, using an iterative procedure such as a Kalman filter to tune the track parameters to minimise the residuals of the track to the drift cylinders

All track-fitting procedures will have to accommodate a number of common challenges, such as rejecting hits resulting from noise or secondary particles, the pile-up of multiple coincident tracks, and the ambiguity in knowing which side of the wire the track passed because of radial degeneracy in the track DCA measurement provided from the drift times, known as left-right ambiguity. The track-fitting procedure must also carefully propagate the position uncertainties due to the finite straw resolution and the amount of material traversed by the particle.

A particular challenge to the track-finding in the Fermilab muon g-2 experiment is that the trackers lie in the fringe field of the storage-ring magnet, and the magnitude and direction of the field varies significantly both vertically and radially. The tracks cannot thus be treated as a simple helix as they would in a constant magnetic field, but are instead constructed as a sequence of short helix segments. The varying magnetic field also means that the curvature of the paths of drifting charges varies from straw to straw, complicating the extraction of the track-to-wire DCA from drift times.

Once a track has been formed, it can be extrapolated upstream to the point of the μ^+ decay, allowing a profile of the stored μ^+ beam to be built up. The tracks can also be extrapolated downstream to determine the location the particle impacts the downstream calorimeter. Additionally, the particle's momentum can be estimated from the track curvature.

4.7 Ageing

The performance of drift chambers has been observed to degrade over time in numerous past experiments. This is often referred to as detector “ageing”. Specific observations include a reduction in the gain achieved for a given wire voltage, and increasing dark current [56]. The cause or causes of drift chamber ageing are not well understood, and as such are difficult to mitigate against. Deposits of impurities on the wire or inner surfaces of the drift chamber are thought to contribute to ageing, and due to this procedures will be developed to ensure the purity of gas supplied to the Fermilab muon g-2 experiment straw tracking detectors. The relatively short expected duration of the experiment, 18 months of integrated physics run time, is expected to reduce the potential impact of ageing, and the modularity of the detector and availability of spare modules will allow faulty units to be replaced easily.

Chapter 5

The tracker data path

Beginning with the induced signal on the sense-wires in each straw and culminating with particle tracks and extrapolated μ^+ decay vertices, the flow of tracker data-processing follows a sophisticated path through both hardware and software. This path includes analog and digital frontend electronics, backend aggregator electronics boards, fast Data Acquisition (DAQ) software, slow-control software, online Data Quality Monitoring (DQM) and many stages of offline data processing. These stages are described in the following sections.

5.1 Readout electronics

The tracker readout electronics consist of a hierarchical system of frontend and backend boards through which the data, clock and control signals propagate. The path of signals and data through the boards is summarised in figure 5.1, and the hierarchy and cardinality of the boards is shown in figure 5.2. The boards are described in the following sections.

5.1.1 Frontend electronics

Frontend electronics refers to the electronics boards used to detect signals in the straw sense-wires and characterise and record these signals as straw-hits. There are two distinct boards in this system; the Amplifier Shaper Discriminator with charge (Q) encoding (ASDQ) board and the Time to Digital Converter (TDC) board.

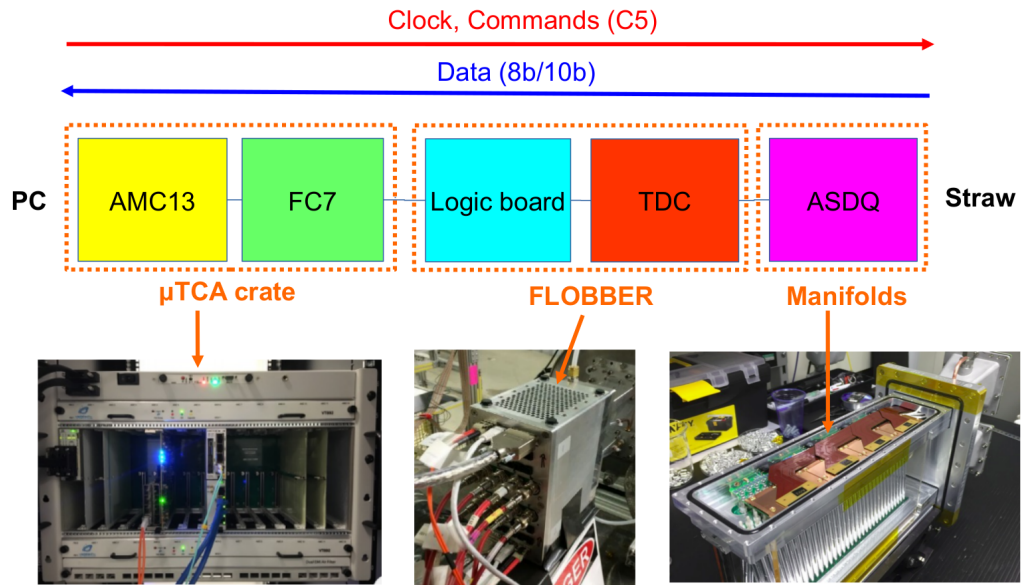


Figure 5.1: The path of clock, control signals and straw-hit data through the various layers of frontend and backend boards in the straw tracker readout system. The physical location of the boards is also indicated.

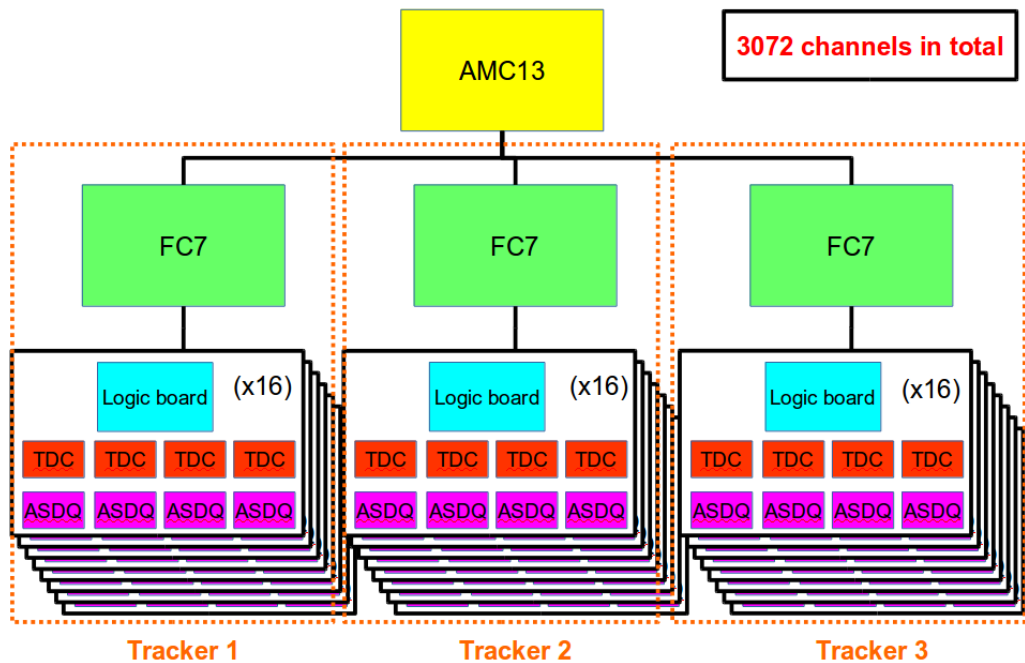


Figure 5.2: The hierarchy of frontend and backend boards and the numbers of each used in the straw tracker readout system.

The role of the ASDQs is to record when a charged particle traverses the straw, known as a straw-hit, which it achieves by triggering when the induced signal on the straw sense-wires passes a threshold [58, 59]. The ASDQs are analog electronics and are mounted directly on end of the straw sense-wires, as can be seen in figure 4.6. These boards also distribute High-Voltage (HV) power to the straw sense-wires.

There are a number of steps in the ASDQ signal processing which are optimised in terms of signal-to-noise, efficiency and timing resolution. The analog signal from the wire is amplified, then shaped to smooth out spikes and other short time-scale features. In doing so, short signals from individual primary ionisation events are integrated into a single pulse. As well as the fast primary electron pulses, the wire signal includes a long tail due to the much slower ion signal. This tail is corrected for and effectively cancelled out during signal shaping in a process known as baseline-restoration. This prevents the long ion-tail overlapping with primary ionisation events from subsequent charged particles and thus allows the primary ionisations to be recorded at higher rates.

The ASDQ then uses a discriminator, which registers when the shaped pulse passes a configurable threshold. This defines the leading and trailing edges of the pulse, which is then output as a digitised signal with transitions at the leading and trailing edges only. Samples of the pulse are not recorded as would be the case for a waveform digitizer. This signal path within the ASDQ is summarised in figure 5.3.




Figure 5.3: Simplified schematic of the signal path in the ASDQ. (a) The raw signal includes a number of fast spikes corresponding to the avalanches from individual primary ionisations. (b) The shaped signal integrates the raw signal into a single pulse. (c) The discriminator records where the shaped pulse crosses a threshold (the red line represents the threshold, the blue lines represent the two crossing points). The time of the first crossing is considered to be the hit time. (d) The output ASDQ signal records only the leading and trailing edges of the shaped pulse following discrimination. Note that the ion tail is not included in this diagram.

The output signal from the ASDQ is carried via a flexicable to a TDC, which is housed outside of the tracker module and storage-ring vacuum chamber in a small custom rack known as the Frontend Low-voltage Optical Box to BackEnd Readout (FLOBBER) rack mounted at the end of the tracker module snouts (see section 4.5.3), which can be seen in figure 5.4. The FLOBBER is used to house the boards such as the TDCs that are required to be close to the straws to minimise signal transmission lengths, but do not need to be directly connected to the straws as the ASDQs do. It is significantly easier to access and cool boards in this external rack compared to those mounted within the manifolds which are within the vacuum chamber, hence the choice of this system. There is one FLOBBER connected to each tracker module.




Figure 5.4: Four FLOBBERs mounted on the outside of the vacuum chamber, along the inside radius of the storage-ring. Each is attached via a snout through the flange to a tracker module within the vacuum chamber.

The TDC is connected to an experiment-wide external 40 MHz clock, and time stamps the transitions in the ASDQ output signal with 625 ps precision. Each transition is written to the output data stream as a hit-word which encodes the hit channel, transition time and whether the edge is leading or trailing. These hit-words are the fundamental unit of all downstream data processing in the tracker, with the leading-edge-hit-words being considered the hit time in a straw.

The TDC accumulates data during triggered time windows, with a single accumulation covering one μ^+ fill in the storage-ring. These triggers are received from the backend

boards via a line carrying both commands and the clock using the Clock-Command Combined Carrier Coding (C5) protocol [60]. At the end of each accumulation all hit-words are packaged into a data-block which is output using 8b10b encoding [61]. Up to 2016 hit-words may be included in each event for each TDC, along with auxiliary information in a header such as the time the trigger was received. Both C5 and 8b10b protocols are DC-balanced to decouple sender and receiver, and can carry both data and clock on a single line.

ASDQs and TDCs are paired one-to-one: each has 16 channels where one channel is connected to one straw. Each straw module is composed of 128 straws, and hence requires 8 ASDQ-TDC pairs. The full tracker readout system comprises 3 trackers each with 8 modules, and thus requires 192 ASDQ-TDC pairs.

5.1.2 Backend electronics

Backend electronics refers to the electronics boards used to interface to and synchronise the large number of frontend electronics boards used by the tracker. In principle a single master board could be designed to provide all the required interfaces to all frontend boards. However, cost savings can be achieved by re-using existing board designs, and this choice for the Fermilab muon g-2 experiment leads to a hierarchical backend electronics system.

The lowest level backend board is the so-called logic board (LB). The LB is a custom board designed to provide an interface to four ASDQ-TDC pairs, and so two LBs are required per tracker module. The LBs are mounted in the FLOBBERS, next to the TDC boards. The LB provides three external interfaces; a data fiber-optic link, a configuration serial line, and a low voltage (LV) supply line. A LB communicates with the TDCs via a backplane in the FLOBBER and with the ASDQs via the flexicable through the snout. The fiber-optic data link is the connection to the higher-level backend boards, and is used to receive the clock and control signals via the C5 protocol and send data via the 8b10b protocol. Fiber-optic cables are selected here as they allow transmission of these signals over large distances, meaning that the higher-level boards can be located far from the storage-ring where they are then free to use magnetic components. Electronics close to the ring such as those in the FLOBBER must be free of magnetic components to avoid perturbing the storage-ring magnetic field. Control signals and the clock are fanned out from the

LB to the connected TDCs, and the data from all connected TDCs is aggregated by an event builder in the LB firmware to provide a single data-block containing data from all TDCs in response to a trigger command. The LB serial line is used to provide read/write access to the LB registers. These registers store configuration parameters and the status of the LB, TDC and ASDQ boards. The LV line provides power to the LB and the frontend boards.

The FC7 [62] and AMC13 [63] boards are housed further downstream in a microTCA crate [64] away from the storage region. These boards are re-purposed from their original design for the Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC) in CERN. Both boards act in a similar manner to the LBs by fanning out clock and control signals to lower-level boards and aggregating the data-blocks from all connected boards into a single data-block using an internal event builder. Each tracker uses a single FC7 board connected to the data-line fibers of all LBs for a given tracker, and hence there are three FC7s in the full readout system. A single AMC13 is connected via the microTCA backplane to all three FC7s, and provides a singular interface to the tracker clock, control and data readout system.

The AMC13 is connected via a Gigabit Ethernet (GbE) fiber-link to a PC network card, allowing the final “per fill” data-block containing hits from all frontend boards to be read out by the DAQ software and stored on disk. The microTCA crate includes a controller board that provides an ethernet connection to the AMC13 and the three FC7s for configuration purposes. Both the FC7 and AMC13 feature DDR3 memory blocks which allow the data from multiple triggers to be buffered in hardware, allowing the periodic long gaps between μ^+ fills to be exploited for readout. The AMC13 is also connected as consumer to the experiment-wide clock and command system.

The C5 and 8b10b protocols are used consistently through-out the frontend and backend board links for clock and control transmission and data-block sending respectively. The use of the microTCA crate and the AMC13 as the high-level backend board and clock and trigger consumer is common to other readout systems in the Fermilab muon g-2 experiment, including the calorimeters and fiber harps. This facilitates the use of common software for tasks such as slow-controls and data unpacking.

The final component of signal aggregation employed in the tracker system is a controller board connected to the serial lines from the LBs, with one controller being used per tracker. These three controller boards connect to a PC via USB, and can route serial commands to the desired LB and return the response, or broadcast to all LBs for that tracker simultaneously. This is also the means by which firmware is loaded on the LBs and TDCs.

Despite the number of layers in this hierarchical system, the final system is self-contained, and presents only the following external interfaces to the user :

- One fiber-optic data line (PC to AMC13)
- One fiber-optic clock and control receive line (experiment clock and control system to AMC13)
- One ethernet configuration line (PC to AMC13)
- Three USB serial configuration lines (PC to serial controller board)

5.2 Data acquisition software

The straw tracker DAQ software is responsible for configuring and controlling the readout electronics, and reading out the event data from the electronics and storing it. The requirements of the DAQ system are thus as follows:

- Store user-configurable run-time configuration parameters in a database
- Translate the run-time configuration parameters into board register values and write these to the boards at the start of each data taking run
- Ensure all boards and interfaces are correctly reset prior to the start of each data taking run
- Asynchronously receive physics data from the AMC13 data fiber and store on disk
- Monitor the status of the system during data taking and either remedy or report errors
- Provide a user-friendly run-control interface for operators

The following sections describe the design of the DAQ system and how this meets these requirements.

5.2.1 DAQ software description

The core of the DAQ software is a number of C++ “controller” classes for each readout board, each implementing a Finite State Machine (FSM) for managing the execution of the software. There is one generic controller class for each readout board type; `AMC13Controller`, `FC7Controller`, `LogicBoardController` and `TDCController`. The exception is the ASDQ, which is not directly communicated with and is instead controlled via other boards. ASDQ control is included in the `TDCController` class, as the TDCs and ASDQs are paired one-to-one. An instance of the controller class is created for each physical instance of the board, and that controller instance is uniquely responsible for the configuration and control of that board.

The hierarchy of readout boards present in the system is stored in the MIDAS Online Data Base (ODB), where MIDAS is the experiment-wide DAQ system [65]. At start-up, the DAQ software parses the ODB and dynamically instantiates a controller class for each board present. The controller instances are created in a hierarchical framework representing the board interfaces shown in figure 5.2, with each controller instance being a child of the controller above it, and itself having children controllers below it. For example, a given `FC7Controller` instance is the child of the `AMC13Controller` instance, and will itself have 16 `LogicBoardController` instances as children. This dynamic, object-oriented system means that the structure of the readout system can be quickly changed to reflect the boards present without the need to modify or re-compile the software, making it easy to configure the DAQ software for test stands with small numbers of boards in addition to the final DAQ system.

All controller classes inherit from a `ControllerBase` base-class. This includes accessor functions for the parent and children of the controller, using base-class pointers to remain generic. The DAQ software FSM is also implemented by this base-class. The FSM features a number of states, as well as defined commands which trigger the transition from one state to another. These commands have corresponding virtual functions in the base-class which can be implemented in the specific board controller classes, allowing the controllers to perform actions upon receipt of these commands. The command execution cascades down through the controller hierarchy,

resulting in the function corresponding to the executed command being called for all boards at all levels of the hierarchy. The command sequence subsequently flows back up to the highest-level board, meaning that each controller can perform actions both before and after their children controllers. Multi-threading is employed to ensure that the DAQ system can still respond to commands whilst other processes are underway.

The DAQ software FSM is described in figure 5.5. The software is immediately initialised at start-up, dynamically creating the controllers. The `Configure` command is used to parse the current state of the ODB and write configuration data to the board registers to configure the system ready for data taking. AMC13 and FC7 register read and write operations are performed using the `IPbus` [66] software and firmware suite of tools developed for CMS, and via serial lines to the LBs and TDCs. The serial lines are also accessed by separate slow-control software, and a software manager system has been implemented to allocate access to the line and distribute signals to consumers via sockets.

The `Start` and `Stop` commands are used to start/stop the readout loop in the software. The readout loop is a function that continuously reads data from the AMC13 via fiber into software and buffers it. The data can then be directly accessed, or written to an output stream such as a file. The readout loop is asynchronous rather than triggered, constantly polling a socket, connected to the AMC13 fiber via a PCIe card in the PC running the DAQ software, for new event data.

The `Send-trigger` and `Send-trigger + readout` commands are used for testing and special run modes, but not for nominal data taking operations. These functions allow artificial triggers to be injected into the system.

The core tracker DAQ software is designed to be as generic and flexible as possible and well encapsulated, meaning that it can be used by various types of high-level software. For nominal data taking operations, the Fermilab muon g-2 experiment run-control system, `MIDAS`, is used. This is covered further in section 5.2.3. `python` bindings have also been implemented for the `C++` core tracker DAQ software using `Boost python` [67], meaning that the DAQ system can be instantiated and used in `python` scripts. This allows rapid development of test scripts, and a number of non-standard run modes such as scans characterising noise in the detector system are implemented using `python`. The use of `python` bindings also allows the use of `PyQt` [68] for rapid development of Graphical User Interfaces (GUIs). For example,

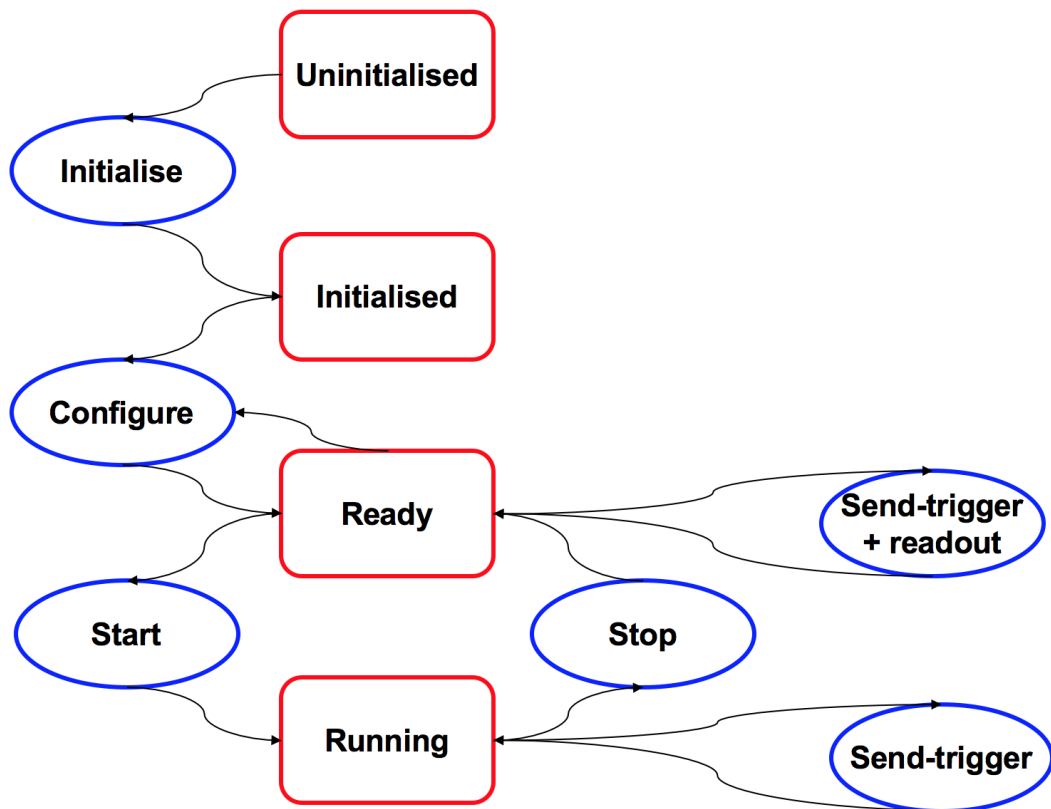


Figure 5.5: The tracker DAQ software FSM. Red boxes represent states. Blue ellipses represent commands, with the arrows indicating the states a command may be executed on and what the resulting state of the command is.

a GUI interface to the FSM and an IPbus read and write GUI, which is shown in figure 5.6, are implemented in such as way.

	Register	Mode	Permissions	Value
2174	STATUS.AMC_TTC.BC0_LOCKED_M...	SINGLE	R	0x0
2175	STATUS.AMC_TTC.TTC_LOCKED_MA...	SINGLE	R	0x0
2176	STATUS.AMC_TTS_STATE	SINGLE	R	0x0
2177	STATUS.DIAG.DDR_RESET_FAILED	SINGLE	R	0x0
2178	STATUS.DIAG.DDR_STATUS_HI	SINGLE	R	0x0
2179	STATUS.DIAG.DDR_STATUS_LO	SINGLE	R	0x0
2180	STATUS.EVB	SINGLE	RW	0x70000009
2181	STATUS.EVB.DATA_READY_MASK	SINGLE	R	0x0
2182	STATUS.EVB.DDR3_WRITE_PORT_F...	SINGLE	R	0x0
2183	STATUS.EVB.DDR3_WRITE_PORT_R...	SINGLE	R	0x7
2184	STATUS.EVB.EVENT_SIZE_MASK	SINGLE	R	0x0
2185	STATUS.EVB.OVERFLOW_WARNING	SINGLE	R	0x0
2186	STATUS.EVB.SYNC_LOST	SINGLE	R	0x0
2187	STATUS.EVB.TCP_BUFFER_AVAILABLE	SINGLE	R	0x1
2188	STATUS.FIRMWARE_VERS	SINGLE	R	0x813f
2189	STATUS.FPGA.DIE_TEMP	SINGLE	R	0x238
2190	STATUS.FPGA.DNA_HI	SINGLE	R	0x242800
2191	STATUS.FPGA.DNA_LO	SINGLE	R	0x1231b05c

Figure 5.6: GUI for board register read and write operations using IPbus. The GUI periodically reads and displays all IPbus registers for a given board. Clicking on a given register allows a new value to be written to the register.

A data manager class provides helper functions for handling data from the tracker DAQ system. Foremost of these are unpacking functions that are used to unpack the raw data byte-stream into its constituent data-blocks from each board, and further break up this data into logical units such as headers and footers, and hit-words. The data is subsequently unpacked into the physics data parameters, such as the time and channel of straw-hits. The unpacked data is stored in C++ structures that are nested in a hierarchical manner analogous to the boards themselves. The data manager and data structures also feature python bindings and can be directly used in python scripts.

The tracker DAQ encompasses error handling and a logging system for displaying messages, warnings and errors, and for handling the safe shutdown of the system in the event of a critical error. The level of output messaging can be tuned, and the messages are displayed both in the terminal and in a viewer accessible from the

top-level MIDAS run-control web interface, as well as being written to disk as log files for persistent storage.

5.2.2 Emulators

A software emulator has been developed for each readout board. These emulators provide similar interfaces to the real boards, allowing the DAQ software to interact with them as if the real DAQ hardware was present. This enables the DAQ software to be tested and developed without the presence of real boards, which proved particularly useful in remotely testing the experiment-wide DAQ without the boards being physically at Fermilab.

Each emulated board responds to a trigger by composing a fake data event. These data events are passed up the board hierarchy and aggregated as is the case for the real boards, yielding a single simulated AMC13 data event per trigger. The format of this data event is the same as that resulting from the real hardware, allowing the data to be unpacked using the standard unpacking functions which also aids software development.

An IPbus emulator was used to mimic the IPbus interface with the boards. Memory is allocated in software representing the board registers, and the IPbus read and write operations conducted by the DAQ software access this software memory via the UDP protocol. The read and write actions trigger callback functions in the emulators, which model various behaviours of the boards such as the data-event generation described above. The fidelity of the emulation in each board is sufficient to allow the DAQ software to run in a manner that is blind to whether the real or emulated boards are present.

5.2.3 MIDAS

The Fermilab muon g-2 experiment run-control and high-level DAQ system is an implementation of the Maximally Integrated Data Acquisition System, or MIDAS, developed at the Paul Scherrer Institute (PSI) in Switzerland and TRIUMF in Canada [65]. MIDAS provides a number of utilities including a run-control system with a web interface, an event builder for gathering data fragments from distributed

hardware and software sources, and a means for writing the data stream to a dedicated MIDAS binary file type. MIDAS also provides the ODB used to store the experiment's DAQ and run-time configuration information. The web interface is shown in figure 5.7.

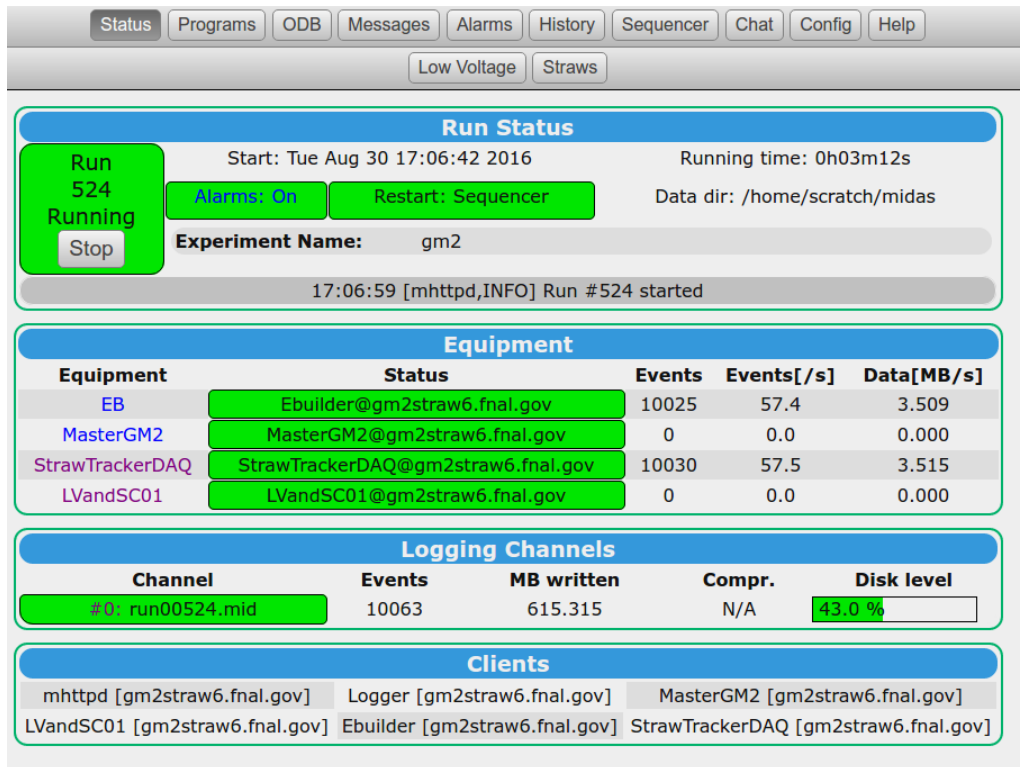


Figure 5.7: The top-level status page of the MIDAS web interface.

To run a given detector system within MIDAS, a software “frontend” for each detector component must be developed. This connects to the MIDAS command signals and run-control FSM, and provides an interface for the detectors to write their data to MIDAS “banks” that are passed to the event builder and subsequently the output file stream. Frontends communicate with one another via a Remote Procedure Call (RPC) service.

The tracker-DAQ frontend instantiates the core tracker software, creating the board controllers and providing an interface to the tracker FSM. Functions in the frontend are called upon receipt of the MIDAS run-control commands, which subsequently trigger commands in the tracker DAQ FSM. The system is developed in a way that minimises the memory overhead, particularly in avoiding making duplicate copies of event data.

Additional frontends handle the configuration and monitoring of the low and high voltage supply systems for the tracker, as well as the readout of data in the tracker slow-control systems.

The ODB stores the configuration parameters for each frontend. The tracker DAQ settings are arranged hierarchically and are parsed by the tracker DAQ software to configure the software and boards at the start of each run. The ODB can be accessed via the same web interface as the run-control system, as can be seen in figure 5.8.

Key	Value	
ID	0 (0x0)	
Present	y	
Enabled	y	
Num FC7s	1 (0x1)	
Readout using IPbus	y	
DAQ link IP address	192.168.41.2	
DAQ link port	4660 (0x1234)	
DAQ link SFP	0 (0x0)	
Fake data	n	
Fake data size	10 (0xA)	
Fake data channel mask	3 (0x3)	
IPbus connection file	`\${GM2TRACKERDAQ_CFG}/connections.xml`	
T1 IPbus device ID	amc13-gm2-ucl.T1	
T1 IPbus emulator device ID	amc13-gm2-emulator.T1	
T1 IPbus emulator UDP port	50001 (0xC351)	
T2 IPbus device ID	amc13-gm2-ucl.T2	
T2 IPbus emulator device ID	amc13-gm2-emulator.T2	
T2 IPbus emulator UDP port	50002 (0xC352)	
Fake trigger server	n	

Figure 5.8: The MIDAS ODB, as viewed using the MIDAS web interface. The page shown contains the run-time settings for the tracker AMC13.

A user operates the experiment DAQ from the MIDAS web interface by enabling the frontends for all desired detectors and starting a data run. The detector readout systems then operate continuously and respond to accelerator signals until the run is stopped, and the resulting data is subsequently available in a MIDAS file.

5.2.4 Data rates

It is crucial that the readout electronics and DAQ software are fast enough so that the readout of new data in the electronics is not inhibited. There are two cases

in which this might occur. The first is if the frontend boards have not finished processing a trigger before the next trigger is received. The second is if the data buffers in the backend boards become full, which would occur if the average rate at which physics data is written to disk is lower than the average rate of data taking.

A model of the data propagation through the readout electronics has been developed to determine the expected performance of the DAQ system. This model encapsulates the board hierarchy, the data transmission rates at each interface between boards as listed in table 5.1, and the inherent parallelisation in the system. It also considers the μ^+ fill structure delivered by the accelerator to the storage-ring, the duration of a single μ^+ fill, and the time over which the TDCs accumulate data.

The fill structure for the Fermilab muon g-2 experiment storage-ring is shown in figure 5.9. There are two features of relevance to the discussion of data rates. The first is that there are regions where fills are closely grouped together with 10 ms gaps between them. This sets the rate at which the frontend electronics must be able to process triggers. The second feature is that there are two large gaps in the cycle where no fills occur. These gaps appear due to the sharing of the accelerator infrastructure with other experiments, notably the Fermilab neutrino programme. These large gaps are available for the DAQ software to transferred buffered data to disk, and so the DAQ software need only meet the average rate of fills within this structure, which is approximately 12 Hz.

Link	Tx type	Payload fraction	Tx rate [Mbps]	Payload Tx rate [Mbps]
TDC to LB	8b10b	80%	25	20
LB to FC7	8b10b	80%	125	100
FC7 to AMC13	8b10b	80%	5000	4000
AMC13 to PC	TCP/IP	Up to 97%	5000	Up to 4870

Table 5.1: Data rates in interfaces between components in the DAQ system. The transmission (Tx) protocols and corresponding data transmission rates are shown. Payload fraction refers to the fraction of each data packet containing payload data, where the remaining packet capacity is occupied by information required by the protocol.

The TDCs will typically accumulate data from the first ~ 1 ms of the fill (μ^+ lifetime in the storage-ring is 64 μ s) and their buffers can record 2016 hits, such that the




Figure 5.9: Proposed fill structure for the Fermilab muon g-2 experiment [2].

maximum average hit rate a single TDC can accommodate without deadtime is 2 MHz. In this worst-case scenario where all TDC buffers are full, the model predicts that the data for a fill will reach the first hardware buffer, in the FC7, in 6.9 ms and the final hardware buffer, in the AMC13, in 8.0 ms following the trigger, within the 10 ms window available.

Dedicated testing was performed to determine the average rate at which the DAQ can write to disk from the AMC13. The software readout is performed in parallel to the electronics readout, and can thus be considered separately. The time taken by the DAQ software to process an event includes the data transfer via fiber from the AMC13 to the PC, buffering in software, writing data to disk and the communication and error handling software overhead. The results of this testing are shown in figure 5.10, which indicates the maximum average event rate that may be processed by the DAQ system versus size of the event. The maximum data load tested corresponds to the maximum possible data load on the system, resulting from saturated TDC buffers as previously described. At this event size, the maximum average event rate the DAQ software can sustain is measured as approximately 90 Hz, significantly greater than the required 12 Hz rate.

Using this combination of testing and modelling, it is concluded that the tracker DAQ system is sufficiently fast to meet the requirements of the Fermilab muon g-2 experiment, even in the worst-case scenario data load. Moreover, the expected number of straw-hits from decay e^+ per fill according to Monte-Carlo simulation is approximately 8% of the data volume considered in the worst-case scenario, leaving ample margin. The DAQ software has routinely been run at 50 Hz using radioactive sources in a test stand at Fermilab.




Figure 5.10: Maximum event rate achievable by the DAQ software versus event data size (blue line), measured at a test stand. Also indicated are the Fermilab muon $g-2$ experiment fill rate (red), expected event size during physics runs (green) and maximum possible event size (magenta). The time taken to readout one event scales linearly with the event size, and as such the number of events that can be readout in 1 s scales as $1/(\text{event size})$.

5.3 Offline processing

The offline processing of data in the Fermilab muon g-2 experiment is performed using the `art` event-processing framework developed by Fermilab [69]. This is a C++ software suite that allows users to develop modules that process data in discrete events, and to populate the events with data, known as `art` records. The event data is stored as a series of `art` records which are persistently stored on disk using the ROOT dictionary system. The `art` records themselves are C++ classes or structures, or vectors thereof.

The offline system is used to process the raw detector data, converting it into meaningful physics data. The raw data must be calibrated and checked for errors, then forwarded on to higher-level processing such as event reconstruction and eventually the analysis modules for measuring ω_a .

5.3.1 MIDAS-to-art

The first stage in the offline data processing chain is to create the initial `art` events and populate them with the raw data from the detector systems. The detector raw data is stored in MIDAS banks, which are blocks of raw data plus header information. An input driver module called `MIDAS-to-art` has been developed to accomplish this, which parses information from the MIDAS output streams such as a MIDAS file to create `art` run and event records in the `art` framework and store relevant metadata. The MIDAS and `art` definition of an event is the same, and corresponds to one μ^+ fill in the storage-ring. Generic data structures have been developed to record MIDAS banks in `art` records regardless of the type of data in the bank.

The contents of the ODB are also recorded by MIDAS at the start and end of each run, and are parsed by the `MIDAS-to-art` module and stored in the `art` run record. In this way the detector configuration data is made available to offline analysis.

`MIDAS-to-art` can be run in two modes: offline or online. In offline mode, the source of MIDAS data is a MIDAS file containing the stored data. This is the main way in which data is processed, and can take place at any time after the run has taken place using archives of the MIDAS files containing the experiment's raw data.

In online mode, `MIDAS-to-art` forms a live socket connection to the top-level `MIDAS` system through which it asynchronously receives `MIDAS` event data such as `MIDAS` banks, as well as the run and event signals and the current `ODB`. This communication system is outlined in figure 5.11. `MIDAS-to-art` can be used in this mode as the input driver for online monitoring processes which are constantly running and processing detector data in real time as the experiment runs.

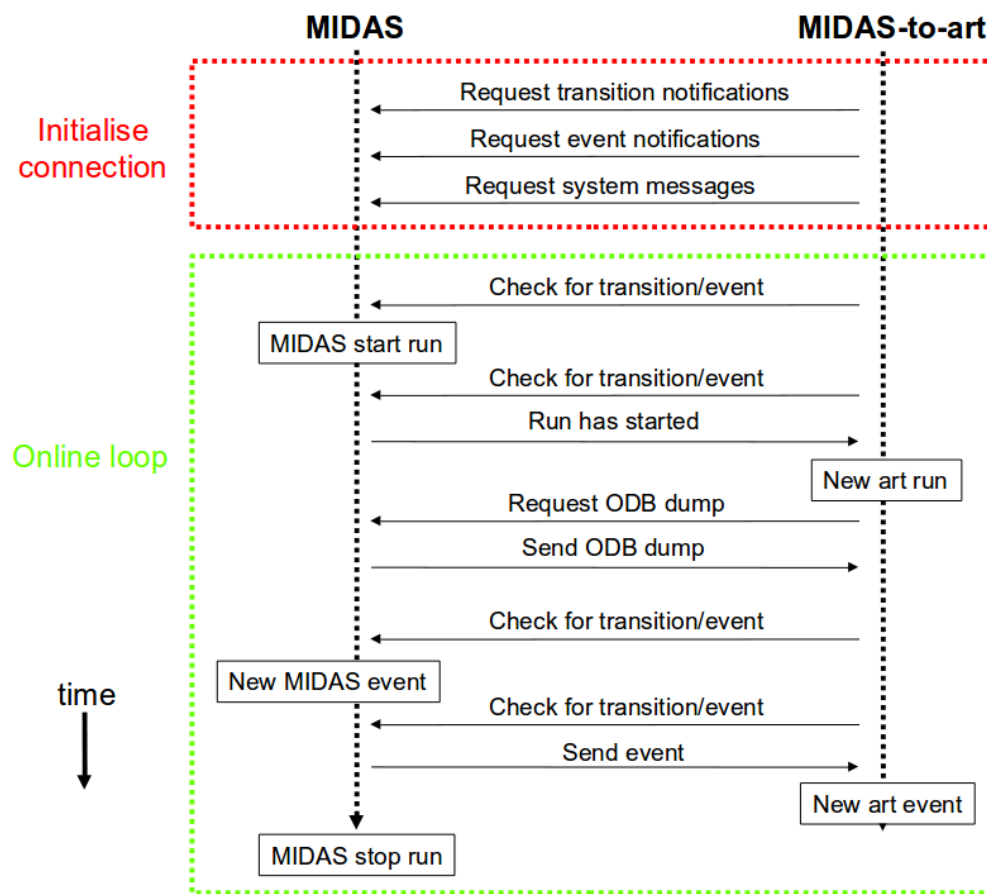


Figure 5.11: Communication between `MIDAS` and `MIDAS-to-art` when running in online mode.

Utilising `MIDAS-to-art` in online mode for online monitoring has the major advantage of using the experiment's offline framework, meaning that all offline data processing, reconstruction and analysis modules may be used in online monitoring. Additionally, `MIDAS` can support multiple connected `MIDAS-to-art` clients, meaning that different monitoring services may be developed and run concurrently. An example would be a system monitoring the low-level status of the boards, running simultaneously with an event display system showing reconstructed straw tracks and calorimeter events.

These various monitoring processes can also be run on different physical workstations, enabling the computational load to be distributed. Breaking up monitoring processes in this way also improves reliability, as failures in one monitoring system do not impact the others.

5.3.2 Tracker offline data processing

MIDAS-to-art does not perform any true processing of the raw data, but simply makes this data available in art. The processing of these raw data-blocks is achieved using a sequence of art modules. For the tracker, the stages implemented to provide straw-hit data from the raw detector data are shown in figure 5.12.

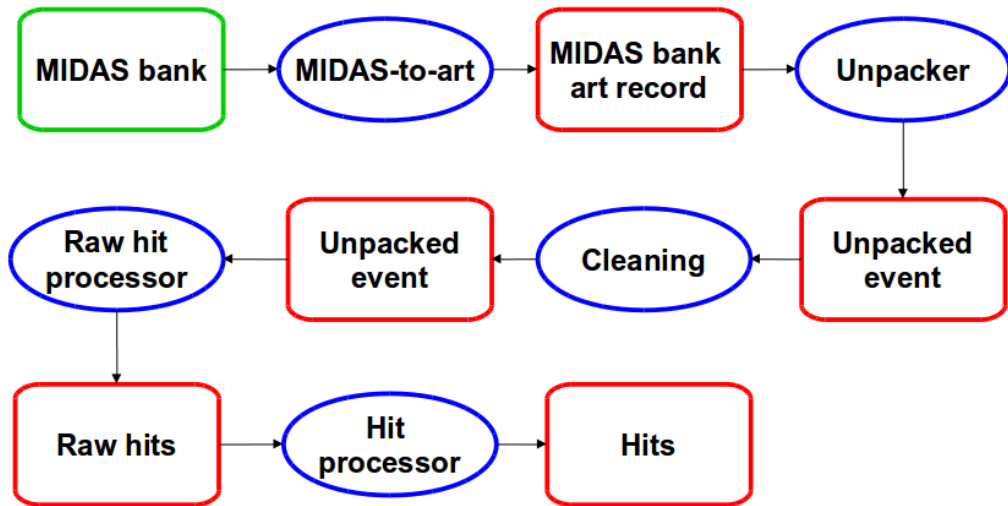


Figure 5.12: Offline data processing stages for the straw tracker data. The green box represents MIDAS data, the red boxes are art records, and the blue ellipses represent art modules performing the data processing.

The unpacking stage uses the data manager class described in section 5.2.1 to extract the hit and header information from the raw data. The subsequent cleaning-stage checks the data for error conditions such as hardware error flags, or missing or out-of-sync data from the readout boards. Data with errors will be rejected at this stage, although the user has the opportunity to override particular errors if desired.

The raw-hit processing stage pairs the leading and trailing edge hit-words and performs some data conversions such as converting from clock ticks to real time units. Synchronisation of the data in different boards and also other detector systems also

takes place here, correcting for effects such as cable delays and the time taken for signals to propagate through the hierarchy of boards.

The final stage of this low-level tracker data processing makes the hits more amenable for physics analysis, primarily via the mapping of electronics channels to the straws and their geometry. Following this stage, the hits are ready for use in track reconstruction algorithms and physics analyses.

5.4 Data monitoring systems

Online data processing and display tools have been developed for the trackers to allow operators to monitor the health and performance of the system in real time whilst it is running. Figures 5.13 and 5.14 show the running tracker online data monitoring system as used at a test stand, with low-level readout diagnostics and hit counts in ASDQ-TDC channels displayed respectively.

Online processing of the straw-hit data is performed using `MIDAS-to-art` as described in section 5.3.1. The data processing chain described in section 5.3.2 is applied online to this raw data, and a dedicated data publishing module delivers plotting data to an online web display. This data transmission is performed using the `ZeroMQ` distributed messaging toolkit to provide reliable high-speed transmission of serialised hit data as a byte-stream to the receiving clients [70].

The web display system uses the `Node.js JavaScript` run-time environment to create a webserver and receive, unpack and process data from the serialised data publisher [71]. `Node.js` is an event-based system that is well suited to the event-based nature of the tracker data path. Plotting on the web display is performed using `plotly`, which supports `JavaScript` and has a large library of plotting functionality, including user interactivity [72]. This system using `MIDAS-to-art`, `ZeroMQ`, `Node.js` and `plotly` for online monitoring is used across multiple g-2 detector and field measurement systems, and is well tested at a number of test stands with representative data loads. `ZeroMQ`, `Node.js` and `plotly` are well used industry tools and have an active support base.

Performance is critical in the online monitoring systems if they are to keep up with the detector readout. The data publisher and plotters are designed with this in mind

to minimise the amount of processing performed and the data volume transmitted. Additionally, the monitoring web display is designed such that only pages currently being accessed by a user are updated.

[CONNECTED home channels](#)

Tracker Test Stand Summary

RUN 610 EVENT 53577

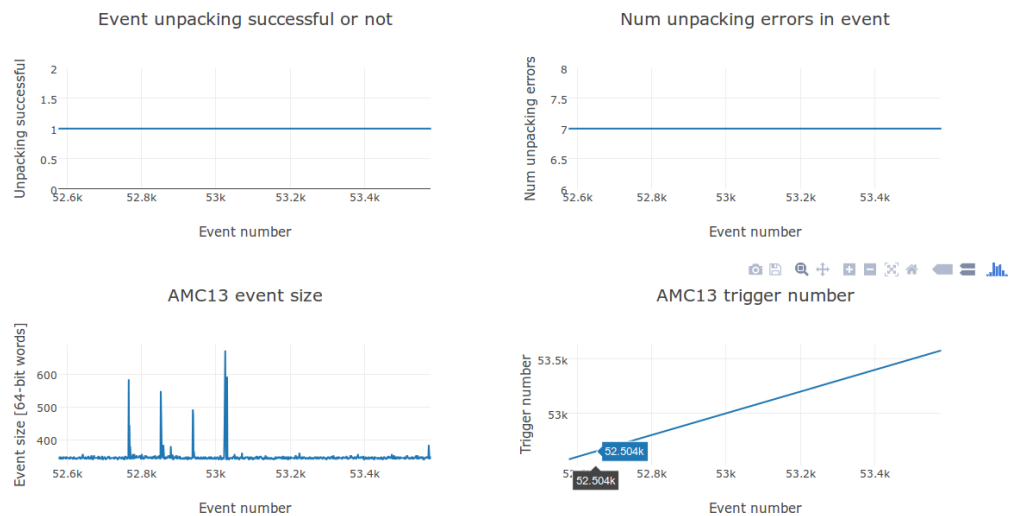


Figure 5.13: Online tracker data monitoring web display. Readout hardware monitoring is shown.

In addition to monitoring detector hit data, it is also desirable to monitor the tracker slow-control data. This allows detector hardware parameters such as temperatures and currents to be observed and potential problems identified. A tracker slow-control monitoring web display has been implemented using the **Flask** light-weight **python** webserver [73]. Slow control data is periodically written by a **MIDAS** frontend to an SQL database. A user can select a parameter of interest from the web display, and the monitor will fetch measurements of that parameter for the specified time window and plot them on the web display. An example is shown in figure 5.15.

These monitoring tools will continue to be updated as the full tracker system is installed in the g-2 storage-ring and commissioned.

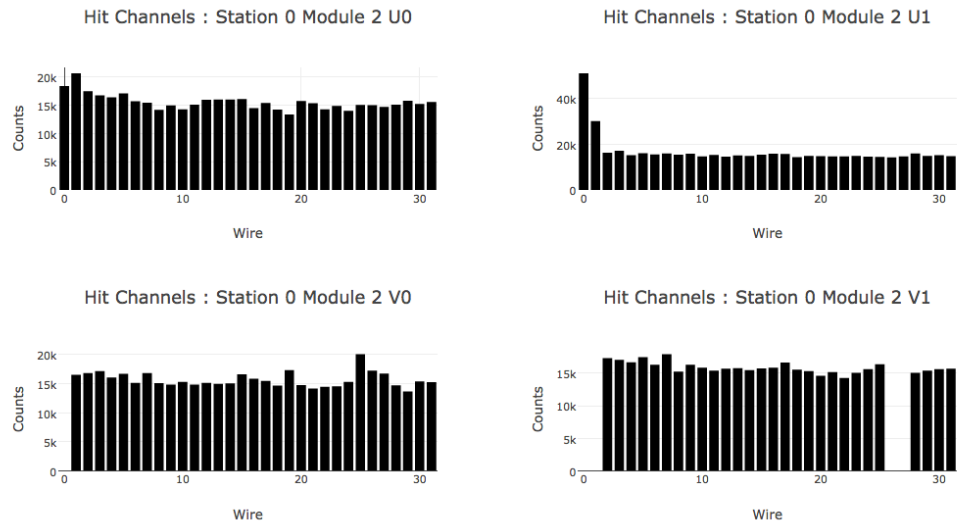


Figure 5.14: Online tracker data monitoring web display. Counts of hits in ASDQ-TDC channels (each corresponding to a single straw) for single tracker module are shown.



Tracker Slow Control DB Monitor

[HOME PAGE](#)

[LAB-3 TEMP](#)

[TRACKER 1](#) »

[TRACKER 2](#) »

[TRACKER 3](#) »




Figure 5.15: Tracker slow-control data monitoring web display. This example shows the temperatures of the boards in the manifold and FLOBBER for a single tracker module.

Chapter 6

Testing the tracker

Prototype and production tracker modules have been tested extensively during the detector development process. These tests include three test beams of prototype modules at the Fermilab Test Beam Facility (FTBF), and at a number of test stands [74]. The purpose of these tests is to verify that the design and construction of the detector hardware, readout electronics and online and offline software will meet the design requirements, and to characterise the detector performance.

The following sections describe the tests performed and the detector performance results obtained.

6.1 Test beams

Three test beams have been performed using straw module prototypes at the MTest beam line at the FTBF. Results from the most recent test beam, conducted in 2015, are presented in the following sections.

The detector arrangement at the 2015 test beam is shown in figure 6.1. A single prototype tracker module was mounted inside a vacuum chamber, which itself was mounted on a motion table that could be remotely translated in vertical and transverse directions relative to the beam. The prototype is very similar to the final production design, in part due to the successful performance of the prototype during the test beam meaning that significant changes were not required.

The prototype module is shown in figure 6.2. The tracker frontend readout electronics were mounted in a prototype FLOBBER attached to the flange via the snouts, again in a system very similar to the final system that will be used at the Fermilab muon g-2 experiment.




Figure 6.1: Detector arrangement at the 2015 test beam. The beam path is indicated by a green line. The location of the tracker module within the vacuum chamber is indicated by a red box.

In order to seed hit-finding and determine the straw resolution and efficiency, two auxiliary detector systems were used. The first was a series of four Multi-Wire Proportional Chambers (MWPC). These tracking detectors consist of gas filled planes containing sense wires with 1 mm spacing. The detection principle is similar to that of the straws, with traversing particles ionising the gas and the resulting mobile charge carriers being collected on the sense wires and triggering readout electronics.




Figure 6.2: Prototype tracker module used for the 2015 test beam.

Each plane measures a 1D position, with each MWPC featuring two planes mounted orthogonally to produce 2D hits. The MWPCs are used to measure the beam profile, and are large enough to cover all particles incident on the tracker module. Two MWPCs were placed upstream of the tracker module and two downstream such that the beam profile could be measured before and after the tracker module and effects such as multiple scattering could be quantified.

The second auxiliary detector system was a silicon telescope. This telescope consists of eight 1 cm^2 silicon strip layers arranged in four planes. Each layer is $120\text{ }\mu\text{m}$ thick. As for the MWPCs, each silicon layer provides a 1D measurement so two orthogonal layers per plane are used for 2D hit position reconstruction. The silicon strip pitch is $80\text{ }\mu\text{m}$, with a position measurement resolution of $23\text{ }\mu\text{m}$. Two planes were upstream of the tracker module and two downstream, allowing tracks to be formed that intersect the straws and can be compared to reconstructed straw-hit positions. Thin scintillators, at the upstream and downstream ends of the telescope, trigger the detector when a coincidence is recorded. Light-tight boxes surround the detectors to reduce noise.

The MWPCs and silicon telescope can be seen in figure 6.3.




Figure 6.3: Auxiliary detectors used in the 2015 test beam

The silicon telescope and tracker module used a common clock and trigger to synchronise the two systems and so facilitate the matching of hits between the two detectors offline.

6.1.1 The MTest beam

The MTest beam line provides 4 s spills of 120 GeV protons once per minute. The number of protons per spill is tune-able, and typically 1×10^5 protons per spill were requested.

The spatial profile of the proton beam is measured by the MWPCs and can be seen in figure 6.4. It is well modelled in both the horizontal and vertical directions by a central core Gaussian with secondary Gaussian tails. This measured beam profile is used as an input to the test beam simulation.

The time distribution of the protons has two distinct features. The protons are extracted from the Fermilab main injector synchrotron ring. This extraction is tuned to deliver single protons, spaced in time by $\mathcal{O}(\mu\text{s})$. However, around 35% of the time the extraction process delivers two or more protons instead, which arrive within a few ns of each other. These are known as multiple-occupancy events, and cause pile-up in the detectors which must be handled by the reconstruction software.

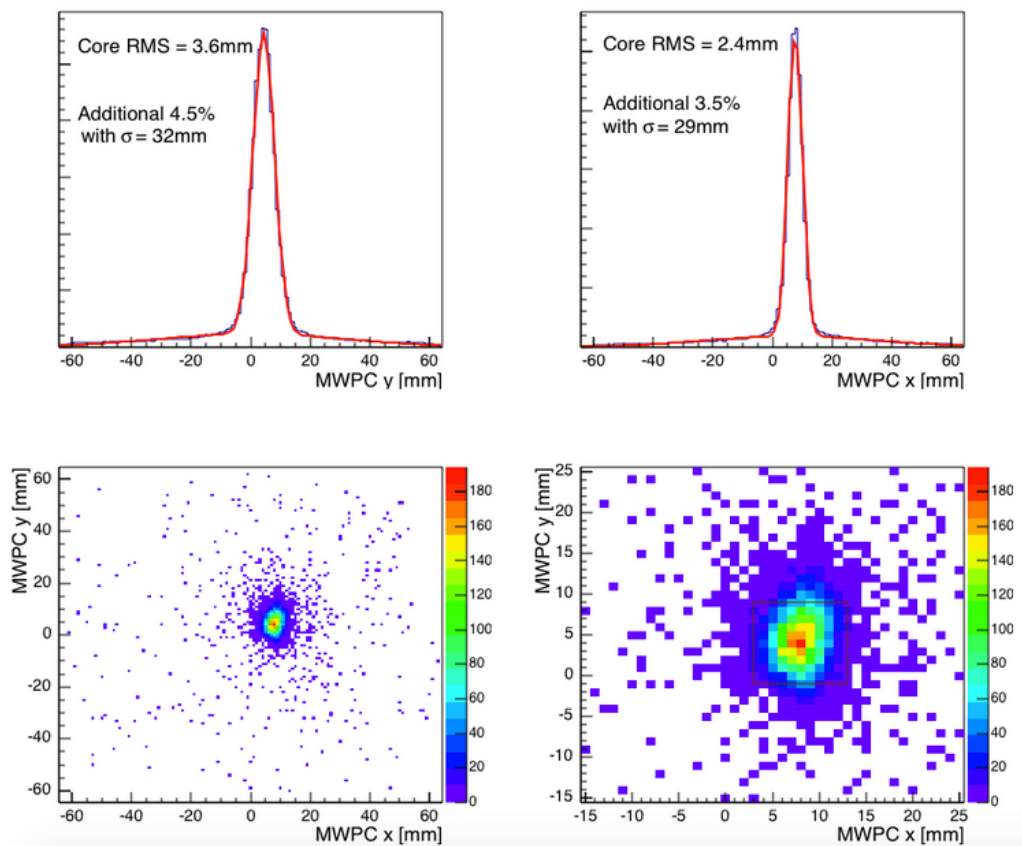


Figure 6.4: Beam profile as measured by the MWPC immediately upstream of the tracker module. The top left and right plots show the fitted vertical (y) and horizontal (x) beam profiles respectively. The bottom left plot shows the spread of the measured hits. The bottom right plot shows a close-up of the central region of the bottom left plot, with a red box indicating the size of the silicon telescope's active area.

6.1.2 Event reconstruction

The general data processing and event reconstruction scheme used for the test beam data is as follows:

- Unpack raw data to form hits
- Reject data when hardware errors are identified
- Reject hits from noisy straws
- Reject hits that are within 150 ns of a previous hit in the same straw
- Group straw hits by time
- Group straw hits spatially into clusters of hits in adjacent straws
- Reconstruct straw hit position
- Reject unphysical straw hit positions
- Reconstruct silicon hit positions
- Fit straight line tracks through the silicon telescope's hits
- Group silicon tracks and straw hits by time

The resulting “event” thus contains a silicon track, a reconstructed straw hit position, and an event time from the silicon telescope's scintillator. The MWPC data is considered separately.

The rejection of hits within a small time window after a previous hit in the same straw is intended to remove hits caused by re-firing in the straws. This refers to events where the electronics are triggered multiple times from the charge deposited by a single traversing particle. This can result from breakdown of the gas due to high gain events. The operating gain of the detector is chosen to minimise this effect, where a suitable gain is found to be $\mathcal{O}(10^5)$ total electrons produced per primary ionisation.

6.2 Test stands

Prototype and production straw modules have been used at a number of test stands. In particular, a test stand was constructed allowing up to three straw modules to be mounted together, orientated vertically to observe cosmic-rays. Scintillators are placed immediately above and below the modules. Coincidences in these scintillators are used to identify events where a charged particle has passed through the tracker modules and provides a time measurement. A photograph of this test stand is shown in figure 6.5.

Radioactive sources such as ^{90}Sr can also be used at this test stand. When a source is placed on the top scintillator, scintillator coincidences can again be used to identify the radioactive decay particles that pass through the full straw system.

A tracker module has also been installed in its real position in the storage-ring and connected to the readout system. This has the advantage of allowing the module to be operated in its design environment, allowing environmental noise to be characterised. The tracker DAQ system has also been integrated with the other detector components and the clock and accelerator control signals distribution system as part of this testing.

6.3 GEANT4 simulation

To support the analysis of the test beam and test stand data and to facilitate the development of reconstruction algorithms, GEANT4 simulations of the test environments have been developed [75]. These simulations utilise the modularity of the full Fermilab muon g-2 experiment GEANT4 simulation to re-use elements such as the geometry of an individual tracker module, with additional elements added as required such as the proton beam spill structure and spatial profile, and the silicon telescope. The simulations are designed such that the output data is in the same format as the real detector data, meaning analysis and reconstruction algorithms can be run without modification on either simulated or real data, allowing a direct comparison of the results. The data processing framework put in place for these test environments served as the prototype from which the full experiment's offline software was developed.


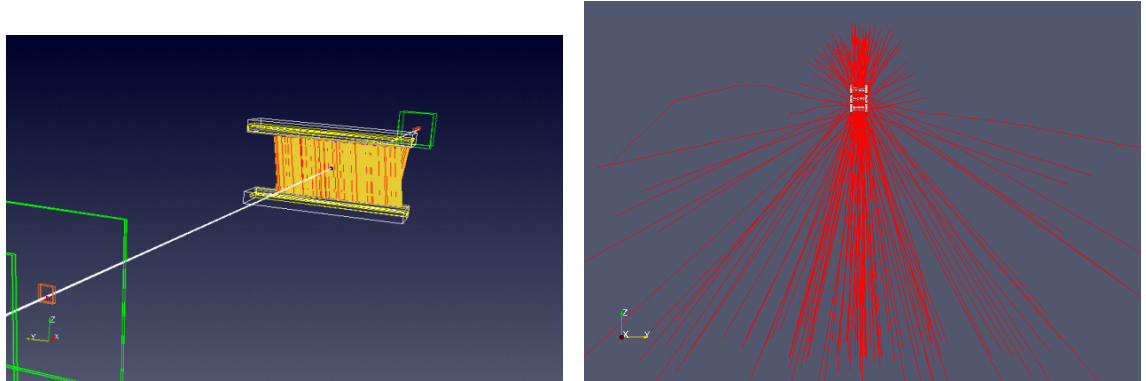


Figure 6.5: A single tracker module mounted in a cosmic-ray test stand. Scintillator paddles clad in black tape can be seen above and below.



(a) GEANT4 simulation of the test beam. The silicon planes are coloured orange, whilst the straws are coloured yellow and red. A particle trajectory is shown in white. The vacuum chamber is not shown.

(b) GEANT4 simulation of the cosmic ray test stand. The red lines are particle tracks, with the straw modules shown in grey.

Figure 6.6

The test beam simulation includes the tracker module and silicon telescope, the vacuum chamber in which the tracker is mounted, and a model of the proton beam. The geometry, materials and detector effects are all modelled. An example simulated event of a particle passing through the silicon telescope and the tracker module straws is shown in figure 6.6 (a). The rendering of this image uses the `ParaView` visualisation software [76]. The MWPCs are not included in the simulation as they are not used in the main event reconstruction analysis.

A simulation of a cosmic-ray test stand is shown in figure 6.6 (b). Between one and three tracker modules are placed in positions and orientations representing the test stand, and a model of the cosmic ray flux is implemented.

6.4 GARFIELD simulation

In addition to GEANT4 simulations, dedicated simulations of the physics in a single straw are performed using the GARFIELD gaseous detector simulation package [57]. This simulation models the properties of the gas in the straws and its interaction with the traversing charged particle, including primary and secondary ionisations, the drifting motion of the liberated charge under the electric and magnetic fields, and the avalanche near the wire. An example event consisting of a single charged

particle traversing the straw is shown in figure 6.7. The induced signal on the straw sense-wire and the response of the readout electronics is also modelled, with examples given in figure 6.8.

GARFIELD simulation results are presented alongside the detector testing results in the following sections.




Figure 6.7: Event display showing a 120 GeV proton traversing a straw as simulated using GARFIELD. The perspective is top-down into the straw. The proton track is shown in red, the locations of primary ionisation events with small black circles, and the path of the drifting primary ionisation electrons is shown in blue. The large black circle represents the straw wall.

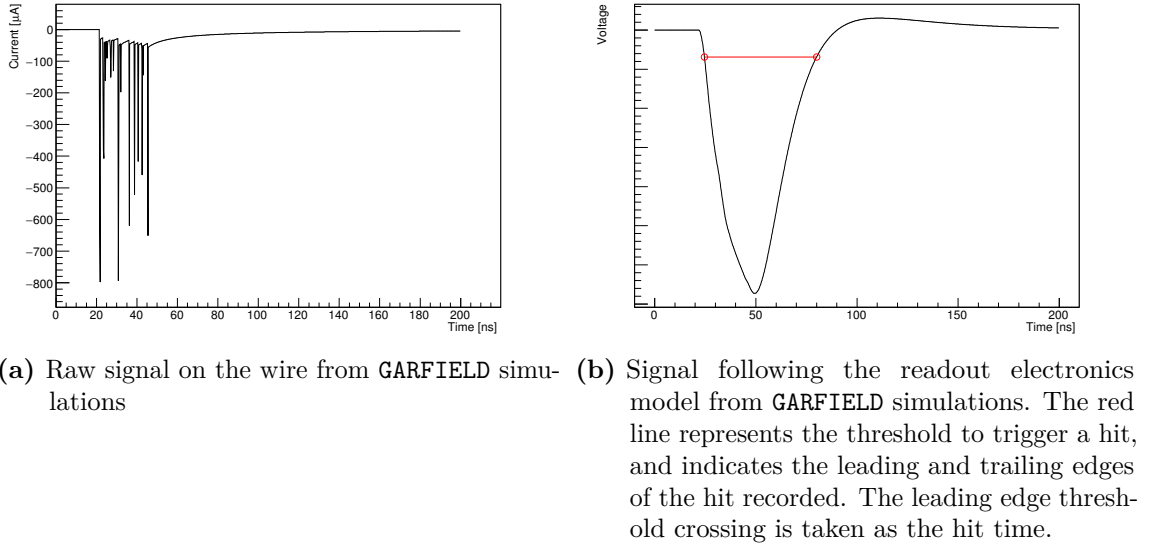


Figure 6.8

6.5 Detector testing results

6.5.1 Straw drift times

As outlined in section 4.4, the relationship between the DCA of a charged particle to the straw wire and the time between the particle passage and the electronics triggering, known as the drift time, is critical to track finding in the straw tracker.

The drift time for a given straw hit is related to the time of the hit as:

$$t_h = t_0 + t_d \quad (6.1)$$

where t_h is the hit time recorded by the straw electronics, t_0 is the time that the charged particle traversed the straw, and t_d is the drift time.

For a full tracker in the Fermilab muon g-2 experiment, t_0 will be determined algorithmically by combining many straw hits for a particular track. For a single straw module with four layers of straws there are not enough hits to accurately perform this procedure. Instead, the t_0 at the test beam and test stands is obtained

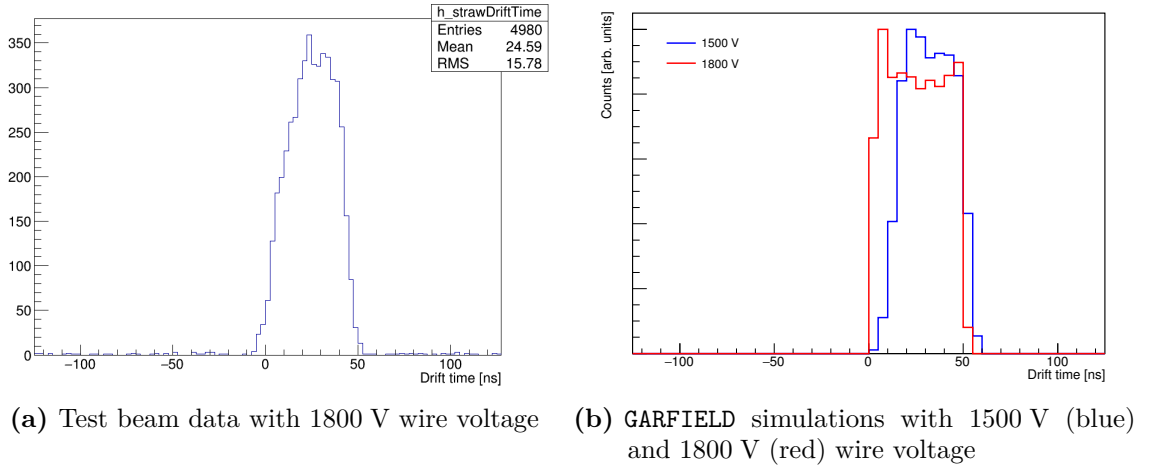



Figure 6.9: Drift times in the straws for the 2015 test beam

using the auxiliary detectors. At the test beam, the time of coincidences in the silicon telescope scintillators is used, and the coincidence time of up and downstream scintillators is also used in the cosmic-ray test stand. These coincidences correspond to the time a particle traversed the system, including the straws. The drift times in the straws are thus determined by subtracting the scintillator t_0 from the straw hit times. The resulting drift time spectrum obtained using this technique at the 2015 test beam is shown in figure 6.9 (a).

The 50:50 argon-ethane gas mixture used in the straws has a saturated drift velocity of $v_d = 50 \mu\text{m}/\text{ns}$ (determined from GARFIELD simulations) that is approximately constant versus radial position within the straw. The drift time is thus expected to linearly increase with track-to-wire DCA, and range from ~ 0 ns for tracks passing near the wire to ~ 50 ns for tracks only just passing within the straw walls. This matches the observed drift time spectrum in figure 6.9 (a) well.

Figure 6.9 (b) shows the drift time spectra predicted from GARFIELD simulations for two different wire voltage values: 1800 V, which is the voltage used at the test beam, and the lower voltage case of 1500 V for comparison. A lower wire voltage results in lower gain, and hence lower charge amplification in the straws (see section 6.5.7 for further details).

The measured drift spectra in 6.9 (a) appears more consistent with the lower-gain GARFIELD simulation in figure 6.9 (b), in particular with respect to the less steep



(a) Test stand data with 1650 V wire voltage (b) GARFIELD simulation with 1650 V wire voltage

Figure 6.10: Drift times in the straws for the cosmic-ray test stand

rising-edge compared to the 1800 V case. As will also be seen in sections 6.5.2, 6.5.4 and 6.5.6, the test beam data consistently favours a lower gain than expected for the 1800 V wire voltage used. The cause of this is believed to be because the measurements were taken after the argon-ethane gas had only been flushing through the detector for 2 hours after a long period of using argon-CO₂, resulting in a contaminated gas mixture.

It is noted also that the test beam data includes additional effects not modelled in the GARFIELD simulation, such as the non-uniform illumination of the straws in the proton beam, and the t_0 measurement resolution of the scintillator. Uneven weighting in the time bins of the TDCs was also discovered after the beam test concluded, resulting in additional smearing of the drift time measurements.

The drift spectrum obtained for cosmic-rays at the test stand is shown in figure 6.10 (a), alongside a GARFIELD simulation of the straw response to cosmic-rays in figure 6.10 (b). The steeper edges of the observed drift time spectrum are consistent with the expectation from higher-gain GARFIELD simulations. The peak at ~ 10 ns seen in the GARFIELD drift time spectrum in 6.10 (b) is not seen in 6.10 (a). This may indicate some discrepancy between the electronics model in GARFIELD and the true performance of the ASDQ, and requires further investigation.

It should be noted that these drift time spectra differ from those expected in the full experiment due to the tracker's location in the fringe of the storage-ring's magnetic

field. This introduces curvature to the path of particle drift, increasing the path length and hence drift times. These measurements serve to broadly validate that the tracker is operating in line with the GARFIELD predictions.

6.5.2 Straw drift velocity

As well as looking at the drift time spectrum, the relationship between drift time and the track-to-wire DCA can be probed using tracks from the silicon telescope at the test beam. The fitted silicon tracks pass through the straws, and the DCA of these tracks to the straw wires can be determined for each straw the track intersects. Figure 6.11 plots the track-to-wire DCA as determined using this technique versus the drift time in that straw.

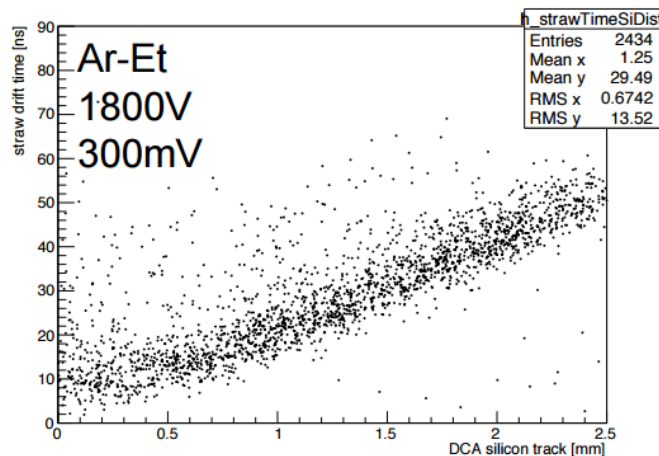


Figure 6.11: Scatter plot of the drift times measured in the straws (y-axis) versus the track-to-wire DCA reconstructed from the silicon track (x-axis). The plot shows test beam data taken with argon-ethane 50:50, a wire voltage of 1800 V, and ASDQ discriminator threshold voltage of 300 mV.

Figure 6.11 indicates that the track-to-wire DCA versus drift time relationship is approximately linear as expected and consistent with a constant drift velocity with respect to radial position in the straw. A linear fit to this data yields a measured drift time of $v_d = 48 \pm 2 \mu\text{m}/\text{ns}$, which is consistent with the GARFIELD prediction of $50 \mu\text{m}/\text{ns}$.

The expected drift time versus track-to-wire DCA relationship from GARFIELD simulations is shown in figure 6.12 for wire voltages of 1500 V and 1800 V, as described




Figure 6.12: Scatter plot of drift times in the straws versus the track-to-wire DCA from GARFIELD simulations of the straw response to the MTest proton beam.

in section 6.5.1. The same linear relationship is observed in simulation and data. The broader distribution of points about this linear trend in the test beam data appears more consistent with the lower gas gain simulations, as was the case for the drift times spectrum. Broadening of this distribution in the test beam data is also expected from the sources of uncertainty in the drift time measurements described in section 6.5.1. Additionally, pile-up in the detectors due to multiple protons within a few ns of each other in the beam can result in the mis-pairing of silicon track-to-wire DCA values and straw drift times from different particles.

One feature that can be observed in both the simulated and test beam data is a larger spread in the drift time values at low track-to-wire DCA, with the spread being asymmetrically above the expected linear trend. This results from the fact that the drift time is related to the DCA of the primary ionisation points along the track, rather than the DCA of the track itself, and the finite number of clusters produced along the track means that there may be a small difference between these two DCA values. Figure 6.13 shows how this effect produces a larger difference between the true and measured DCA values for tracks close to the wire than for those only just passing within the straw walls. The effect is reproduced in GARFIELD simulations as can be seen in 6.14, which shows the mean DCA to the wire for the first, second and third closest clusters produced for each track versus the track-to-wire DCA in the GARFIELD simulation. There is significant scatter about these mean curves, with the DCA of the third closest cluster to the wire being up to 500 μm larger than

the track-to-wire DCA. This corresponds to a scatter in drift time of up to 10 ns, consistent with the scatter seen in 6.12 (b).

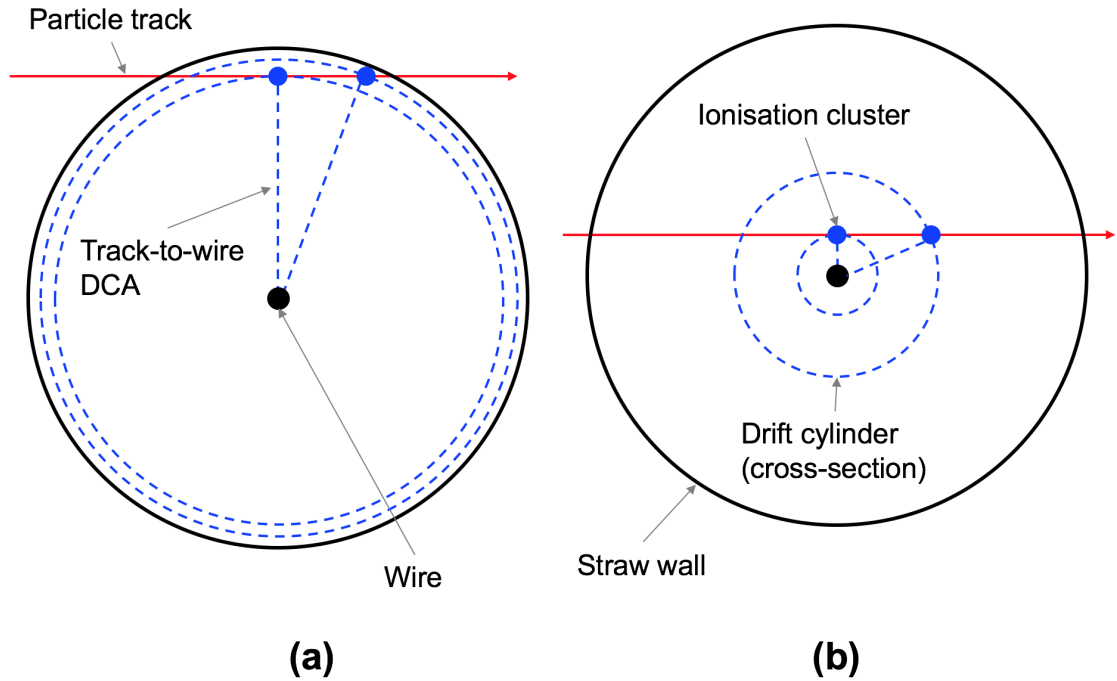


Figure 6.13: Top-down view of a straw cross-section, showing the DCA to the wire for a primary ionisation cluster produced at the track-to-wire DCA, and of a cluster offset from this position. Two cases are shown: in (a) the track is far from the wire, and the difference between the DCA of the two clusters to the wire is small, whilst in (b) the track is close to the wire, and the difference between the DCA of the two clusters to the wire is large. In both cases the separation of the two clusters along the track is the same.

6.5.3 Single module hit reconstruction

A particle traversing the tracker module with normal incidence, as is the case in the test beam, will pass through all four straw layers. Although a single module is insufficient for track formation, it is possible to reconstruct a single hit position using hits in overlapping straws in the four layers. Figure 6.15 (a) gives a schematic of the charged particle traversing the module and the resulting DCA to the straw wires in each layer.

The first step in the hit reconstruction is to subtract the t_0 for the track from the hit times in each straw to obtain the drift times. The t_0 can be obtained either





Figure 6.14: The mean DCA of primary ionisation clusters to the wire in the straws versus the track DCA from GARFIELD simulations of the response on the straws to the MTest proton beam. The simulated wire voltage is 1800 V. The coloured lines represent the closest (black), second closest (red) and third closest (green) cluster to the straw wire liberated by a single traversing particle.



- (a) Diagram of a charged particle traversing the four straw layers in the tracker module at the test beam. The perspective is top-down. The blue line represents the charged particle, and the red lines represent the DCA of the particle to the straw wires.
- (b) Diagram showing the anti-correlation between the DCA to the two wires in a straw doublet for particles normally incident on the straw layers. The particle trajectory is shown in blue. The DCA in each straw is shown by solid red lines, whilst the corresponding drift cylinder (seen as a circle in this top-down perspective) is shown in the red dashed line.

Figure 6.15

using the silicon telescope scintillators, or from the straw data itself. There is an anti-correlation between the track-to-wire DCA and hence drift times for two overlapping straws at the same stereo angle. Overlapping pairs of straws such as these are referred to as doublets. For normally incident particles, the sum of the track-to-wire DCA for the two straws in the doublet equals the pitch between wires in the two layers, p , as can be seen in figure 6.15 (b). Combined with the assumption of a constant drift velocity, t_0 can thus be determined as:

$$t_0 = \frac{1}{2} \left(t_{h1} + t_{h2} - \frac{p}{v_d} \right) \quad (6.2)$$

Once t_0 is determined, the drift time in each straw is calculated using equation 6.1. The track-to-wire DCA in each straw is then determined from the drift time using the constant drift velocity assumption. The DCA in each straw defines a cylinder centred on the wire through which the particle intersected at one, currently unknown, point. The point of closest approach of the particle to the straw wire must be orthogonal to both the wire and the beam direction, but it is unknown on which side of the wires the particle passed. Combining the DCA values for two straws in a doublet removes this ambiguity, as seen in figure 6.16. The radial vector from the wire to the track's point of closest approach is now known, but the height of this point within the straw is not. However, the track also intersects a second doublet in the next two layers, where the straws are at a different stereo angle relative to the first doublet. Combining the two doublets yields a single unique height that satisfies all four DCA values, uniquely defining the 2D hit position of the particle transverse to the plane of the straws, and effectively a 3D position given that the the location of the straw planes in the downstream detection is known. This hit reconstruction process is shown in figure 6.16.

Before hit reconstruction can be performed, suitable events where hits are registered in all four straws in two overlapping doublets are selected in a process shown in figure 6.17. The straws are not perfectly packed in each layer due to limitations in the manifold machining, resulting in a geometric acceptance in each layer of 85%. Additionally, the straws are not 100% efficient at detecting hits from charged particles, particularly near the straw walls. These two effects mean that not all charged particles traversing a module leave hits in all four straw layers.




Figure 6.16: The 2D hit reconstruction process in a single tracker module.

Following event selection, the hit reconstruction is performed, and cuts on the resulting hit positions are made to remove unphysical results. In particular, this serves to remove selected events where not all hits resulted from the same charged particle, for example due to pile-up in the straws from multiple protons.

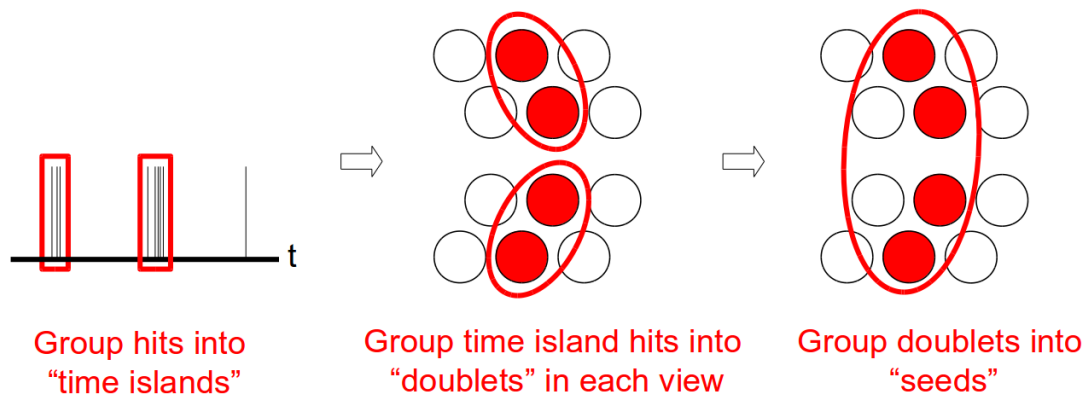


Figure 6.17: The event selection process for hit reconstruction in a single tracker module.

The reconstructed hit positions from the 2015 test beam data are shown in figure 6.18, showing a clear beam spot as expected. The beam is broader in the vertical direction than the horizontal direction, as can be seen in the MWPCS in figure 6.4. Additionally, the hit resolution in the vertical direction in the straws is worse than the horizontal direction owing to the relatively small stereo angle, and this along with reconstruction errors from pile-up also results in the reconstructed beam spot being broader in the vertical direction than the horizontal direction.




Figure 6.18: The reconstructed 2D hit positions in the straw module.

6.5.4 Straw resolution

One of the key straw parameters to determine is the resolution. In a single straw, the resolution is the uncertainty in determining the DCA of the particle track to the straw wire from the measured drift time.

The straw resolution is measured at the test beam by comparing the track-to-wire DCA reconstructed from straw-hit data to the track-to-wire DCA determined from the silicon tracks. The silicon DCA measurement has a resolution of $10\ \mu\text{m}$, which is significantly more accurate than expected for the straws owing to the small strip pitch in the silicon, and the fact that the silicon tracks are reconstructed from four silicon planes, as opposed to a single plane in the tracker module. Figure 6.19 (a) plots these two DCA measurements against each other. Only straw-hits passing the hit-reconstruction event selection are used in order to mitigate against pile-up and noise. The residuals of the straw-reconstructed DCA values to the silicon-reconstructed DCA values indicate the error in measuring the radial position of particles traversing the straws. The resolution of this straw radial hit reconstruction is defined as the standard deviation of a Gaussian function fitted to these residuals, and is found to be $200 \pm 6\ \mu\text{m}$. Both the residuals and the fitted Gaussian are shown in figure 6.19

(b). A consistent measurement is obtained when the resolution is instead defined as the Root-Mean-Square (RMS) spread of the residuals in figure 6.19 (b).


The tracker design specification is a straw resolution of 300 μm , derived from the requirement to measure the beam profile to sub-cm precision so that 10 ppb corrections can be made to the magnetic field sampled by the stored beam [2]. This requirement on the extrapolated vertex position measurement following track reconstruction leads to the required straw resolution quoted, as determined by simulation. The straw resolution measured at the test beam is comfortably within the tracker design specification. Mis-paired straw-hits and silicon tracks, e.g. resulting from different particles, due to pile-up contributes to this measured spread, suggesting that the true resolution is superior to the value measured here.

An alternative technique for estimating the straw resolution using the anti-correlation in the drift times in a straw doublet produces a result consistent with this value.

The straw resolution depends on a number of factors. The collisional motion by which the liberated charges move through the gas introduces a spread in the time taken by charge to drift a given distance in the straw. Additionally, the number of avalanches from primary ionisation events required before the integrated shaped-signal in the ASDQ passes the discriminator threshold and triggers a straw-hit varies, as both the energy of the primary electrons and the electron-multiplication gain of the avalanche are stochastic processes. The geometric difference between the DCA to the straw wire of the track versus the closest primary ionisation cluster also contributes to the uncertainty, as shown in figure 6.13.

The straw resolution can be predicted using `GARFIELD`, with the difference between the reconstructed and true track-to-wire DCA values being shown in figure 6.20. The residuals plotted in figure 6.20 have an asymmetric form, as effects such as the number of avalanches required to trigger the electronics, and the difference between the track-to-wire and cluster-to-wire DCAs, only increase the drift times, never decreasing them, and so the DCA values reconstructed from the drift times will be systematically overestimated. This asymmetry is also seen in the test beam data in figure 6.19 (b).

A number of the effects that determine the resolution are radially dependent within the straw, meaning that the straw resolution is a function of the track-to-wire DCA.



- (a) Scatter plot of the track-to-wire DCA reconstructed in the straws (y-axis) versus that reconstructed by the silicon (x-axis). This data was taken with argon-ethane 50:50 with a wire voltage of 1800 V. This plot consists of approximately 4000 proton tracks.
- (b) Residuals of the track-to-wire DCA measurement made with the straws with respect to the silicon track-to-wire DCA measurement (blue). A Gaussian function (red) is fitted to the data, where the fit parameter $\sigma = 200 \pm 6 \mu\text{m}$ is defined as the straw resolution. This plot consists of approximately 4000 proton tracks.

Figure 6.19

The presented resolution value represents the radially-averaged resolution in the straws.

Taking the RMS spread of the distributions in figure 6.20 as the straw resolutions, the GARFIELD simulations predict straw resolutions of $\sim 200 \mu\text{m}$ for a 1500 V wire voltage and $\sim 45 \mu\text{m}$ for 1800 V. The strong dependence of the straw resolution on the wire voltage results from the exponential increase in gain in the straw with increasing wire voltage [56]. Increasing gain reduces the average number of primary ionisation avalanches that are required to reach the wire before the straw electronics trigger, and thus decreases the spread in the time between the charged particle traversing the straw and the hit time being recorded, which directly determines the straw resolution. As with other test beam data, the resolution measured from figure 6.19 (b) appears consistent with lower-than-expected gas gain, given that the straw wires were operated at 1800 V. If the gas gain can be increased, this suggests that the resolution in the straws can be improved. Additionally, the track-to-wire DCA reconstruction procedure in the straws will be subject to further improvements.




Figure 6.20: Residuals of the reconstructed versus true values of the track-to-wire DCA from GARFIELD simulations of the response of the straws to the MTest proton beam. Wire voltages of 1500 V (blue) and 1800 V (red) are shown.

6.5.5 Reconstruction resolution

As well as the resolution of the track-to-wire DCA measurement within a straw, it is also possible to estimate the resulting reconstructed hit position resolution. This is done by comparing the reconstructed straw-hit positions to the expected position as determined from the silicon track, and plotting the residuals between the two position measurements. This is performed separately in the vertical and horizontal directions, with the results shown in figure 6.21. The statistics are limited in these plots due to the low efficiency of the silicon telescope scintillator trigger and also in the straws (see section 6.5.6), resulting in only a small number of cases where both silicon tracks and reconstructed straw hits can be matched. Nonetheless, the observed distributions of the residuals are consistent with a Gaussian, and are fitted with Gaussian curves shown in red.

As the reconstructed hit positions are determined by the straw geometry and the measured track-to-wire DCA values, the horizontal and vertical reconstructed hit position resolutions are related geometrically to the track-to-wire DCA resolution in a single straw. The horizontal hit position uncertainty, σ_x , is reduced by the straw stereo angle whilst the vertical uncertainty, σ_y , is increased. If the resolution in the straw track-to-wire DCA measurement is σ_r and the straw wire stereo angle from the vertical is θ , then these can be expressed as:

$$\sigma_x = \frac{\sigma_r}{2 \cos \theta}, \quad \sigma_y = \frac{\sigma_r}{2 \sin \theta}. \quad (6.3)$$

For the stereo angle of $\pm 7.5^\circ$, the expected reconstructed hit resolution in the horizontal and vertical directions with respect to σ_r are thus $\sigma_x \approx 0.5\sigma_r$ and $\sigma_y \approx 3.8\sigma_r$, respectively. Given the measured straw DCA resolution $\sigma_r \approx 200 \mu\text{m}$, the reconstructed hit position resolutions are expected to be $\sigma_x \approx 100 \mu\text{m}$ and $\sigma_y \approx 750 \mu\text{m}$. Gaussian curves with these widths are shown in blue in figure 6.21 and agree well with the observed distributions and fitted Gaussians given the low statistics. Note that the height of the blue curves is normalised to the largest histogram bin.

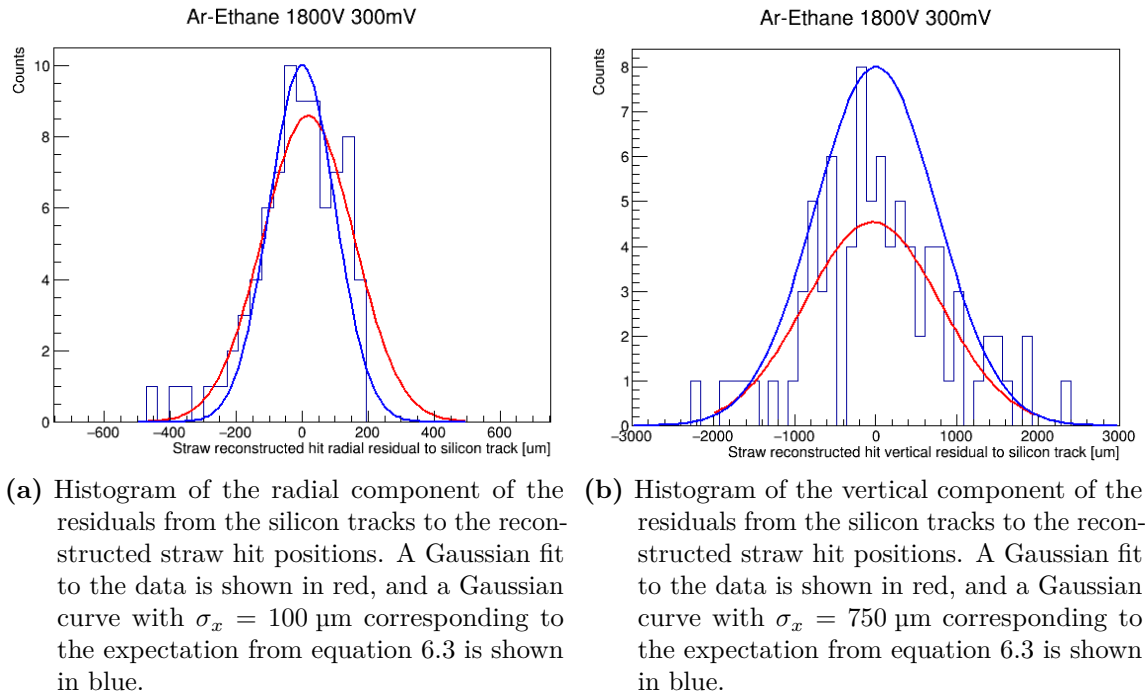
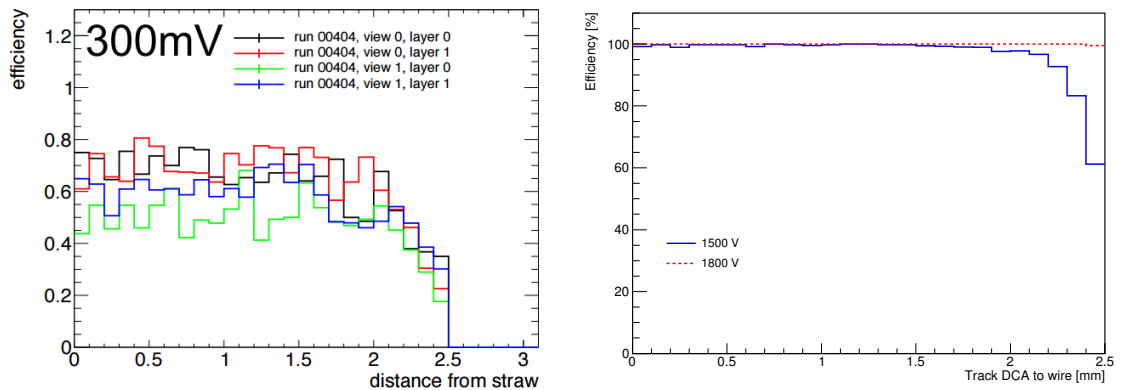


Figure 6.21

6.5.6 Straw efficiency

Efficiency is a measure of the fraction of charged particles traversing a straw that result in a straw-hit being recorded. One method of probing the straw efficiency is to determine which straws the silicon tracks pass through at the test beam and check if those straws recorded a coincident hit. Figure 6.22 (a) shows the efficiency measured using this technique as a function of the DCA of the silicon track to the straw wire. Figure 6.22 (b) shows the predicted straw efficiency at the test beam according to GARFIELD simulations.

The form is consistent between all plots, showing approximately constant efficiency at low radii before falling off at high radii. This fall-off results from the decreasing path length through the gas subtended by tracks with increasing track-to-wire DCAs. The number of primary ionisation clusters produced is proportional to the path length in the gas, and so tracks passing just within the straw walls may not deposit enough charge to trigger the electronics.



- (a) Straw efficiency versus the track-to-wire DCA measured by the silicon at the test beam. The discriminator threshold voltage used is 300 mV. Each coloured curve represents a single straw plane. The straw planes are identified by their “view” number, referring to a pair of adjacent straw planes with the same stereo angle in a straw module, and “layer” number, which discriminates between the two straw planes within a given view.
- (b) GARFIELD simulation of the straw efficiency versus track-to-wire DCA at the test beam, shown for the wire voltages 1500 V (blue) and 1800 V (red)

Figure 6.22

The drop off beyond a track-to-wire DCA of ~ 1.5 mm observed at the test beam is consistent with the lower gain GARFIELD result. However, the overall scale of the efficiency at the test beam is low, significantly below the expected $\sim 100\%$ efficiency for low track-to-wire DCA tracks, and it varies significantly between straw layers.

Given this unexpectedly low straw efficiency at the test beam, further testing was performed using the cosmic-ray test stand. Another measure of straw efficiency is the number of straw layers registering a hit in a module in response to a traversing particle. Ideally all four straw layers are hit. However, as stated earlier the straws are not perfectly packed, resulting in each layer having a straw coverage of approximately 85%. Given this, the expected number of layers in a module registering hits according to simulation if the straws themselves are 100% efficient is shown in figure 6.23. For comparison, the number of straw layers hit for each silicon track at the test beam is shown in figure 6.24 (a). The test beam data shows significantly lower numbers of layers hit on average than would be expected for 100% efficiency. At the cosmic-ray test stand, the number of straw layers hit for a given scintillator coincidence (indicating a particle has traversed the straws) is shown in figure 6.24 (b).




Figure 6.23: Number of straw layers hit for a particle traversing a single straw module in Monte-Carlo simulation, where 100% straw efficiency is assumed.

The test stand data indicates a significantly higher straw efficiency: much closer to the 100% efficient simulation case than the test beam data. It is noted that the test stand scintillators are noisy, and randomly coincident noise in the two scintillators with no cosmic-ray present is considered the likely cause of the significant number of events with no layers hit in the test stand data.

Figure 6.25 shows the ratio of the numbers of particles hitting 4 vs 3 straw layers in a Monte-Carlo simulation, plotted versus the simulated intrinsic straw efficiency. As this ratio varies with the straw efficiency, comparing this ratio for the test beam and test stand data to the trend predicted by simulation can be used to estimate the straw efficiency. This indicates a radially averaged straw efficiency of 57% for the test beam data, consistent with figure 6.22 (a), and 97% for the test stand data.

This test stand straw efficiency measure is consistent with **GARFIELD** predictions, and gives confidence that the straw design is capable of achieving high efficiency. The straw efficiency will be measured as a function of straw radius at the cosmic-ray test stand in the future by fitting a track to the straw-hits in two tracker modules and checking if the straws intersected by the track in a third module register hits. This is a similar analysis to that used to produce figure 6.22 using the silicon telescope at the test beam.




Figure 6.24: Number of straw layers hit for a particle traversing a single straw module




Figure 6.25: The ratio of the number of tracks leaving hits in 3 vs 4 layers of straws from a Monte-Carlo simulation, plotted as a function of the simulated straw efficiency. The value of this ratio is also indicated for the test beam and test stand data.

6.5.7 Optimising wire voltage

The multiplication of the primary electrons liberated in the gas to the eventual number of electrons detected on the straw sense-wire following the avalanches is known as the gas gain [56]. Increasing the wire voltage increases the electric field strength within the straw, which in turn increases the energy gained by electrons between collisions, increasing the chance they will cause further ionisation. Increasing the voltage thus increases the gain, resulting in a larger signal on the wire.

The gain in the straws can be measured by exposing the straws to an ^{55}Fe radioactive source, which is an approximately monochromatic source of 5.9 keV photons [77]. The energy required to liberate an electron in argon is 26 eV [78], so a 5.9 keV photon depositing all its energy in the gas will on average liberate ~ 225 primary electrons. The straw hit rate resulting from exposure to the ^{55}Fe radioactive source can be measured as a function of the ASDQ discriminator threshold. Once the threshold is raised above the noise level, the hit rate will be approximately constant at lower-thresholds as all photons will trigger a hit, but when the threshold is raised to a level that is above the total amplified charge from a single photon then the hit rate drops to zero. The threshold where this drop-off occurs thus indicates the total amplified charge collected by the sense-wire in response to a single 5.9 keV photon, and dividing this measured charge by the initial charge deposited in the form of the primary electrons gives the gas gain.

This measurement was made using the g-2 straws, with the results shown in figure 6.26. The number of primary electrons liberated has a spread about the mean value, as does the gain for a given electron. As such, the signal induced on the wire due to the photons also has a spread, as can be seen in a **GARFIELD** simulation shown in figure 6.27, and so the drop in hit rate is not instantaneous but instead falls across the range 700 – 900mV. The half-height of the falling-edge of the hit rate curve is taken as corresponding to the mean case where 225 primary electrons are deposited. The discriminator threshold is determined by a configurable voltage, and is converted to a charge using data from dedicated characterisation testing of the ASDQs. The fit to the data includes additional contributions to the straw hit rate, including noise, additional ^{55}Fe decay channels, and an argon “escape peak”. This escape peak is the result of when a lower-shell electron in the argon is liberated, and the resulting transition of a higher-shell electron into the empty state results in a photon being

emitted that escapes the straw without liberating further electrons, reducing the amount of primary charge deposited in the straw [77]. The escape peak can be seen in GARFIELD simulations of the straw response to a 5.9 keV photon, as shown in figure 6.27.




Figure 6.26: The straw hit rate when exposed to an ^{55}Fe source, measured across a range of discriminator threshold voltages. The charge corresponding to these discriminator threshold voltages, as calibrated by lab measurements, is shown on the x-axis as the top of the plot in red.

The straw gain measurement was made as described above at three different wire voltages; 1600 V, 1650 V and 1700 V. The relationship between the gain and wire voltage is given by Diethorn's formula [56]:

$$\ln G = \frac{\ln 2}{\ln(b/a)} \frac{V}{\alpha} \ln \frac{V}{\ln(b/a) a \beta}. \quad (6.4)$$

where G is gain, a is wire radius, b is straw inner radius, V is wire voltage, and α, β are constants. The values of α and β are determined by fitting equation 6.4 to the three measured data points, allowing the expected straw gain to be extrapolated across a range of wire voltages. The result of this extrapolation is shown in figure 6.28.

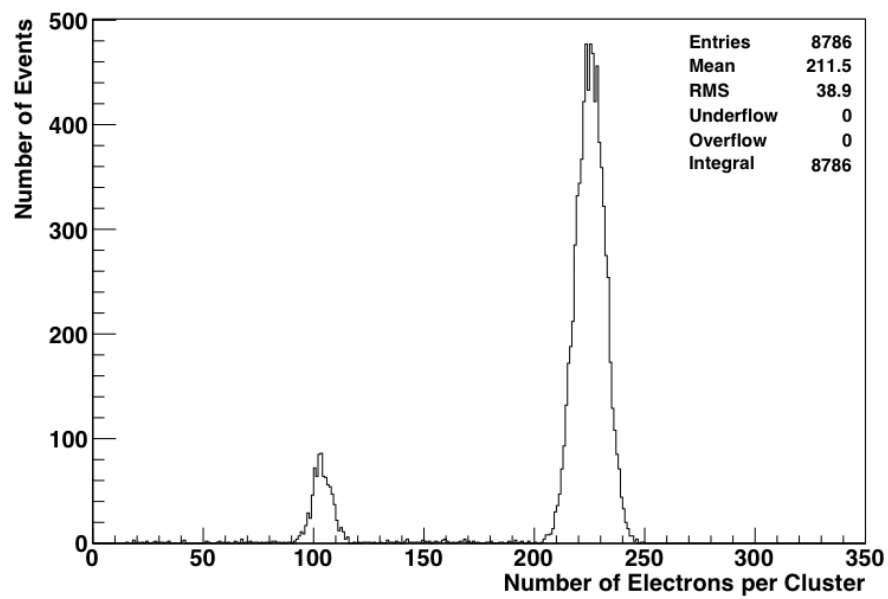


Figure 6.27: Histogram of the number of primary electrons liberated in a straw by a 5.9 keV photon (characteristic of ^{55}Fe radioactive decay) from a GARFIELD simulation. The larger peak is due to the photon depositing all energy in the gas. The smaller peak results from the case where a secondary photon is emitted by an argon ion and escapes the straw, reducing the total amount of energy available to liberate electrons. This is known as an argon escape peak.

The measured gain as a function of wire voltage is used to calibrate the GARFIELD simulations of the straws.




Figure 6.28: Gain in the g-2 straws versus wire voltage, determined by extrapolation from gain measurements using an ^{55}Fe source at wire voltages of 1600 V, 1650 V and 1700 V.

The choice of wire voltage used for the straw tracker is important. If the wire voltage is too low, then the gain will be too low and the small amount of charge liberated by the primary ionisations from the traversing particles will not be sufficiently amplified to induce a large enough signal on the wire to trigger the readout electronics, leading to inefficiency in the straws. If the voltage is too high however then the gain is so high that breakdown of the gas can occur. This can result in dead-time and saturation in the readout electronics. The wire voltage must be selected to balance these effects.

In order to determine the optimum wire voltage, the straws were exposed to a ^{90}Sr beta source, and the rate of hits in an individual straw was measured as a function of the wire voltage. The results are shown in figure 6.29. At low voltages, the gain is too low to trigger straw-hits, but at ~ 1250 V the hit-rate begins to rise rapidly as the gain increases, and the electronics begin to trigger on an increasing fraction of the beta particles. At ~ 1550 V a plateau is reached, which corresponds to approximately 100 % efficiency for a beta particle traversing the straw to trigger a

hit. Beyond ~ 1700 V the hit rate rises again. This is due to the large gain resulting in the breakdown of the gas, which produces multiple hits in the straw electronics for a single traversing charged particle. The optimum wire voltage is thus within the plateau region where efficiency is maximum and gas breakdown is minimised. The resolution of the straws improves with gain, and hence with wire voltage, and so it is desirable to select the highest wire voltage possible within the plateau region, which for figure 6.29 is ~ 1650 V.

Figure 6.29 also shows the expected straw hit rate versus wire voltage from **GARFIELD** simulations for a number of ASDQ discriminator threshold settings. The triggering of multiple hits due to gas breakdown beyond the plateau region is not modelled, and hence this behaviour is not observed in the **GARFIELD** results. The 300 mV curve from simulation agrees well with the data, which was taken using a 300 mV discriminator threshold, giving confidence that the straws are well-modelled by the **GARFIELD** simulation. There are a number of effects that can account for differences between the simulation results and real data, including uncertainty in the straw radius, wire radius, wire offset relative to the centre of the straw, the ratio of argon and ethane in the gas mixture, the purity of the gas, and the temperature. Additionally, the **GARFIELD** simulation does not model the ^{90}Sr beta spectrum but instead generates all electrons at a single energy near the median point of the spectrum. These effects can significantly affect the straw gain, and further studies of their impact will be performed in the future.

Note that the observation of the plateau region in figure 6.29 indicates that maximum efficiency was achieved by the straws at these wire voltages, which supports the findings in section 6.5.6 that the high efficiency in the straws is now being achieved at test stands, and the low efficiency seen in test beam data is confined to that case.

The relationship between the straw hit rate and wire voltage depends on the species and energy of particles traversing the straws, as this affects the amount of primary ionisation charge liberated by the particle track, and correspondingly the size of the signal on the wire for a given gain. Figure 6.30 shows the plateau curves for the straws under the operating conditions at the test beam and test stands, and during the g-2 experiment, determined from **GARFIELD** simulations. The differences between these curves can be used to relate the measured ^{90}Sr results to other scenarios, for example selecting the optimum wire voltage for the muon g-2 experiment physics




Figure 6.29: The hit rate from a ^{90}Sr source in a single straw versus wire voltage, as measured at a test stand using an ASDQ discriminator threshold of 300 mV (black). Also shown are the expected straw hit rates from GARFIELD simulations using ASDQ discriminator thresholds of 200 mV (red), 300 mV (green) and 400 mV (blue).

runs. This comparison will be improved in the future once the response of the straw readout electronics to gas breakdown is included in the GARFIELD simulation.

6.5.8 Optimising ASDQ discriminator threshold

The discriminator threshold voltage in the ASDQ determines the signal size required to trigger a hit. This threshold must be set high enough to minimise hits from electronics noise or low energy secondary particles, but should be low enough to maximise the number of hits from decay e^+ in the g-2 storage-ring.

To determine the discriminator threshold value at which the noise in the system is sufficiently reduced, background data without a particle source is taken using the tracker modules for a range of discriminator thresholds. The results of such a test for a single ASDQ are shown in figure 6.31, which shows the average number of hits recorded in each channel at each threshold value. Each channel is plotted individually as the noise level can vary significantly between channels.




Figure 6.30: The straw hit rate versus the wire voltage from GARFIELD simulations. The results for simulations of cosmic-ray muons (black) and a ^{90}Sr source (green) at the test-stand, 120 GeV protons at the test beam (red), and 1.5 GeV e^+ at the Fermilab muon g-2 experiment (blue) are shown. Note that the simulated results for the Fermilab muon g-2 experiment use the lower discriminator threshold of 200 mV, compared to 300 mV for the other cases. This is because the low-voltage supplies to be used in the final experiment have improved noise characteristics than those used during testing, meaning a lower discriminator threshold can be used.

For all channels the number of hits saturates the TDC buffer at low thresholds, before falling off sharply. Beyond the fall-off, the mean number of hits becomes significantly less than one noise hit per straw per TDC accumulation window, or equivalently per μ^+ fill. The optimal discriminator threshold is as low as possible just beyond this fall-off, in the range 200-300 mV, to minimise noise whilst also minimising the number of particle hits that fail to trigger the electronics.

The data in figure 6.31 was taken using a tracker module installed in the Fermilab muon g-2 experiment storage-ring, but with no source of particles with the exception of background sources such as cosmic rays. Specifically, no stored μ^+ beam was present. Potential sources of environmental noise such as nearby vacuum pumps were running so as to give a realistic picture of the noise environment during physics runs. This scanning process is automated, and can be repeated easily to ensure that the noise in the system has not increased, for example when a module is replaced in the tracker, or the environment around the tracker changes.




Figure 6.31: The number of background hits in each ASDQ channel in the absence of a particle source, measured as a function of the ASDQ discriminator threshold voltage. The curves for each channel are offset vertically for clarity.

6.6 Conclusions

Extensive testing of the tracker modules has successfully demonstrated the ability to interpret the straw-hit data and reconstruct hit positions, and allowed optimised operating parameters to be determined, such as the wire voltage and discriminator threshold. The readout and offline data processing systems have also received extensive testing and development throughout this ongoing process.

The test beam data is consistent with a lower gain than expected for the wire voltages used when compared to other measurements and **GARFIELD** simulations. In particular, this resulted in low efficiency being observed in the straws at the test beam. Subsequent work at test stands however seems to indicate that this was an anomaly, possibly caused by insufficient flushing of the detector gas prior to operation, and test stand data is consistent with high-gain and near 100% straw efficiency.

The straw resolution has been measured to be significantly better than the design requirements, even at the lower gain seen at the 2015 test beam, and it is expected to improve when the straws are operated at higher gain for the Fermilab muon g-2 experiment.

Good agreement between **GARFIELD** simulation and straw data is seen in many cases, giving confidence in the **GARFIELD** straw model developed. This model can be used going forwards to characterise the operation of the straws in the Fermilab muon g-2 experiment, including the effects of the storage-ring magnetic field.

Chapter 7

Simulations of tracker performance

7.1 GEANT4 simulation of the g-2 experiment

A Monte Carlo simulation of the g-2 storage-ring and detector systems has been developed using the `GEANT4` and `art` software frameworks. This simulation provides a broad range of opportunities, such as optimising the placement of detectors and beam control equipment, studying the dynamics of the stored beam and decays to identify and quantify systematic errors in the ω_a measurement, and developing reconstruction and physics analysis algorithms in advance of the first beam arriving in the experiment.

In the simulation, the initial state of the μ^+ in the storage-ring is defined immediately after exiting the inflector magnet. These states are derived and parameterised from dedicated simulations of the upstream beam lines and the inflector magnet itself. The μ^+ are then transported around the storage-ring, taking into account the magnetic field, the electric field from the electrostatic quadrupoles, the kicker magnet during the first orbit, and the collimators. As the μ^+ orbit, the precession of their spins is modelled and tracked. After a number of turns in the ring the μ^+ decay, and the decay phase-space is modelled and the resulting e^+ is tracked as it leaves the storage region and passes through the fringe field into the detector stations.

The detector geometry, detection characteristics and readout are modelled. For the straw tracker, particles interacting with the straw gas produce truth hits recording

their position, momentum and energy deposited in the interaction. The detector model, using parameterised inputs from the test beam and test stand data and GARFIELD simulations, produces corresponding drift times and hit times, and hit parameters are smeared and hits discarded to model the straw resolution and efficiency respectively.

The simulation also features a number of supporting tools, including a persistent record of the truth trajectories of all particles, and a library of all coordinate systems with accompanying transformation tools.

Examples of μ^+ decays in the simulation are given in figures 7.1 and 7.2. In figure 7.1, e^+ can be seen leaving the storage region and either entering the detectors or leaving the ring region altogether. Figure 7.2 has examples of e^+ trajectories passing through the straw tracker and subsequently stopping in the downstream calorimeter.

The following sections present results from tracker performance studies using this simulation.

7.2 Straw and tracker occupancy

Straw occupancy is defined here as the number of hits in a single straw in a single μ^+ fill. It is important for two reasons. Firstly, each TDC can only store a limited number of hits per fill in its buffer, and so it is important to verify that this buffer is larger than the expected number of hits to ensure that the electronics do not saturate and miss data. Secondly, multiple particles crossing the same straw in a small time window cause pile-up, which may not be resolved as distinct particles in the readout electronics.

The mean number of hits in each straw in a single tracker station per fill is shown in figure 7.3. This can be broken down into two distinct components; hits from the decay e^+ , shown in figure 7.4 (a), and hits from secondary particles such as those caused by showering when the decay e^+ passes through material, shown in figure 7.4 (b). For decay e^+ the highest occupancy straws are those at the top of the plot near the storage region, with the occupancy decreasing towards to the lower left of the plot due to the shadow of the upstream calorimeter and the curvature of the tracks. The secondary particles on the other hand are most prevalent in the upstream straws




Figure 7.1: Top-down view of the storage-ring in the g-2 GEANT4 simulation. The design μ^+ orbit is clockwise and shown in green, with e^+ trajectories shown in red. White lines indicate the outline of geometric elements. 24 calorimeters can be seen equally spaced around the ring, as well as three tracker stations. The vacuum chambers are not shown for clarity except at the three tracker stations. All other geometry has been hidden for clarity. The diameter of the ring is 14 m




Figure 7.2: Top-down view of a detector station in the g-2 GEANT4 simulation, using the same colouring scheme described for figure 7.1. Calorimeters can be seen in the far left and right, with eight tracker modules immediately upstream of the right-hand calorimeter. To the left of the tracker modules is an inlet to the vacuum chamber.

to the left of the plot, immediately behind a vacuum chamber inlet that can be seen in figure 7.2 through which the decay e^+ may pass through and shower. There is also a smaller increase in secondaries towards the downstream end of the track to the right of the plot which result from decay e^+ at large vertical angles hitting the straw manifolds above or below the straws and showering. It is advantageous that the secondary hits are concentrated in a region mostly separated from the decay e^+ hits.

The number of hits in the straws and TDCs per fill is shown in figure 7.5, with each TDC being connected to 16 straws. Both the mean and maximum observed numbers of hits in the TDCs are far below the TDC buffer size of 2016 hit words, leaving plenty of margin.

The pile-up of tracks in a station and hits in straws can both be estimated. The simulation shows that pile-up of multiple decay e^+ in the same station takes place for 2.6% of all decay e^+ . The time window for pile-up is defined as 100 ns. Track finding algorithms will need to be able to resolve these piled-up tracks, although many will result from low energy decay e^+ which pass through only a few straws.

Pile-up in individual straws is found to affect 2.8% of hits, and is largely concentrated in regions where hits from secondaries are prevalent as can be seen in figure 7.6.

These occupancy estimates will be revised as the simulation evolves. In particular, GARFIELD simulations will be used to improve the characterisation of the response of the straws to different particle types and energies in the simulation.






Figure 7.3: Straw occupancy heat map showing the mean number of hits per fill in each straw in a single tracker station, where the colour scale is the number of hits. The storage region is just beyond straws at the top of this plot, with μ^+ orbiting from left to right. The downstream calorimeter is just beyond the straws in the right-most layer on the plot. Each of the 8 tracker modules has 4 layers, and the last module before the downstream calorimeter is layers 28-31 in this figure. Each tracker layer has 32 straws.



(a) Mean straw occupancy per fill from decay e^+ hits, where the colour scale is the number of hits. (b) Mean straw occupancy per fill from secondary particle hits, where the colour scale is the number of hits.

Figure 7.4



(a) Histogram of the number of hits in an individual straw during a single μ^+ fill of the storage-ring. The plot includes data from all straws in all three tracker stations, and so averages across the exposure experienced by different straws at different locations. (b) Histogram of the number of hits in an individual TDC during a single μ^+ fill of the storage-ring. The plot includes data from all TDCs in all three tracker stations, and so averages across the exposure experienced by different TDCs at different locations.

Figure 7.5


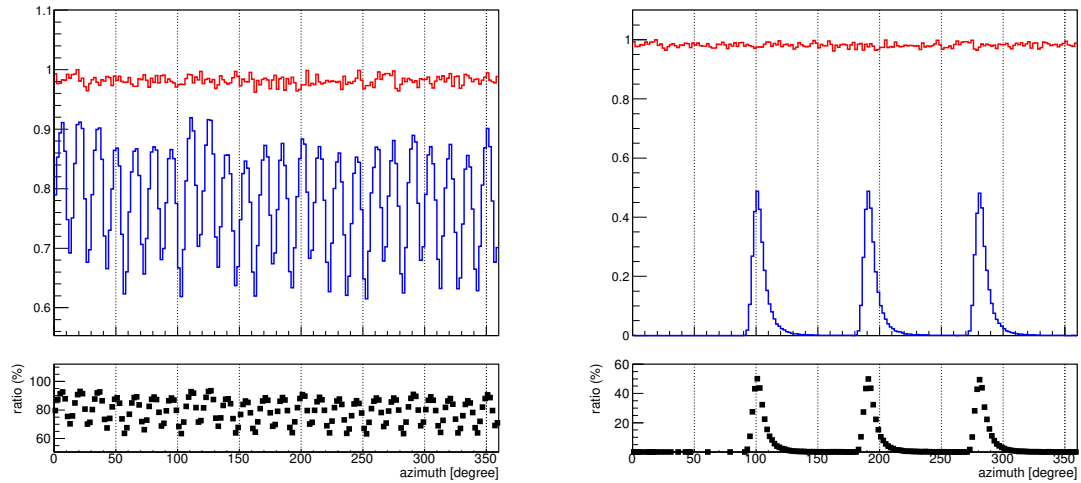


Figure 7.6: Straw heat map showing the mean number of pile-up hits per fill in each straw in a single tracker station, where the colour scale is the number of hits.



(a) Azimuthal acceptance of the 24 calorimeters (b) Azimuthal acceptance of the 3 trackers

Figure 7.7: Detector azimuthal acceptance. In the top panel of each plot, the red curve shows the true μ^+ distribution whilst the blue curve shows the distribution accepted by the detector. The bottom panels show the ratio of accepted to total μ^+ for the detector.

7.3 Tracker acceptance

The tracker acceptance is defined here as the fraction of μ^+ for which the tracker can track the decay e^+ and reconstruct the μ^+ vertex. The acceptance depends on a number of correlated stored μ^+ and decay e^+ parameters, including the azimuthal, radial and vertical decay position of the μ^+ , and the e^+ energy and angle relative to the orbit direction.

The overall azimuthal coverage of the ring by the trackers is relatively low as there are only trackers in 3 of the 24 possible calorimeter locations, as shown in figure 7.7. However, the measured beam profile will be azimuthally averaged and measurements are made at well spaced azimuthal locations. The resulting beam profile is thus not significantly affected beyond reduced statistics, assuming the observed slices of azimuth are representative of the whole ring. For reference, a top-down view of the decay vertices of μ^+ whose decay e^+ are tracked in a single tracker is shown in figure 7.8. The dependence of the tracker acceptance on any parameter has significant contributions from other parameters due to the correlations between them, and azimuthal effects strongly influence all observed tracker acceptance distributions.




Figure 7.8: Heat map showing the decay vertex positions of those μ^+ whose decay e^+ are tracked in a single tracker station located at $x \sim 6800$ mm, $z \sim 100$ mm. The perspective is top-down on the ring.

The acceptance of the calorimeters and trackers with respect to the e^+ energy is shown in figure 7.9. Only e^+ with $E_{e^+} > 1.8$ GeV are counted by the calorimeters for the ω_a measurement, with this energy cut being chosen as a trade-off between maximising the correlation between the e^+ momentum direction and the μ^+ spin vector direction, which increases with e^+ energy, whilst retaining sufficient statistics. The detectors should thus be optimised to maximise the acceptance to decay e^+ with energies above this cut, and figure 7.9 confirms that this is the case.

Ultimately, the goal of the tracker is to reconstruct the stored beam profile of μ^+ whose decay e^+ contribute to the calorimeter ω_a analysis. The beam profile will be measured in the radial and vertical directions, and averaged in azimuth. This profile can then be convoluted with the azimuthally-averaged field map in the ring to determine the effective magnetic field experienced by the orbiting μ^+ . The simulation is used to understand the relative acceptance of the trackers versus the calorimeters, and a translation function must be found to map those μ^+ accepted by the tracker to those accepted by the calorimeters. The design of the trackers maximises the overlap of acceptance phase-space of the two detector systems, minimising the magnitude of this translation.

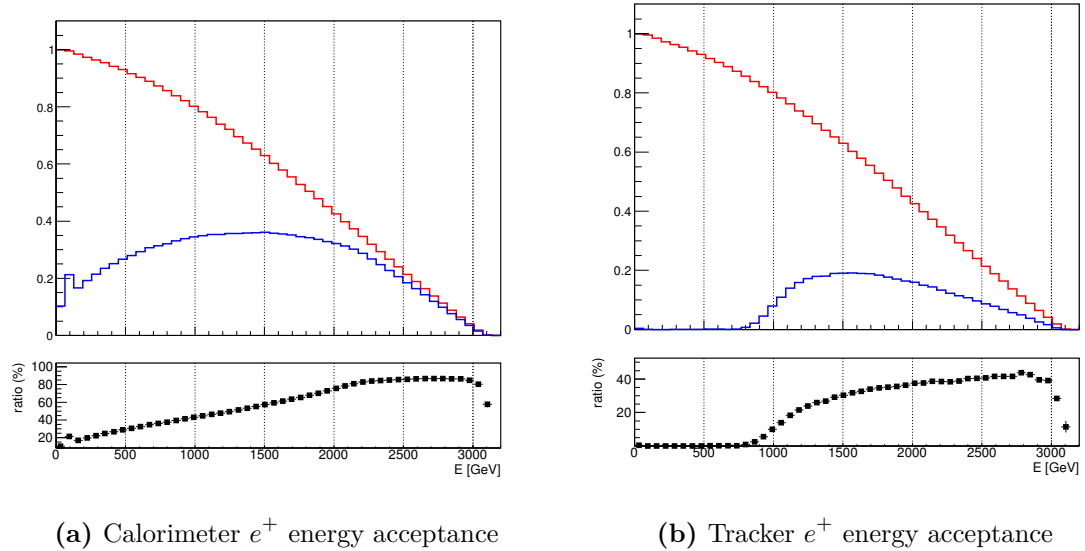


Figure 7.9: Detector e^+ energy acceptance. In the top panel of each plot, the red curve shows the true μ^+ distribution whilst the blue curve shows the distribution accepted by the detector. The bottom panels show the ratio of accepted to total μ^+ for the detector. The distributions are normalised to a single detector station.

The stored μ^+ beam profile in the ring is shown in figure 7.10, as determined from the simulation. The acceptance of the detectors with respect to the μ^+ decay vertex position, both radial, x , and vertical, y , is shown in figures 7.11 and 7.12. The radial position is defined relative to the design orbit radius, whilst the vertical position is designed relative to the mid-plane of the storage region. Both radial and vertical acceptances have a similar form between the calorimeters and trackers, minimising the required translation to the beam profile measured by the tracker, and hence reducing the contribution of any uncertainty in this translation.

The relative radial and vertical acceptances between the calorimeter and the tracker for μ^+ whose decay e^+ pass the energy cut are shown in figure 7.13 (a). The corresponding 2-dimensional ratio of the calo-over-tracker acceptance for these same μ^+ is shown in figure 7.13 (b). This ratio map gives the translation that must be applied to the beam profile measured by the trackers in the Fermilab muon g-2 experiment to determine the beam profile seen by the calorimeters, and hence to correct the ω_a fitted value. This translation will be refined as the simulation continues to develop, and higher statistics simulated datasets are generated.


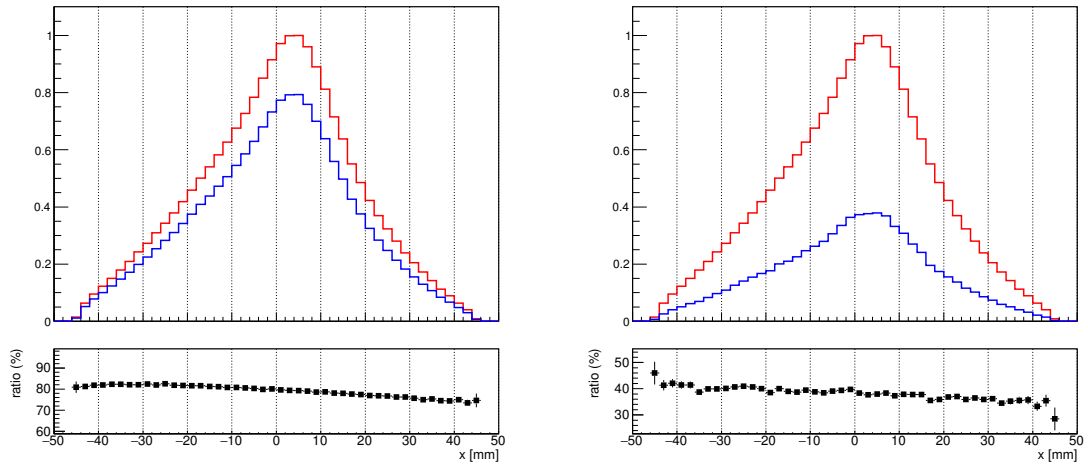



Figure 7.10: The stored μ^+ beam profile in the simulation, as determined from the decay vertices of the μ^+ . The colour scale is the number of μ^+ in each $[x, y]$ bin, normalised to 1.



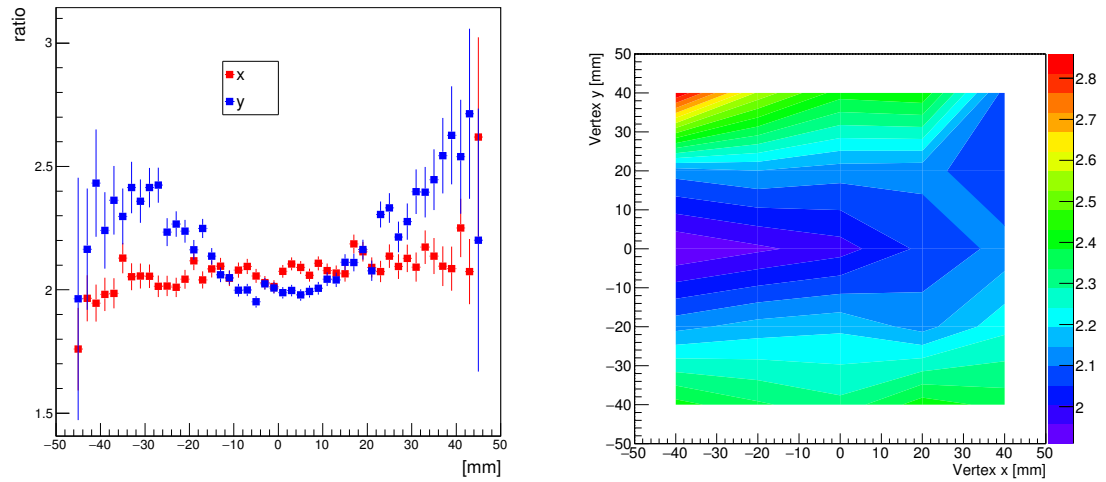
(a) Calorimeter radial decay vertex acceptance (b) Tracker radial decay vertex acceptance

Figure 7.11: Detector radial decay vertex acceptance for e^+ above the energy cut. In the top panel of each plot, the red curve shows the true μ^+ distribution whilst the blue curve shows the distribution accepted by the detector. The bottom panels show the ratio of accepted to total μ^+ for the detector. The distributions are normalised to a single detector station.



(a) Calorimeter vertical decay vertex acceptance (b) Tracker vertical decay vertex acceptance

Figure 7.12: Detector vertical decay vertex acceptance for e^+ above the energy cut. In the top panel of each plot, the red curve shows the true μ^+ distribution whilst the blue curve shows the distribution accepted by the detector. The bottom panels show the ratio of accepted to total μ^+ for the detector. The distributions are normalised to a single detector station.



(a) Calorimeter versus tracker relative radial (red) and vertical (blue) decay vertex acceptance. The error bars show the statistical uncertainty in these distributions, which increases away from the centre of the storage region due to the fall off in numbers of μ^+ orbiting in these fringe regions.

(b) Tracker-to-calorimeter beam profile translation map, where the colour scale shows the magnitude of the translation to be applied to the μ^+ counts measured by the tracker at a given $[x, y]$ coordinate.

Figure 7.13

The total number of μ^+ expected to be stored during the Fermilab muon g-2 experiment is 1.5×10^{12} [2]. Using the tracker acceptance determined from the simulation, the fraction of μ^+ producing e^+ that can be tracked is estimated to be 4.8%, resulting in 1.3×10^{10} tracks in total for the experiment. This number can be used as the basis for determining the statistical precision of the measurements and the systematic uncertainties whose determination are predominantly statistically limited.

7.4 Tracker and calorimeter e^+ matching

The placement of each straw tracker immediately upstream of a calorimeter allows tracks for particles passing through the straws to be extrapolated forwards into the calorimeter and matched to reconstructed calorimeter hits. This has a number of uses, including resolving pile-up in the calorimeters using straw tracks, comparing reconstructed calorimeter hit energy with track momentum, using the calorimeter hit times as an independent t_0 in straw track finding, and relative alignment of the two detectors.

The fraction of calorimeter e^+ hits surviving the $E_{e^+} > 1.8$ GeV cut for which a track can be reconstructed in the immediately upstream straw tracker, plotted against the hit position on the calorimeter front-face, is shown in figure 7.14. There is a strong bias towards the radially inner side of the calorimeter, with e^+ hitting the radially outer side of the calorimeter front-face near the storage region not having traversed enough tracker planes for track fitting to be reliably performed. In total, 41% of calorimeter hits that pass the energy cut have matching tracks (e.g. those used for the ω_a measurement), and 18% of all calorimeters hits have matching tracks.

Estimates of pile-up in the calorimeter in this measurable region can be extrapolated across the full face of the calorimeter, and to all calorimeters including those without upstream trackers. Additionally, hits in the final few straw layers can be used to form simplified tracks, known as track-stubs, which can be used to discern calorimeter pile-up across a greater fraction of the calorimeter front-face. Use of a track stub formed from hits in last two modules only for example would allow up to 69% of the total calorimeter hits above the energy cut to be matched to tracks or stubs, and 48% of hits from all decay e^+ .




Figure 7.14: Probability of a calorimeter e^+ hit having a matching track in the upstream tracker, plotted as a function of hit position on the calorimeter front-face. The colour scale is the probability of a hit at that position. The view is downstream into the calorimeter, with the storage region on the left.

Chapter 8

Outlook

Following the successful tracker test campaigns described in this thesis, including a number of beam tests, production of the tracker modules is now nearing completion. Several tracker modules are installed in the Fermilab muon g-2 experiment storage-ring, with the rest under test at Fermilab and the University of Liverpool. These tests have confirmed that the straw tracker can operate at the required wire voltage for optimum performance, that the noise can be reduced to less than one hit per straw per μ^+ fill, and that the design resolution of 300 μm has been achieved and even surpassed. It has also been demonstrated that the drift time of liberated charge in the straws can be successfully interpreted from the time recorded by the triggering of the straw readout electronics, and can be used to reconstruct 3D positions of traversing charged particles.

As part of the commissioning of tracker modules, they will be placed in a test frame with a motorised C-shaped arm featuring a radioactive source and a scintillator at opposing sides of the straws. This is used to scan all modules in fine steps, mapping the wire positions and calibrating the straw uniformity. During installation, the tracker modules are integrated with the low- and high-voltage supplies, the gas delivery system, the manifold water cooling and FLOBBER air cooling systems.

Installation of the detectors in the g-2 storage-ring will continue through the first quarter of 2017, in parallel with the installation of the full tracker DAQ system and its integration with the experiment-wide DAQ and clock and control system.

In parallel to the construction and testing efforts, the offline data processing will continue to be developed. In particular, the track-finding and extrapolation algo-

gorithms will be developed and extensively tested using Monte Carlo simulation data. The simulations will also be used to quantify systematic effects in the trackers and throughout the experiment, including the pitch and electric field corrections.

The first μ^+ beam will arrive in mid-2017, allowing commissioning and calibration of the trackers, with a subsequent period for the analysis of commissioning data and the refinement of algorithms.

Once the experimental subsystems are fully commissioned, two nine-month physics runs will be performed between 2017 and 2019. These will collect ten times the statistics of the BNL a_μ measurement in total, yielding new world's best measurements of a_μ and the muon EDM by the end of the decade, and determining whether or not the BNL a_μ measurement was the harbinger of BSM physics.

Bibliography

- [1] E821 Muon (g-2) Collaboration, G. W. Bennett *et al.*, Phys. Rev. Lett. 92 (2004).
- [2] E989 Muon g-2 Collaboration, J. Grange *et al.*, arXiv:1501.06858 (2015).
- [3] F. Jegerlehner, *The Anomalous Magnetic Moment of the Muon*, Springer Tracts in Modern Physics (Springer, 2008).
- [4] F. Jegerlehner and A. Nyffeler, Physics Reports 477 (2009).
- [5] J. P. Miller, E. de Rafael, and B. L. Roberts, Reports on Progress in Physics 70 (2007).
- [6] R. Frisch and O. Stern, Zeits. Physik 85 (1933).
- [7] I. Etermann and O. Stern, Zeits. Physik 85 (1933).
- [8] I. Etermann and O. Stern, Phys. Rev. 45 (1934).
- [9] I. I. Rabi, J. M. B. Kellogg, and J. R. Zacharias, Phys. Rev. 46 (1934).
- [10] P. Kusch and H. M. Foley, Phys. Rev. 74 (1948).
- [11] D. Hanneke, S. Fogwell, and G. Gabrielse, Phys. Rev. Lett. 100 (2008).
- [12] Y. Ueno *et al.*, Hyperfine Interactions 238 (2017).
- [13] S. Eidelman and M. Passera, Modern Physics Letters A 22 (2007).
- [14] The DELPI Collaboration, J. Abdallah *et al.*, The European Physical Journal C 35 (2004).
- [15] J-PARC New g-2/EDM experiment collaboration, H. Iinuma *et al.*, Journal of Physics: Conference Series 295 (2011).

- [16] T. Mibe, Measurement of muon $g-2$ and EDM with ultra-cold muon beam at J-PARC, 12th International Conference on Low Energy Antiproton Physics (2016).
- [17] H. Nunokawa, S. Parke, and J. W. Valle, Progress in Particle and Nuclear Physics 60 (2008).
- [18] E821 Muon ($g-2$) Collaboration, G. W. Bennett *et al.*, Phys. Rev. D 80 (2009).
- [19] Particle Data Group, K. A. Olive *et al.*, Chin. Phys. C38 (2014).
- [20] T. Kinoshita and M. Nio, Phys. Rev. D 70 (2004).
- [21] T. Aoyama, M. Hayakawa, T. Kinoshita, and M. Nio, Phys. Rev. Lett. 109 (2012).
- [22] T. Aoyama, M. Hayakawa, T. Kinoshita, and M. Nio, Phys. Rev. Lett. 109 (2012).
- [23] C. Gnendiger, D. Stöckinger, and H. Stöckinger-Kim, Phys. Rev. D 88 (2013).
- [24] K. Hagiwara, R. Liao, A. D. Martin, D. Nomura, and T. Teubner, Journal of Physics G: Nuclear and Particle Physics 38 (2011).
- [25] M. Davier *et al.*, The European Physical Journal C 66 (2010).
- [26] B. Chakraborty *et al.*, Phys. Rev. D 93 (2016).
- [27] J. Prades, E. de Rafael, and A. Vainshtein, *The Hadronic Light-by-Light Scattering Contribution to the Muon and Electron Anomalous Magnetic Moments*, (World Scientific, 2011).
- [28] F. Jegerlehner, EPJ Web of Conferences 118 (2016).
- [29] T. Blum, The muon $g-2$ from lattice QCD, 22nd International Spin Symposium (2016).
- [30] B. Shwartz, Nuclear Physics B - Proceedings Supplements 253 (2014).
- [31] A. Denig, Recent Results on Light Meson Decays at BES III, 12th Conference on Quark Confinement and the Hadron Spectrum (2016).
- [32] A. Nyffeler, Hadronic light-by-light scattering in the muon $g-2$: impact of

proposed measurements of the $\pi^0 \rightarrow \gamma\gamma$ decay width and the $\gamma^*\gamma \rightarrow \pi^0$ transition form factor with the KLOE-2 experiment., 7th International Workshop on Chiral Dynamics (2012).

- [33] D. K. Hong and D. H. Kim, *Physics Letters B* 758 (2016).
- [34] T. Moroi, *Phys. Rev. D* 53 (1996).
- [35] D. Stöckinger, *Journal of Physics G: Nuclear and Particle Physics* 34 (2007).
- [36] A. Czarnecki and W. J. Marciano, *Phys. Rev. D* 64 (2001).
- [37] J. H. Heo, *Phys. Rev. D* 80 (2009).
- [38] M. Krawczyk, arXiv:hep-ph/0103223 (2001).
- [39] A. Cherchiglia, D. Stöckinger, and H. Stöckinger-Kim, arXiv:1711.11567 (2017).
- [40] M. Pospelov, *Phys. Rev. D* 80 (2009).
- [41] H. Davoudiasl, H.-S. Lee, and W. J. Marciano, *Phys. Rev. Lett.* 109 (2012).
- [42] F. Combley, F. J. M. Farley, and E. Picasso, *Physics Reports* 68 (1981).
- [43] T. Kibble and F. Berkshire, *Classical Mechanics*, (Imperial College Press, 2004).
- [44] L. H. Thomas, *Nature* 117 (1926).
- [45] S. Nagaitsev, Fermilab antiproton source, recycler ring, and main injector, 2014, arXiv:1408.0759.
- [46] D. Stratakis *et al.*, Fermilab Report No. FERMILAB-CONF-17-176-AD, 2016 (unpublished).
- [47] T. D. Lee and C. N. Yang, *Phys. Rev.* 104 (1956).
- [48] C. Joseph and J. Soffer, *High-Energy Physics with Polarized Beams and Polarized Targets*, Proceedings of the 1980 International Symposium, Lausanne, September 25 to October 1, 1980 (Springer, 1980).
- [49] A. Fienberg *et al.*, *Nucl. Instrum. Meth.* A783 (2015).
- [50] A. Anastasi *et al.*, *Nucl. Instrum. Meth.* A788 (2015).
- [51] G. T. Danby *et al.*, *Nucl. Instrum. Meth.* A457 (2001).

- [52] R. Prigl, U. Haebleren, K. Jungmann, G. zu Putlitz, and P. von Walter, Nucl. Instrum. Meth. A374 (1996).
- [53] X. Fei, V. W. Hughes, and R. Prigl, Nucl. Instrum. Meth. A394 (1997).
- [54] W. Liu *et al.*, Phys. Rev. Lett. 82 (1999).
- [55] P. J. Mohr, D. B. Newell, and B. N. Taylor, Rev. Mod. Phys. 88 (2016).
- [56] W. Blum *et al.*, *Particle Detection with Drift Chambers*, (Springer-Verlag Berlin Heidelberg, 2008).
- [57] R. Veenhof, Garfield - simulation of gaseous detectors, <http://garfield.web.cern.ch/garfield/>.
- [58] T. Affolder *et al.*, Nucl. Instrum. Meth. A526 (2004).
- [59] W. Bokhari *et al.*, The ASDQ ASIC for the Front End electronics of the COT, https://www.researchgate.net/publication/265995498_The_ASdq_ASIC_for_the_Front_End_electronics_of_the_COT.
- [60] J. Wu and Z. Shi, C5 IEEE, www-ppd.fnal.gov/EEDOffice-W/Projects/ckm/comadc/C5_IEEE.doc.
- [61] A. X. Widmer and P. A. Franaszek, IBM J. Res. Dev. 27 (1983).
- [62] M. Pesaresi and P. Vichoudis, FC7: AMC FMC carrier for CMS, https://indico.cern.ch/event/300897/contributions/1664696/attachments/567234/781330/fc7_2014.03.17_xTCA_IG.pdf, 2014.
- [63] AMC13 documentation, <http://www.amc13.info>.
- [64] MicroTCA Overview, http://www.vadatech.com/media/article_MicroTCA_Overview.pdf, 2014.
- [65] MIDAS, https://midas.triumf.ca/MidasWiki/index.php/Main_Page.
- [66] C. G. Larrea *et al.*, Journal of Instrumentation 10 (2015).
- [67] Boost.Python, http://www.boost.org/doc/libs/1_62_0/libs/python/doc/html/index.html.
- [68] What is PyQt, <https://riverbankcomputing.com/software/pyqt/intro>.

-
- [69] C. Green *et al.*, J. Phys. Conf. Ser. 396 (2012).
- [70] ZeroMQ Distributed Messaging, <http://zeromq.org/>.
- [71] Node.js, <https://nodejs.org/en/>.
- [72] Plotly, <https://plot.ly/>.
- [73] Flask (A Python Microframework), <http://flask.pocoo.org/>.
- [74] Fermilab Test Beam Facility, <http://ftbf.fnal.gov/>.
- [75] S. Agostinelli *et al.*, Nucl. Instrum. Meth. A506 (2003).
- [76] ParaView, <http://www.paraview.org/>.
- [77] Z. Ahmed *et al.*, Journal of Instrumentation 9 (2014).
- [78] G. Alkhazov, A. Komar, and A. Vorob'ev, Nucl. Instrum. Meth. A48 (1967).