A direct measurement of $P_\mu^\pi \xi$ from muon decay

by

James Bueno

M.Phys., The University of Oxford, 2005

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

The Faculty of Graduate Studies

(Physics)

The University of British Columbia

(Vancouver)

April, 2010

Abstract

The asymmetry of the positron distribution in polarised muon decay provides confirmation that the weak interaction maximally violates parity conservation. Since 1957 the quantity $P_\mu \xi$ has been measured with increasing precision, where P_μ is the polarisation of the muon, and ξ is a parameter describing the asymmetry. Thus far the results have been consistent with the standard model using a $(V - A)$ interaction.

A new measurement of $P_\mu^\pi \xi$ using the TRIUMF Weak Interaction Symmetry Test (TWIST) spectrometer is presented in this thesis. The result is $P_\mu^\pi \xi = 1.00084 \pm 0.00035\,(\text{stat.})^{+0.00165}_{-0.00063}\,(\text{syst.})$, which is a factor of 3.2 more precise than a previous TWIST direct measurement, and a factor of 7.0 more precise than the pre-TWIST value. New limits are set on physics beyond the standard model, including the weak decay of right-handed muons, and left-right symmetric models where a $(V + A)$ current is introduced to conserve parity at higher energies.

Table of Contents

Abstract ii

Table of Contents iii

List of Tables ix

List of Figures x

Acknowledgements xiv

1 Introduction 1
- 1.1 The standard model of particle physics 2
- 1.2 Spin and polarisation 4
- 1.3 Muon production 5
- 1.4 Muon decay 7
  - 1.4.1 Decay modes 7
  - 1.4.2 Matrix element 7
  - 1.4.3 Muon decay parameters 9
  - 1.4.4 Differential decay rate 10
  - 1.4.5 Theoretical spectrum and radiative corrections 10
- 1.5 Standard model extensions 11
  - 1.5.1 Right-handed muons 15
  - 1.5.2 Left-right symmetric electroweak models 16
- 1.6 Muon depolarisation mechanisms 17
  - 1.6.1 Depolarisation in a magnetic field 18
  - 1.6.2 Depolarisation while decelerating 19
  - 1.6.3 Depolarisation after thermalisation 20
- 1.7 Previous $P_\mu^\pi \xi$ measurements 24
  - 1.7.1 Measurements of $P_\mu^\pi \xi$ 24

1.7.2 Measurements of $P_\mu^K \xi$ . . . 26

2 Apparatus . . . 28
2.1 Overview . . . 28
2.2 Muon production and delivery . . . 28
2.2.1 Location . . . 28
2.2.2 Muon production . . . 28
2.2.3 M13 beam line . . . 29
2.2.4 Beam line momentum calibration . . . 30
2.2.5 Particles delivered . . . 32
2.2.6 M13 configurations . . . 32
2.3 Measuring the muon beam . . . 32
2.4 Quadrupole steering . . . 36
2.5 The detector . . . 37
2.6 Coordinate system . . . 38
2.7 Solenoid magnet . . . 39
2.8 Wire chambers . . . 41
2.9 The upstream "beam package" . . . 44
2.10 Muon stopping target . . . 45
2.11 Muon ranging . . . 47
2.12 Electronics . . . 47
2.13 Alignments . . . 48

3 Analysis . . . 51
3.1 Overview . . . 51
3.2 Track reconstruction . . . 53
3.2.1 "Unpacking" . . . 53
3.2.2 Crosstalk removal . . . 55
3.2.3 Windowing . . . 55
3.2.4 Classification . . . 55
3.2.5 Pattern recognition . . . 56
3.2.6 Fitting the helical trajectory of the decay positron . . . 59
3.2.7 Improved drift chamber space-time-relationship . . . 61
3.2.8 Improved resolution function . . . 63
3.3 Track selection and cuts . . . 65

3.3.1 Trigger particle time-of-flight . . . 66
3.3.2 Event classification cut . . . 66
3.3.3 Muon cuts . . . 67
3.3.4 Decay particle trajectories . . . 68
3.3.5 Kinematic fiducial region . . . 71
3.4 Extraction of muon decay parameters . . . 72
3.5 Energy calibration . . . 77
3.6 Time dependence of depolarisation . . . 78
3.7 Muon beam spots . . . 81

4 Simulation . . . 84
4.1 Introduction . . . 84
4.2 Overview of the simulation . . . 84
4.3 Primary particle generation . . . 85
4.4 Depolarisation . . . 86
4.4.1 Initial polarisation . . . 86
4.4.2 Electromagnetic field . . . 88
4.4.3 Time-dependent depolarisation . . . 89
4.4.4 Other depolarisation . . . 89
4.5 Muon decay . . . 89
4.6 Wire chamber response . . . 90

5 Data . . . 91
5.1 Overview . . . 91
5.2 Muon beam tuning . . . 91
5.3 Data sets . . . 92
5.4 Data quality checks . . . 95
5.4.1 Data acquisition system . . . 95
5.4.2 Chamber signals . . . 95
5.4.3 Rates . . . 96
5.4.4 Beam line stability . . . 96
5.4.5 Chamber foil bulging . . . 97
5.4.6 Muon stopping distribution . . . 97
5.4.7 Muon beam stability . . . 97

6 Systematic Uncertainties and Corrections . . . 100
6.1 Introduction . . . 100
6.2 Polarisation . . . 103
6.2.1 Overview . . . 103
6.2.2 Measures of polarisation . . . 104
6.2.3 Initial beam position and angle alignment . . . 104
6.2.4 Magnetic field position and angle alignment . . . 107
6.2.5 Alignment uncertainties from beam and magnetic field . . . 107
6.2.6 Overview of the magnetic field map . . . 110
6.2.7 Muon beam tunes for evaluating the uncertainty . . . 113
6.2.8 Muon beam angular distribution width . . . 119
6.2.9 Summary of muon beam and fringe field uncertainties . . . 122
6.2.10 Stopping material . . . 123
6.2.11 Muon production target . . . 130
6.2.12 Background muon contamination . . . 131
6.3 Chamber response . . . 136
6.3.1 Drift chamber space-time relationship . . . 136
6.3.2 Drift chamber geometric effects . . . 138
6.3.3 Upstream-downstream efficiency . . . 138
6.3.4 Crosstalk . . . 139
6.3.5 Wire time offsets . . . 139
6.3.6 Negligible chamber response uncertainties . . . 140
6.4 Detector alignment . . . 141
6.5 Positron interactions . . . 142
6.5.1 Bremsstrahlung and δ-electron rates . . . 142
6.5.2 Outside material . . . 146
6.6 Resolution . . . 148
6.7 Momentum calibration . . . 151
6.7.1 Magnetic field shape . . . 151
6.7.2 Use of the kinematic endpoint . . . 152
6.8 Beam stability . . . 154
6.8.1 Muon beam intensity . . . 154
6.9 External . . . 156
6.9.1 Radiative corrections . . . 156
6.9.2 Correlation with η . . . 156

7 Results . . . 157
7.1 Blind results . . . 157
7.2 Result revealed . . . 161
7.3 "White box" consistency test . . . 161
7.4 Physics implications . . . 162
7.4.1 Global analysis of muon decay data . . . 162
7.4.2 Left-right symmetric models . . . 162
7.5 Future experiments . . . 165
7.5.1 Statistical uncertainty . . . 165
7.5.2 Magnetic field map uncertainty . . . 165
7.5.3 Stopping material depolarisation uncertainty . . . 167
7.5.4 Other uncertainties . . . 167
7.6 Conclusions . . . 170

Bibliography . . . 171

Appendices

A Personal contributions . . . 181

B History of the muon . . . 183

C Naming of the muon . . . 185

D Magnetic field mapping and simulation . . . 188

E New downstream trigger . . . 191

F Bias in $P_\mu^\pi \xi$ extraction technique . . . 193

G Time expansion chambers analysis . . . 194
G.1 Introduction . . . 194
G.2 Typical raw events . . . 194
G.3 Analysis . . . 196
G.3.1 Unpacking and wire time offsets . . . 196
G.3.2 Reject multiple trigger events . . . 196
G.3.3 Discriminator amplitude walk . . . 197

G.3.4 Track candidates . . . 197
G.3.5 Final multiple track removal . . . 199
G.3.6 Other algorithms . . . 200
G.4 Sense plane calibration . . . 201
G.4.1 Wire time offsets . . . 201
G.4.2 Discriminator amplitude walk . . . 202
G.4.3 Space-time relationship . . . 203
G.5 Performance . . . 204
G.5.1 Resolution . . . 204
G.5.2 Number of hits per event . . . 205
G.6 Summary . . . 205

H E1111: Subsidiary μ^+SR experiment . . . 207
H.1 Introduction . . . 207
H.2 Theory . . . 207
H.3 μ^+SR Apparatus . . . 209
H.3.1 M20 Beam line . . . 209
H.3.2 Detector . . . 209
H.3.3 Sample preparation . . . 212
H.4 Data . . . 212
H.5 Analysis . . . 213
H.5.1 Rebinning . . . 213
H.5.2 Calibration using spin glass . . . 214
H.5.3 Metal samples . . . 218
H.6 Correction due to scintillator stops . . . 224
H.7 Other systematic uncertainties . . . 228
H.8 Conclusions . . . 228
H.9 Future experiments . . . 229

I The M13 beam line and $(g-2)$ . . . 230

J Estimating muon stops in PC6 gas . . . 232

List of Tables

1.1 Fermions in the standard model . . . 2
1.2 Muon decay modes . . . 7
1.3 Weak coupling constants: experimental limits . . . 9
1.4 Right-handed muon decay probability . . . 15

2.1 Materials traversed by the muons . . . 48

3.1 Event types in the fiducial. . . . 67

5.1 Data sets for this measurement. . . . 94

6.1 Summary of uncertainties . . . 102
6.2 Summary of μ^+ beam and fringe field uncertainties . . . 103
6.3 Beginning and end of set muon beam measurements . . . 105
6.4 Validation of the magnetic field map . . . 116
6.5 Summary of μ^+ beam and fringe field uncertainties . . . 120
6.6 Relaxation rate λ for each data set . . . 125
6.7 Resolution difference between data and simulation . . . 149
6.8 Magnetic field shape systematic: fit parameters . . . 151
6.9 Scale factors for beam intensity systematic uncertainty . . . 155

7.1 Blind results, after corrections . . . 158
7.2 Blind results, before corrections . . . 159
7.3 Corrections and set-dependent uncertainties . . . 159

A.1 Personal contributions: internal analysis reports. . . . 182

H.1 μ^+SR data sets. . . . 213
H.2 μ^+SR calibration : spin glass fit results . . . 215
H.3 μ^+SR scintillator contribution: fit results . . . 216
H.4 μ^+SR SRIM simulation results . . . 226

List of Figures

1.1 Surface muon production . . . 6
1.2 Dominant muon decay mode . . . 7
1.3 Feynman diagrams for radiative corrections . . . 12
1.4 Effect of Radiative Corrections . . . 13
1.5 Effect of $P_\mu^\pi \xi$ on the decay spectra . . . 14
1.6 Previous measurements of $P_\mu^\pi \xi$ . . . 27

2.1 Graphite production target . . . 29
2.2 M13 beam line . . . 30
2.3 Momentum calibration of beam line . . . 31
2.4 Time expansion chambers: positioning . . . 34
2.5 Time expansion chambers: module location . . . 34
2.6 Cross section of a time expansion chamber module . . . 35
2.7 Principle of quadrupole steering . . . 37
2.8 Muon beam sensitivity to Q6 steering . . . 38
2.9 TWIST spectrometer . . . 39
2.10 Side view of TWIST spectrometer . . . 42
2.11 Schematic of a wire chamber . . . 42
2.12 Upstream "beam package" . . . 45
2.13 Muon stopping target . . . 46
2.14 Summary of alignments . . . 49

3.1 Analysis procedure . . . 52
3.2 Wire time offsets . . . 54
3.3 Determining helix fit parameters. . . . 57
3.4 Determining the helix wavelength . . . 58
3.5 Event display . . . 60
3.6 `GARFIELD` and refined space-time-relationships . . . 62
3.7 Analysis using refined STRs . . . 63

3.8 Resolution function for helix fitting . . . 64
3.9 Events removed for each cut . . . 65
3.10 Time-of-flight selection . . . 66
3.11 Pulse widths in target proportional chambers . . . 69
3.12 Theoretical and reconstructed muon decay spectra . . . 73
3.13 Theoretical derivative spectra . . . 76
3.14 Kinematic endpoint in data and simulation . . . 77
3.15 Dependence of the endpoint position on $|1/\cos\theta|$. . . . 78
3.16 Theoretical predicted for muon decay asymmetry . . . 80
3.17 Exponential fit to the asymmetry . . . 81
3.18 Muon beam with large transverse momentum . . . 83

4.1 Dependence of range on x position . . . 87

5.1 Tuned muon beam . . . 92
5.2 μ^+ beam stability from time expansion chambers . . . 98
5.3 μ^+ beam stability from wire chambers . . . 99

6.1 Quadratic dependence of $P_\mu(0)$ on initial position/angle. . . . 109
6.2 `Opera` longitudinal field components . . . 110
6.3 `Opera` transverse field components . . . 111
6.4 Magnetic field measurements compared to `Opera` . . . 113
6.5 Muon beam envelopes . . . 115
6.6 Comparison I sensitivity to transverse field components . . . 117
6.7 Data-simulation agreement for scaled magnetic field . . . 118
6.8 $P_\mu(0)$ sensitivity to multiple scattering correction . . . 119
6.9 Number of hits in final TEC track . . . 121
6.10 Muon pulse widths in PC5 and PC6 . . . 126
6.11 Asymmetry analysis applied to simulation . . . 129
6.12 Background muon contamination . . . 134
6.13 Muon stopping distribution in data and simulation . . . 135
6.14 Time residuals after refining the drift cell STR . . . 137
6.15 Reconstruction efficiency: dependence on p and $\cos\theta$ . . . 139
6.16 Validation of δ-electron measurement . . . 143
6.17 Measure of the δ-electron rate . . . 145
6.18 Measure of the Bremsstrahlung rate . . . 145

6.19 Outside material systematic uncertainty 147
6.20 Comparison of resolution in data and simulation 150
6.21 Difference in kinematic endpoint for nominal data/simulation 153

7.1 $\Delta P_\mu^\pi \xi$ consistency 158
7.2 Normalised fitting residuals 160
7.3 Exclusion region for left-right symmetric models 164

C.1 Muon naming: frequency in abstracts 185

D.1 Field mapper device: schematic 189
D.2 Field mapper device: photograph 189
D.3 `Opera` simulation components 190
D.4 Difference between measured and `Opera` map 190

E.1 New downstream trigger: photograph 191
E.2 New downstream trigger: front view 192

G.1 Examples of TEC events 195
G.2 Track candidate construction 198
G.3 Collimator used to calibrate time expansion chambers 201
G.4 Stability of wire time offsets for the time expansion chambers 202
G.5 Correction parameters for TEC discriminator amplitude walk 203
G.6 Single hit resolution for time expansion chambers 204
G.7 Time expansion chambers: sense plane aging 206

H.1 μ^+SR experimental setup 211
H.2 Sensitivity of empirical asymmetry to binning 214
H.3 Second exponential term for scintillator 217
H.4 μ^+SR fit results for silver mask 219
H.5 μ^+SRfit results for TWIST metal samples 221
H.6 μ^+SR example fit to Al target 222
H.7 μ^+SR example fit to Ag target 223
H.8 SRIM simulation results 225
H.9 Sensitivity of λ to fraction of muons stopping in scintillator 227

I.1 On-axis particle passing through the beam line 230
I.2 Extreme trajectories through the M13 beam line 231

J.1 Estimating of gas contamination in zone 1 234

Acknowledgements

I would like to thank Glen Marshall and Richard Mischke for working closely with me on all aspects of the $P_\mu^\pi \xi$ measurement. They were always available for consultation, and followed each step of the measurement and this thesis in detail. Thank you to Michael Hasinoff and Alexander Grossheim for their extensive comments and feedback on this thesis. The whole TWIST collaboration has been of enormous help for this measurement, through the group weekly meetings, the student meetings organised by Art Olin, and the annual collaboration meetings.

The experiment would not have acquired data so smoothly without the assistance of the expert technical staff at TRIUMF. This $P_\mu^\pi \xi$ measurement was greatly helped by Grant Sheffer's efforts to regularly replace the time expansion chamber sense planes, and being available at all hours to insert and remove these chambers. A special thanks to my fellow graduate students, especially Anthony Hillairet for always being willing to help with any aspect of the TWIST analysis. Also thank you to Jess Brewer for his assistance with the μ^+SR analysis and theory, and to the TRIUMF CMMS group for making the subsidiary μ^+SR experiment work.

The TWIST experiment is supported by the Natural Sciences and Engineering Research Council and the National Research Council of Canada. Our foreign collaborators receive funding from the U.S. Department of Energy and the Russian Ministry of Science. Computing resources were provided by the WestGrid computing facility.

Chapter 1

Introduction

This thesis describes the TWIST collaboration's final direct measurement of $P_\mu^\pi \xi$, where P_μ^π is the polarisation of the muon from pion decay, and ξ describes the asymmetry of the positrons from muon decay. This measurement is a high precision test of the standard model of particle physics.

Muons decay to a positron and two neutrinos. The distribution of the positron's energy and angle gives information on the fundamental interaction in muon decay. The product $P_\mu^\pi \xi$ is a measure of the forward-backward asymmetry in muon decay. The TWIST experiment determines $P_\mu^\pi \xi$ by stopping muons at the centre of highly symmetric detector, and then reconstructing the decay positrons. The experiment remains blind to the value of $P_\mu^\pi \xi$ in data; this is accomplished by running a simulation of the detector with a hidden value of ξ, and then fitting *data minus simulation*. The experiment is limited by systematic uncertainties, and $P_\mu^\pi \xi$ is limited by how well we can determine the depolarising processes that the muon experiences. After the systematic uncertainties are evaluated, the hidden value of ξ in the simulation is revealed, and the result for data is then known.

The current chapter will describe the physics under investigation and the previous measurements of $P_\mu^\pi \xi$. Chapter 2 describes the delivery of muons, the time expansion chambers that measured the muon beam, and the low mass spectrometer that was used to measure the positron tracks. Chapter 3 describes the analysis that identified the particles and reconstructed their trajectories. The analysis of the time expansion chambers can be found separately in Appendix G. The detailed simulation of the particles and the spectrometer is covered in Chapter 4. A subsidiary μ^+SR experiment to determine $P_\mu(t)$ is described in Appendix H; this project was part of the author's thesis proposal, but its results were not competitive with those from the TWIST detector. Chapter 5 describes the data accumulated in 2006 and 2007 that were analysed for this measurement. The uncertainties that dominated the $P_\mu^\pi \xi$ measurement are described in Section 6. Lastly, the results and their physics implications are considered in Chapter 7; this chapter includes guidelines for an improved measurement at the level of 1×10^{-4}.

Appendix A details the author's personal contributions to the experiment. Appendices B and C describe the discovery and naming of the muon.

1.1 The standard model of particle physics

The standard model (SM) describes the fundamental particles that make up all matter, and the interactions between these particles[1]. The model is very successful, but has known limitations; an extension is needed to accommodate neutrino oscillations, gravity is not included, and the fundamental interactions are not unified under a common symmetry. The SM uses arbitrary parameters (masses, couplings, mixing angles, *etc.*) that must be measured, rather than being predicted by the model itself.

In the SM, all matter is composed of fundamental spin-1/2 particles[1] called fermions. There are six leptons, which exist as free particles, and six quarks, which have not been observed as free particles. They are grouped into three generations (I, II, III) of increasing mass scale; Table 1.1 lists the particles and their charges. The leptons are the electron (e^-), muon (μ^-), and tau lepton (τ^-), all with charge -1 (in units of elementary charge), and their associated neutrinos that have no electric charge. The quark flavours are up (u), down (d), charm (c), strange (s), top (t) and bottom (b), and in each generation there is a quark of charge ($+\frac{2}{3}$) and ($-\frac{1}{3}$). Quarks have an extra degree of freedom, "colour charge", which can be red, green or blue. For each fermion there is an associated antiparticle with the same mass but opposite charge. Antiparticles are denoted by their opposite charge (*e.g.* μ^+) or a bar (*e.g.* $\bar{p}$).

Table 1.1: Fundamental fermions in the SM, in generations of increasing mass scale[2]. Charge, Q, is given in units of elementary charge ($e \approx 1.60 \times 10^{-19}$ C).

Particle	Generation			$Q/\lvert e\rvert$
	I	II	III	
leptons	e^-	μ^-	τ^-	-1
	ν_e	ν_μ	ν_τ	0
quarks	u	c	t	+2/3
	d	s	b	-1/3

[1]Spin is an intrinsic property, such as mass or charge. More detail will be given in Section 1.2.

The SM describes three of the four fundamental interactions between fermions, which are mediated by particles of integral spin, namely the bosons. The strong interaction binds quarks, and is mediated by spin-1 massless gluons that also carry a colour charge; the interactions are described by quantum chromodynamics (QCD). Leptons do not carry colour, and are therefore unaffected by the strong interaction. The electromagnetic interaction is mediated by massless photon exchange, and both quarks and charged leptons can interact. The weak interaction is mediated by three massive charged bosons, the $W^{\pm}$ and Z_0, each with a mass of order 100 protons. Gravity is not included in the SM, but is supposedly mediated by a spin-2 boson called the graviton. Relative to the strong interaction, the strength of the force between two protons is 10^{-2} for the electromagnetic interaction, 10^{-7} for the weak interaction, and 10^{-39} for gravity[2].

The charged weak interaction (WI) can convert charged leptons into neutral leptons, and vice-versa, but only within a single generation. The WI can convert quarks between generations, by defining WI eigenstates that are a mixture of mass eigenstates. The mixing is then characterised by the Cabbibo-Kobayashi-Maskawa (CKM) matrix, V_{CKM}, defined by

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = V_{\mathrm{CKM}} \begin{pmatrix} d \\ s \\ b \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}, \tag{1.1}$$

where (d', s', b') are the WI eigenstates and (d, s, b) are the mass eigenstates. The elements of V_{CKM} are determined experimentally, and V_{CKM} is found to be close to diagonal with the latest values[3]

$$V_{\mathrm{CKM}} = \begin{pmatrix} 0.97419 \pm 0.00022 & 0.2257 \pm 0.0010 & 0.00359 \pm 0.00016 \\ 0.2256 \pm 0.0010 & 0.97334 \pm 0.00023 & 0.0415^{+0.0010}_{-0.0011} \\ 0.00874^{+0.00026}_{-0.00037} & 0.0407 \pm 0.0010 & 0.999133^{+0.000044}_{-0.000043} \end{pmatrix}, \tag{1.2}$$

under the assumption that only three generations exist.

The weak interaction has been experimentally determined to violate parity; this is the symmetry that physical laws are the same after an improper rotation ($\vec{r} \to -\vec{r}$), which is a conserved quantity under the strong and electromagnetic interactions. The operator for the weak interaction was found to be vector minus axial-vector, or $(V - A)$. This combination resulted in maximal parity violation since $(V - A)$ projects out the left-handed part of the wave function. The operator $(V - A)$ is called the SM prediction for the weak interaction since there is no experimental data to disfavour this combination of operators.

The SM includes the Glashow Weinberg Salam (GWS) model of electroweak interactions, which unifies the electromagnetic and weak interactions. At higher energies the interactions have the same strength, and the symmetry is only broken at lower energies. The GWS model addresses maximal parity violation by making left-handed particles (right-handed antiparticles) transform as doublets, and the right-handed particles (left-handed antiparticles) as singlets, so that in the first generation of fermions one has,

$$\begin{pmatrix} u \\ d \end{pmatrix}_L, \begin{pmatrix} \nu_e \\ e^- \end{pmatrix}_L, u_R, d_R, e^-_R. \tag{1.3}$$

The symmetries of the theory are weak isospin (I) and hypercharge (Y); the latter is defined as $Y = Q - I_3$, where Q is the electric charge and I_3 is the third component of weak isospin. The weak isospin determines the behaviour under the weak interaction; the left-handed fermions have $I_3 = \pm\frac{1}{2}$, and these doublets behave the same under the weak interaction. The right-handed leptons have $I = 0$, and do not participate in the weak interaction.

Although the SM includes the unification of the electromagnetic and weak interactions through the GWS theory, it does not unify the electroweak and strong interactions. This is the subject of Grand Unified Theories (GUTs). For more detail on the SM, the reader is referred to Refs. [2, 4], which contain explanations that are accessible to the experimentalist.

1.2 Spin and polarisation

The spin of a lepton is an intrinsic property, such as its mass or charge. Spin has no classical analogue, and the spin operator cannot be defined in terms of physical observables, in contrast to the orbital angular momentum operator, which is $(\vec{r} \times \vec{p})$. However, the spin operators do obey the same commutation relations as the orbital angular momentum, and therefore spin is considered to be an "intrinsic angular momentum", although nothing is actually "spinning".

Spin is quantised, and the component along a direction can only take on the values $\hbar m_i$, where $m_i = -s, -s+1, ..., s-1, s$ and $\hbar s$ is the total spin. Muons are leptons, which are spin-1/2, so that the spin in a given direction can be $\pm\frac{\hbar}{2}$.

A "spin vector" can be defined for a single particle, as the expectation of the spin along each axis. This is a useful concept since the spin vector can receive a torque in a magnetic field that results in classical precession about the field direction. For an ensemble of particles with spin, the "spin polarisation" can be introduced, which describes the degree to which the spins are aligned in a particular direction; this is a scalar quantity. From here on, the muon (spin) polarisation is denoted P_μ, and this is always defined with respect to the magnetic field at the point of decay, which is equivalent to the z-axis in the experiment's coordinate system.

1.3 Muon production

The muons in the experiment were produced from pions that decayed while at rest, in the lab frame. A high-energy proton beam incident on a stationary carbon target produced π^+, which then slowed down and decayed with a branching ratio of $> 99.98\%$[3] into

$$\pi^+ \to \mu^+ \nu_\mu. \tag{1.4}$$

In the π^+ rest frame, conservation of energy and momentum leads to a μ^+ momentum,

$$p_\mu = \frac{m_\pi^2 - m_\mu^2}{2m_\pi} \approx 29.79\,\mathrm{MeV/c}, \tag{1.5}$$

where the neutrino mass is assumed to be zero. The neutrino has negative helicity[2] (it is "left-handed"), so that its polarisation and momentum vectors are opposite[5]. The π^+ has spin-0 and the ν_μ and μ^+ are spin-1/2; therefore conservation of angular momentum ensures the muon also has negative helicity, as described in Fig. 1.1.

Thus, the muons from π^+ decay are 100% polarised, with the exception of the following mechanisms:

- Finite neutrino mass reduces the neutrino's helicity by a factor $(1 - p_\nu/E_\nu)$. Even with a conservative upper mass limit[3] of 0.19 MeV, this changes the muon's polarisation by just 2×10^{-5}, which is an order of magnitude below the experimental sensitivity. Note that cosmological data suggests the sum of the neutrino mass eigenstates is $< 2.0\,\mathrm{eV}$[3].

[2]The neutrino has negative chirality. Under the assumption of massless neutrinos (which is adequate for this measurement), $v = c$ so that the neutrino also has negative helicity.

[3]This is the 90% C.L. from muon based neutrino mass measurements[3].

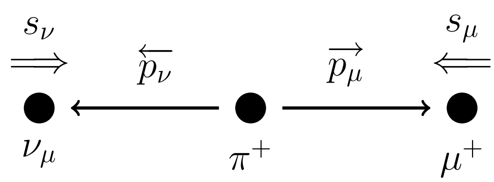

Figure 1.1: Neutrinos have their momentum ($\vec{p}$) and spin ($\vec{s}$) vectors in opposite directions (they are "left-handed"). Conservation of angular momentum requires the muon to also be left-handed.

- The π^+ has a radiative decay mode with a branching ratio of 0.02%,

$$\pi^+ \to \mu^+ \nu_\mu \gamma. \tag{1.6}$$

  In this mode, the muon's longitudinal polarisation is a function of the photon and muon energies. The present experiment selects a limited range of muon momenta, and the branching ratio for this process is already at the 10^{-4} level, so that the loss of polarisation due to this radiative decay mode is negligible.

- If the SM is incomplete, the weak interaction may allow for right-handed neutrinos in pion decay, which would force the muon to also be right-handed. The possibility of right-handed muons is part of the physics motivation for measuring $P_\mu^\pi \xi$ at the level of 10^{-4} (see Section 1.5.2).

1.4 Muon decay

1.4.1 Decay modes

The muon decays with a lifetime of $2.197\,\mu$s into the three modes listed in Table 1.2, with the most probable mode shown in Fig. 1.2. The positron is emitted with a range of energies, up to a kinematic maximum of $W_{e\mu} = (m_\mu^2 + m_e^2)/2m_\mu \approx 52.83\,\mathrm{MeV}$, which provides an energy reference feature.

Table 1.2: Muon decay modes, from the Particle Data Group[3].

Decay mode	Fraction (Γ_i/Γ)
$\mu^+ \to e^+\nu_e\bar{\nu}_\mu$	$(98.6 \pm 0.4)\%$
$\mu^+ \to e^+\nu_e\bar{\nu}_\mu\gamma$	$(1.4 \pm 0.4)\%$
$\mu^+ \to e^+\nu_e\bar{\nu}_\mu e^+e^-$	$(3.4 \pm 0.4) \times 10^{-5}$

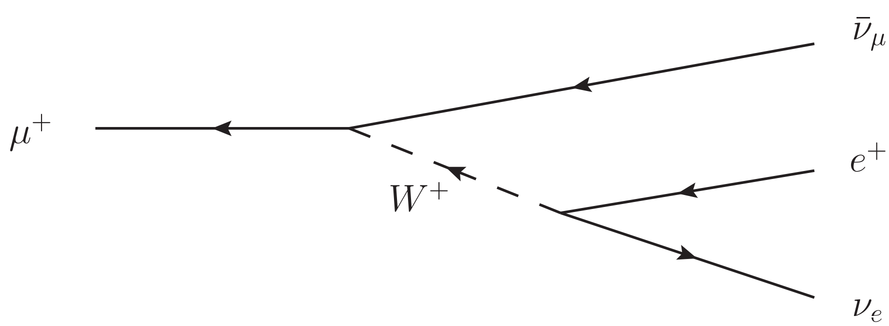

Figure 1.2: Feynman diagram for the most probable muon decay mode[6].

1.4.2 Matrix element

The Particle Data Group regularly reviews the theoretical form for muon decay[3]. For energies much less than m_W, muon decay can be considered a four-fermion point interaction. The most general, local[4], Lorentz-invariant, derivative-free, lepton-number-conserving matrix element M can be written in terms of helicity-preserving amplitudes

[4]The range of the W^+ makes the interaction non-local, but this contributes a negligible deviation $O(m_\mu^2/m_W^2) \approx 2 \times 10^{-6}$[7].

as

$$M = \frac{4G_F}{\sqrt{2}} \sum_{\substack{\kappa=S,V,T \\ i,j=L,R}} g^{\kappa}_{ij} \langle \bar{\psi}_{e_i} | \Gamma^{\kappa} | \psi_{\nu_e} \rangle \langle \bar{\psi}_{\nu_\mu} | \Gamma_{\kappa} | \psi_{\mu_j} \rangle, \tag{1.7}$$

where G_F is the Fermi coupling constant ($1.17 \times 10^{-11}\,\mathrm{MeV}^{-2}$), i and j label the electron and muon chiralities, g^{κ}_{ij} are complex amplitudes, and Γ^{κ} are the possible interactions (scalar-pseudoscalar (S), vector-axialvector (V), tensor (T)), which are given by

$$\Gamma^S = 1, \quad \Gamma^V = \gamma^\mu, \quad \Gamma^T = \frac{1}{\sqrt{2}} \sigma^{\mu\nu} \equiv \frac{i}{2\sqrt{2}} (\gamma^\mu \gamma^\nu - \gamma^\nu \gamma^\mu). \tag{1.8}$$

The amplitudes g^T_{RR} and g^T_{LL} are both zero, leaving 10 complex values of g^{κ}_{ij}, one of which is constrained by normalisation[8],

$$\begin{aligned} \frac{1}{4}\left(|g^S_{RR}|^2 + |g^S_{LR}|^2 + |g^S_{RL}|^2 + |g^S_{LL}|^2 \right) + |g^V_{RR}|^2 + |g^V_{LR}|^2 \quad + |g^V_{RL}|^2 + |g^V_{LL}|^2 \\ +3\left(|g^T_{LR}|^2 + |g^T_{RL}|^2 \right) = 1 \end{aligned} \tag{1.9}$$

Since these are complex amplitudes, there are 18 independent parameters to determine, in addition to G_F. In the SM, where the weak vertex factor has the operator combination $(V-A)$, the amplitude $g^V_{LL} = 1$ and all others are zero. The values of g^{κ}_{ij} are experimentally determined from a global analysis of several inputs:

- The muon lifetime to determine the Fermi coupling constant, G_F.
- The energy and angle of the e^+ from $\mu^+ \to e^+ \nu_e \bar{\nu_\mu}$, such as the experiment described here.
- The longitudinal polarisation of the e^+ from $\mu^+ \to e^+ \nu_e \bar{\nu_\mu}$. Note that this measurement also determines G_F.
- Inverse muon decay, $\nu_\mu e \to \mu^- \nu_e$, to place strict limits on the scalar terms.

A recent global analysis is described in more detail in Ref. [7], where 11 parameters from the above experiments are used to set confidence limits on the magnitude of the coupling constants. Table 1.3 gives the results from the last two global analyses, showing the impact of the TWIST experiment's published results for ρ and δ.

Table 1.3: 90% confidence limits on the weak coupling constants. Limits on $|g^S_{LL}|$ and $|g^V_{LL}|$ are from Ref. [3].

	Prior to TWIST[9]	First TWIST ρ, δ publication[7]	Recent TWIST results MacDonald[10]
$\lvert g^S_{RR}\rvert$	< 0.066	< 0.067	< 0.062
$\lvert g^V_{RR}\rvert$	< 0.033	< 0.034	< 0.031
$\lvert g^S_{LR}\rvert$	< 0.125	< 0.088	< 0.074
$\lvert g^V_{LR}\rvert$	< 0.060	< 0.036	< 0.025
$\lvert g^T_{LR}\rvert$	< 0.036	< 0.025	< 0.021
$\lvert g^S_{RL}\rvert$	< 0.424	< 0.417	< 0.412
$\lvert g^V_{RL}\rvert$	< 0.110	< 0.104	< 0.104
$\lvert g^T_{RL}\rvert$	< 0.122	< 0.104	< 0.103
$\lvert g^S_{LL}\rvert$	< 0.550	< 0.550	< 0.550
$\lvert g^V_{LL}\rvert$	> 0.960	> 0.960	> 0.960

1.4.3 Muon decay parameters

The muon decay parameters[5], which describe the energy and angle of the e^+ from $\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$, are defined as

$$\begin{aligned}
\rho &= \frac{3}{4} - \frac{3}{4}[\left|g^V_{RL}\right|^2 + \left|g^V_{LR}\right|^2 + 2\left|g^T_{RL}\right|^2 + 2\left|g^T_{LR}\right|^2 \\
&\quad + Re\left(g^S_{RL}g^{T*}_{RL} + g^S_{LR}g^{T*}_{LR}\right)], && (1.10) \\
\eta &= \frac{1}{2}Re[g^V_{RR}g^{S*}_{LL} + g^V_{LL}g^{S*}_{RR} + g^V_{RL}(g^{S*}_{LR} + 6g^{T*}_{LR}) + g^V_{LR}(g^{S*}_{RL} + 6g^{T*}_{RL})], && (1.11) \\
\xi &= 1 - \frac{1}{2}\left|g^S_{LR}\right|^2 - \frac{1}{2}\left|g^S_{RR}\right|^2 - 4\left|g^V_{RL}\right|^2 + 2\left|g^V_{LR}\right|^2 - 2\left|g^V_{RR}\right|^2 \\
&\quad + 2\left|g^T_{LR}\right|^2 - 8\left|g^T_{RL}\right|^2 + 4Re(g^S_{LR}g^{T*}_{LR} - g^S_{RL}g^{T*}_{RL}), && (1.12) \\
\xi\delta &= \frac{3}{4} - \frac{3}{8}\left|g^S_{RR}\right|^2 - \frac{3}{8}\left|g^S_{LR}\right|^2 - \frac{3}{2}\left|g^V_{RR}\right|^2 - \frac{3}{4}\left|g^V_{RL}\right|^2 - \frac{3}{4}\left|g^V_{LR}\right|^2 \\
&\quad - \frac{3}{2}\left|g^T_{RL}\right|^2 - 3\left|g^T_{LR}\right|^2 + \frac{3}{4}Re(g^S_{LR}g^{T*}_{LR} - g^S_{RL}g^{T*}_{RL}). && (1.13)
\end{aligned}$$

Under the SM where $g^\kappa_{ij} = 0$, except for $g^V_{LL} = 1$, the muon decay parameters are $\rho = \xi\delta = 3/4$, $\xi = 1$ and $\eta = 0$. The TWIST experiment fixes η to the result of a global analysis, and then measures ρ, δ and $P^\pi_\mu\xi$ simultaneously.

[5] These are sometimes referred to as the "Michel parameters", after the late theoretical physicist Louis Michel. He introduced the ρ and probably the η parameter, but δ and ξ were introduced by Alberto Sirlin[11].

1.4.4 Differential decay rate

The differential decay rate is proportional to $|M|^2$, and for a detector insensitive to the e^+ polarisation, this rate is given by

$$\frac{d^2\Gamma}{dx\, d\cos\theta} = k(x)\left\{F_{IS}(x) + P_\mu \cos\theta F_{AS}(x)\right\}, \tag{1.14}$$

where x is the reduced energy ($= E_e/W_{e\mu}$), θ is the angle between the muon spin and positron momentum vectors, $P_\mu = |\vec{P}_\mu|$ (the degree of muon polarisation), $k(x)$ is defined for convenience as

$$k(x) = \frac{m_\mu}{4\pi^3} W_{e\mu}^4 G_F^2 \sqrt{x^2 - x_0^2}, \tag{1.15}$$

$$W_{e\mu} = E_{\max} = \frac{m_\mu^2 + m_e^2}{2m_\mu}, \tag{1.16}$$

$$x_0 = m_e/W_{e\mu}, \tag{1.17}$$

and separate terms for the isotropic (IS) and anisotropic (AS) contributions are written in terms of Eqs. (1.10) to (1.13) as

$$F_{IS}(x) = x(1-x) + \frac{2}{9}\rho\left(4x^2 - 3x - x_0^2\right) + \eta x_0(1-x) + F_{IS}^{RC}(x), \tag{1.18}$$

$$F_{AS}(x) = \frac{1}{3}\xi\sqrt{x^2 - x_0^2}\left[1 - x + \frac{2}{3}\delta\left(4x - 3 + \left(\sqrt{1 - x_0^2} - 1\right)\right)\right] + F_{AS}^{RC}(x). \tag{1.19}$$

The superscript “RC” refers to radiative corrections, which are described in the next section.

1.4.5 Theoretical spectrum and radiative corrections

Radiative decays with internal and external particles are treated as spectrum corrections. The lowest order Feynman diagrams are shown in Fig. 1.3. The radiative corrections have a significant effect close to $x = 1$, where the rate is changed by up to 10% (see Fig. 1.4(a)). The current analysis includes the following levels of correction: full first order [$O(\alpha)$][12], leading-logarithmic second order [$O(\alpha^2L^2)$, where $L = \ln(m_\mu^2/m_e^2) \approx 10.7$][13], next-to-leading-logarithmic second order [$O(\alpha^2L)$][14, 15], and leading-logarithmic third order [$O(\alpha^3L^3)$][15], where α is the fine structure constant. The relative contributions to the spectrum for $\cos\theta = +1$ are shown in Fig. 1.4(b). The recent radiative correction publications cite the TWIST experiment as significant motivation for their calculations; radiative corrections for the total decay rate calculations have been in existence for a longer time to help with muon lifetime measurements, but these calculations are simpler since the electron mass can be neglected. The current analysis does not include $O(\alpha^2 L^0)$ corrections, which only became available in 2007[16].

The radiative corrections assume the SM coupling for the weak interaction, $(V - A)$. If the muon decay parameters are found to be different from the SM values, the radiative corrections will have to be recalculated using a more general form.

The TWIST experiment simultaneously extracts ρ, δ and $P_\mu^\pi\xi$. The change in spectrum shape due to $P_\mu^\pi\xi$ is shown in Fig. 1.5, where the large $|\cos\theta|$ and higher energy positrons have the greatest sensitivity. The experimental asymmetry of the spectrum is defined as

$$A = \frac{N_F - N_B}{N_F + N_B}, \tag{1.20}$$

where N_F is the number of forward counts and N_B is the number of backward counts. After integration over x, this quantity depends only on $P_\mu^\pi\xi$ and η, and for this reason $P_\mu^\pi\xi$ is sometimes called "the integral asymmetry parameter".

The TWIST experiment does not make a precision measurement of η. This parameter is most sensitive to lower energy e^+, and a precision measurement would have required about two months of dedicated running with a reduced magnetic field, and careful validation of the hard scattering interactions ($> 1\,\mathrm{MeV}$ energy loss) in the simulation[6]. Even with these two improvements, the statistical uncertainty would have been uncompetitive with measurements derived from the e^+ transverse polarisation[18]. In practice, TWIST fixed η to the latest global analysis value of $\eta = (-3.6 \pm 6.9) \times 10^{-3}$[7], and included the correlation with η as a systematic uncertainty.

1.5 Standard model extensions

Standard model extensions with right-handed muons and left-right symmetry will now be described, since an improved $P_\mu^\pi\xi$ measurement has sensitivity to these. There are other extensions such as supersymmetry, but the expected modifications to $P_\mu^\pi\xi$ are beyond the

[6]The hard scattering is needed since a high energy positron may experience significant energy loss in the target (a few MeV), and end up being reconstructed at a lower energy; if the simulation does not behave in the same way, this would result in an incorrect η measurement.

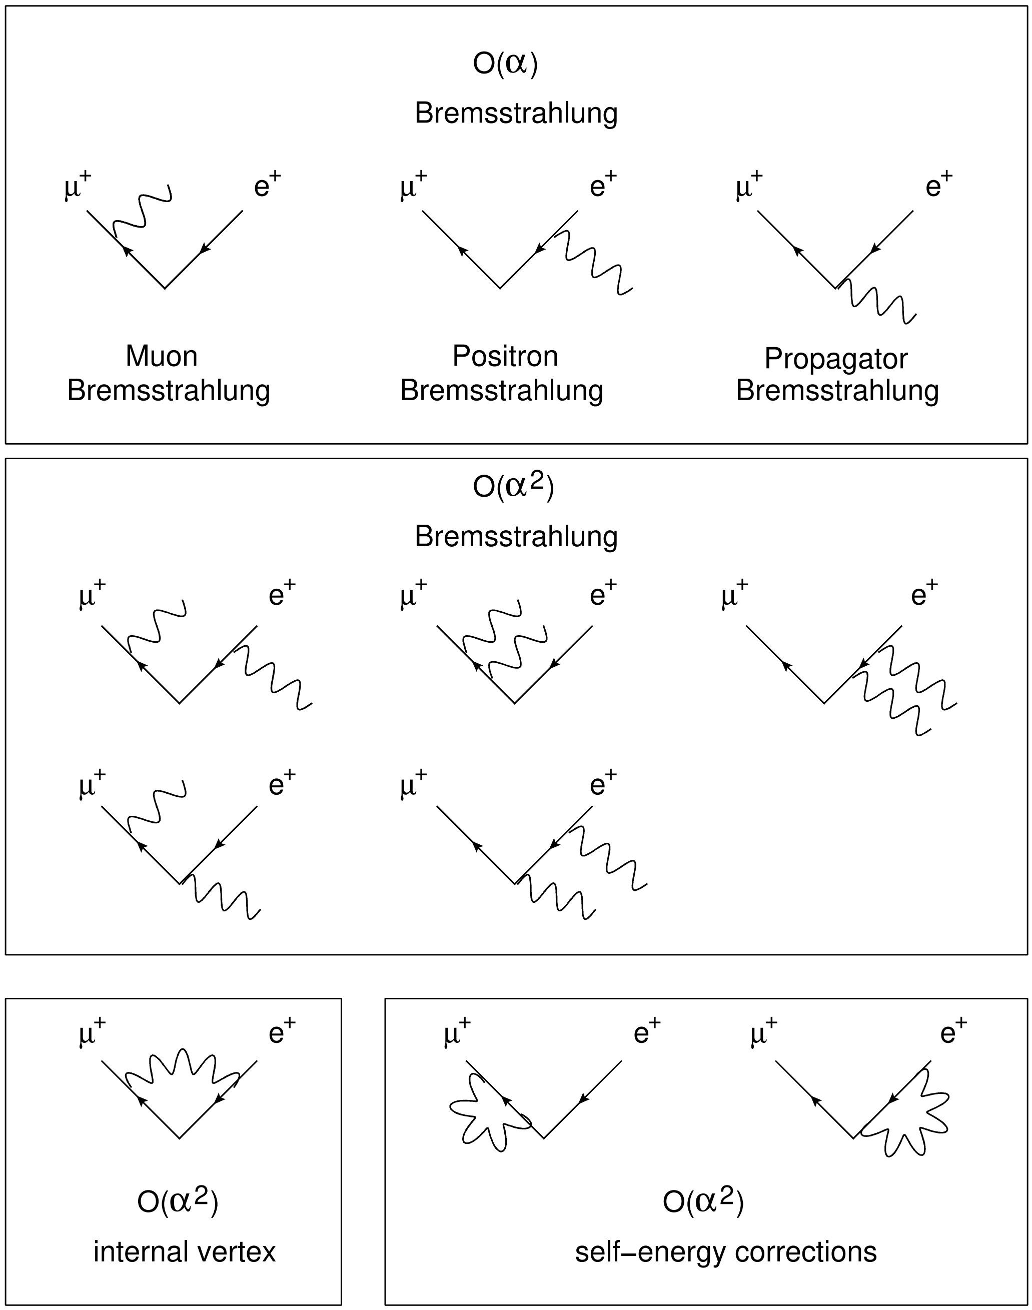

Figure 1.3: Feynman diagrams for $O(\alpha)$ and $O(\alpha^2)$ radiative corrections[17]. The neutrinos are not shown. These diagrams have a significant effect near the spectrum endpoint.

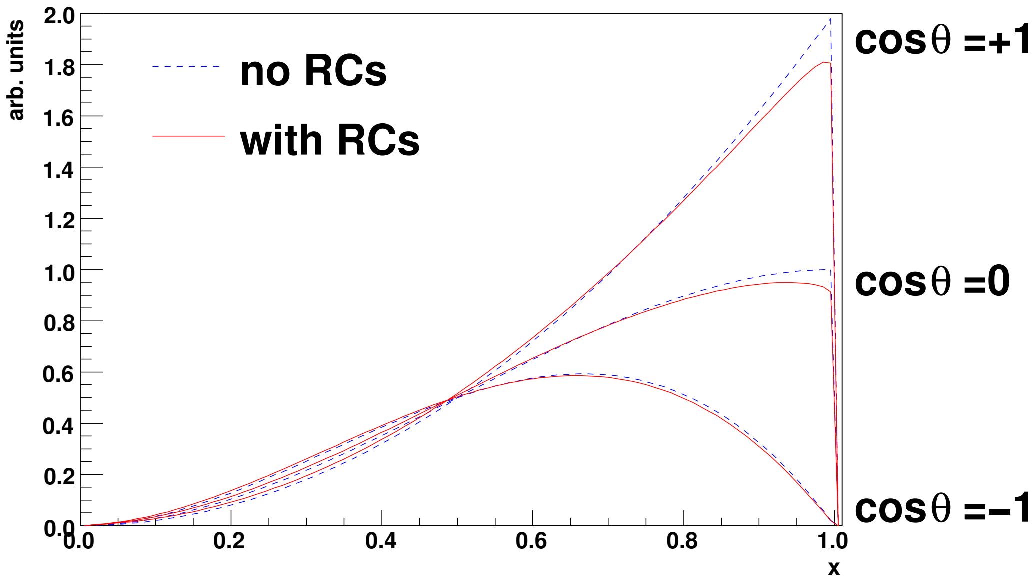

(a) Slices of the theoretical decay spectrum in $\cos\theta$. The radiative corrections (RCs) are as large as 10% at the maximum positron energy.

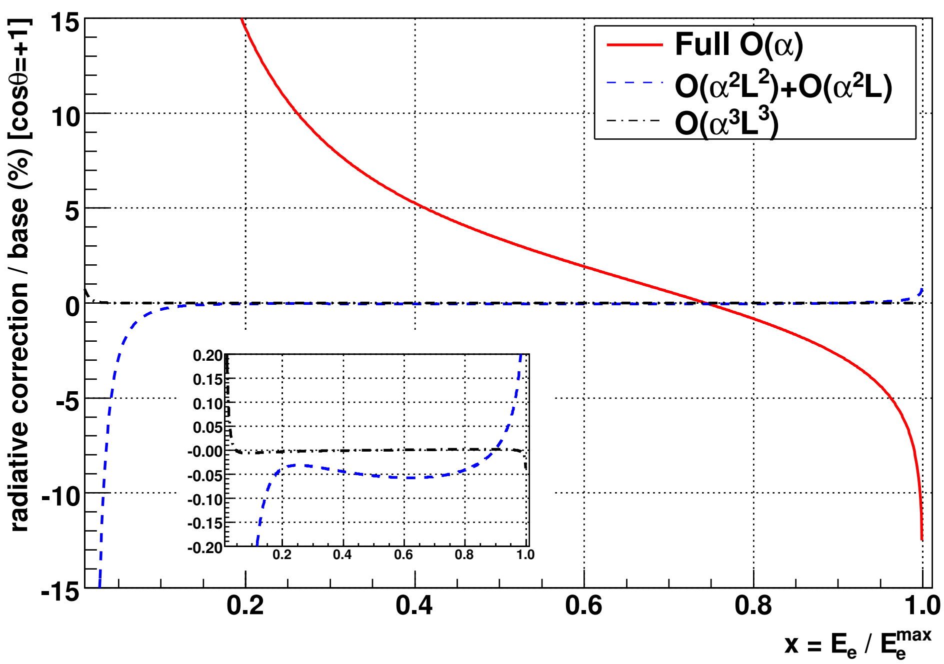

(b) Contribution of RCs for $\cos\theta = +1$

Figure 1.4: Effect of radiative corrections on the muon decay spectrum.

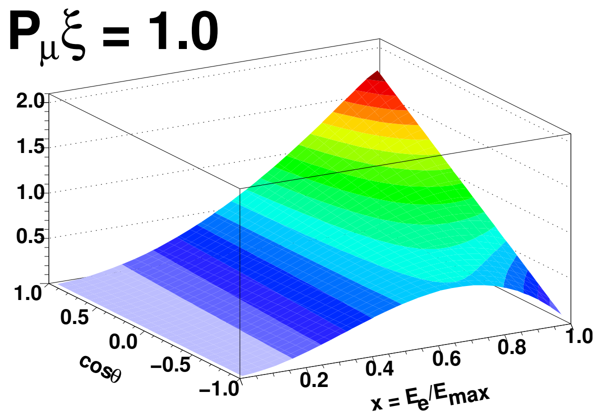

(a) Standard model, $P^{\pi}_{\mu}\,\xi = 1$.

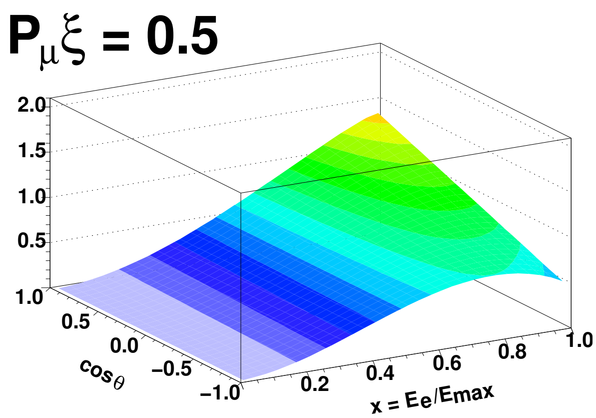

(b) $P^{\pi}_{\mu}\,\xi = 0.5$

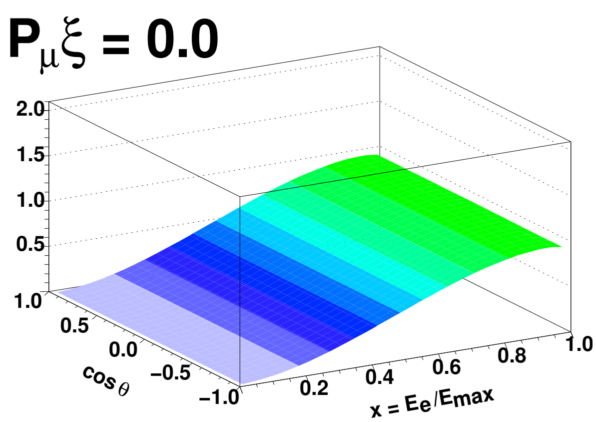

(c) $P^{\pi}_{\mu}\,\xi = 0$

Figure 1.5: Theoretical decay positron spectra (arbitrary units) for three different $P^{\pi}_{\mu}\,\xi$ values.

current experimental precision.

1.5.1 Right-handed muons

Following the notation of Eq. (1.7), the total probability for a $j-$handed muon to decay into an $i-$handed electron, Q_{ij}, can be defined as[8]

$$\begin{aligned} Q_{RR} &= \frac{1}{4}\left|g^S_{RR}\right|^2 + \left|g^V_{RR}\right|^2, & (1.21)\\ Q_{LR} &= \frac{1}{4}\left|g^S_{LR}\right|^2 + \left|g^V_{LR}\right|^2 + 3\left|g^T_{LR}\right|^2, & (1.22)\\ Q_{RL} &= \frac{1}{4}\left|g^S_{RL}\right|^2 + \left|g^V_{RL}\right|^2 + 3\left|g^T_{RL}\right|^2, & (1.23)\\ Q_{LL} &= \frac{1}{4}\left|g^S_{LL}\right|^2 + \left|g^V_{LL}\right|^2. & (1.24) \end{aligned}$$

The coefficients of g^κ_{ij} follow from the normalisation condition, Eq. (1.9). The total probability of a right-handed muon decaying into a left or right-handed electron is then

$$\begin{aligned} Q^\mu_R &= Q_{RR} + Q_{LR} \\ &= \frac{1}{4}\left|g^S_{LR}\right|^2 + \frac{1}{4}\left|g^S_{RR}\right|^2 + \left|g^V_{LR}\right|^2 + \left|g^V_{RR}\right|^2 + 3\left|g^T_{LR}\right|^2, \\ &= \frac{1}{2}\left(1 + \frac{1}{3}\xi - \frac{16}{9}\xi\delta\right). \qquad (1.25) \end{aligned}$$

where the final step used Eqs. (1.12) and (1.13). The probability Q^μ_R is equal to zero for the SM values of $\xi = 1$ and $\xi\delta = 3/4$. The most precise experimental values are shown in Table 1.4, where the published TWIST measurements of ρ, δ and $P^\pi_\mu\xi$ have already improved the constraint on Q^μ_R by a factor of six.

Table 1.4: Probability of a right-handed muon decaying into a positron of any handedness, 90% confidence limits.

Description	Q^μ_R
Best result before TWIST[19].	0.51%
Global analysis using first TWIST ρ, δ results[7].	0.31%
Global analysis using newest TWIST ρ, δ, and $P^\pi_\mu\xi = 1.0003 \pm 0.0038$[10].	0.23%

1.5.2 Left-right symmetric electroweak models

The electromagnetic and strong interactions conserve parity, yet the weak interaction has maximal parity violation. In left-right symmetric (LRS) electroweak models, a $(V+A)$ current is introduced that couples right-handed wave functions, restoring parity conservation at high-energies[20]. The gauge group is extended from $SU(2)_L \times U(1)$ to $SU(2)_L \times SU(2)_R \times U(1)$, so that right-handed fermions also transform as doublets. The distinct vector-boson fields for the $(V-A)$ and $(V+A)$ currents are then mediated by a W_L and W_R, which are related to the mass eigenstates W_1 and W_2 by

$$\begin{aligned} W_L &= W_1 \cos\zeta + W_2 \sin\zeta, \\ W_R &= e^{i\omega}(-W_1 \sin\zeta + W_2 \cos\zeta) \end{aligned} \tag{1.26}$$

where ζ is a mixing angle, and ω is a CP violating phase[7]. The W_R is much heavier than the W_L, and parity violation at low energies is a result of this mass difference. The LRS models also introduce two additional massive neutral gauge bosons. The left and right-handed interactions have separate coupling constants g_R and g_L, which correspond to g^V_{RR} and g^V_{LL} in Eq. (1.7).

The relationship between the LRS parameters and muon decay parameters has been established[20]. Assuming that neutrinos are Dirac fermions (*i.e.* particle and antiparticle are different), and there is no mixing in the leptonic sector, then P_μ, ξ, ρ have the relationships

$$P_\mu \simeq 1 - 2t_\theta^2 - 2\zeta_g^2 - 4t_\theta\zeta_g\cos(\alpha+\omega), \tag{1.27}$$

$$\xi \simeq 1 - 2(t^2 + \zeta_g^2), \tag{1.28}$$

$$\rho \simeq \frac{3}{4}(1 - 2\zeta_g^2), \tag{1.29}$$

where

$$t = \frac{g_R^2 m_1^2}{g_L^2 m_2^2}, \tag{1.30}$$

[7]CP is the product of two transformations; charge conjugation (C), which converts a particle to its antiparticle, and parity (P), an improper rotation ($\vec{r} \to -\vec{r}$). CP violation refers to a change of physical laws under the combined C and P transformations.

$$t_\theta = \frac{g_R^2 m_1^2 |V_{ud}^R|}{g_L^2 m_2^2 |V_{ud}^L|}, \tag{1.31}$$

$$\zeta_g = \frac{g_R}{g_L}\zeta, \tag{1.32}$$

α is a CP violating phase in the right-handed CKM matrix, and $V_{ud}^{L,R}$ are elements of the left and right-handed CKM matrices. Therefore $P_\mu\xi$ and ρ allow limits to be set on the mass ratio, t, and the mixing angle, ζ_g. Note that if muons are produced by K^+ decays, then the substitutions $V_{ud}^R \to V_{us}^R$ and $V_{ud}^L \to V_{us}^L$ must be made.

There are specific cases of LRS models that make further assumptions. In manifest LRS models the right- and left-handed CKM matrices are assumed to be the same, $g_R = g_L$ and $\omega = 0$ so that $t_\theta = t$ and $\alpha = 0$. Equations (1.27) and(1.28) then reduce to

$$P_\mu\xi \approx 1 - 4t^2 - 4\zeta^2 - 4t\zeta, \tag{1.33}$$

so that

$$\zeta = \frac{1}{2}\left(-t \pm \sqrt{1 - P_\mu\xi - 3t^2}\right). \tag{1.34}$$

In pseudo-manifest LRS models, $\omega \neq 0$ (CP violation is still present), and

$$P_\mu\xi \approx 1 - 4t^2 - 4\zeta^2 - 4t\zeta\cos(\alpha+\omega). \tag{1.35}$$

For a more detailed discussion of the LRS models, see Ref. [20].

Direct searches for an additional heavy gauge boson have been made at the Tevatron. The most stringent lower mass limit is 1.00 TeV at 95% C.L., by the D0 collaboration[21]. These experiments must assume a manifest LRS model, and they are insensitive to ζ. The Particle Data Group regularly reviews these searches[3].

The published lower mass limits from TWIST are $m_2 > 360\,\mathrm{GeV/c^2}$ (90% C.L., manifest LRS) and $(g_L/g_R)m_2 > 325\,\mathrm{GeV/c^2}$ (90% C.L., general LRS)[22]. The TWIST best limit on the left-right mixing angle is currently $|\zeta_g| < 0.022$[10].

1.6 Muon depolarisation mechanisms

In order to measure the polarisation at the time of muon production, all sources of depolarisation must be accurately known. This section will describe changes to P_μ during the muon's transport through a magnetic field, and after thermalisation in a metal foil.

1.6.1 Depolarisation in a magnetic field

Non-relativistically, the propagation of a spin vector is governed by the equation

$$\frac{d\vec{s}}{dt} = \frac{g}{2}\frac{e}{m}\left(\vec{s} \times \vec{B}\right), \tag{1.36}$$

where g is the Landé g factor, which is measured as $g = 2.00234$ for the muon[3], (e/m) is the charge to mass ratio, $\vec{s}$ is the spin vector, and $\vec{B}$ is the magnetic field. The motion of a particle in a magnetic field is given non-relativistically as

$$\frac{d\vec{v}}{dt} = \frac{e}{m}\left(\vec{v} \times \vec{B}\right), \tag{1.37}$$

where $\vec{v}$ is the velocity. Equations (1.36) and (1.37) have precession frequencies $\omega_p = e/m$ and $\omega_s = eg/2m$ that differ by

$$\frac{\omega_p - \omega_s}{\omega_p} = \frac{g-2}{2} = 1.17 \times 10^{-3}, \tag{1.38}$$

which demonstrates that if the momentum and spin vectors start anti-parallel, they will remain anti-parallel to a high degree while passing through a magnetic field.

The relativistic treatment of spin in electromagnetic fields can be found in Jackson[23]. In summary, the propagation is governed by the Thomas-BMT (Bargmann, Michel, Telegdi) equation[24, 25],

$$\frac{dS^\alpha}{d\tau} = \frac{e}{mc}\left[\frac{g}{2}F^{\alpha\beta}S_\beta + \frac{1}{c^2}\left(\frac{g}{2} - 1\right)U^\alpha(S_\lambda F^{\lambda\mu}U_\mu)\right], \tag{1.39}$$

where S^α is the particle's spin 4-vector, c is the speed of light, τ is the proper time, $F^{\alpha\beta}$ is the electromagnetic field tensor and U^α is the 4-velocity. Jackson manipulates this expression into the Thomas equation,

$$\frac{d\vec{s}}{dt} = \frac{e}{mc}\vec{s}\times\left[\left(\frac{g}{2} - 1 + \frac{1}{\gamma}\right)\vec{B} - \left(\frac{g}{2} - 1\right)\frac{\gamma}{\gamma+1}(\vec{\beta}\cdot\vec{B})\vec{\beta} - \left(\frac{g}{2} - \frac{\gamma}{\gamma+1}\right)\vec{\beta}\times\vec{E}\right], \tag{1.40}$$

where γ is the Lorentz factor and $\vec{\beta} = \vec{v}/c$. This is the form that we use to determine P_μ at the time of decay (see Section 4.4). There are theoretical limitations to Eq. (1.39), since it is derived for spatially homogeneous electromagnetic fields, and can only be used in inhomogeneous situations when the field gradients are "sufficiently small and the relevant effects are of first order in the spin variable"[26]; such inhomogeneous effects are completely negligible for TWIST ($\sim 10^{-12}$)[27]. In addition, quantum mechanical contributions are not a concern since the magnetic field does not change on a distance scale comparable to the de Broglie and Compton wavelengths of the muon.

1.6.2 Depolarisation while decelerating

When a muon encounters material, there is an electric field within the material's atoms, which is Lorentz transformed to a magnetic field in the muon's rest frame. This magnetic field rotates the polarisation vector, and for muons with momentum 29.6 MeV/c that entirely slow down and stop in aluminium, the resulting depolarisation is $\approx 4 \times 10^{-5}$[28], which is negligible for this measurement.

Inside material the muon is multiple scattered, which changes the momentum vector but not the polarisation vector. This changes their relative orientation, but this is not a depolarisation for TWIST, since the polarisation is unchanged relative to a fixed axis.

As the muon velocity becomes comparable to the atomic electron velocity, the time between collisions is reduced, and electron capture and subsequent electron loss can repeatedly take place. The muon forms muonium (μ^+e^-)[8], which is a hydrogen-like state that depolarises due to the hyperfine interaction between the muon and electron spins (there is a transfer of the muon's spin to the electron)[30]. This only becomes important at a few keV, and does not occur in metals since the conduction electron concentration effectively screens the muon from interactions with individual electrons[31]. For the current measurement, muons were rejected if they stopped outside of the metal target, so that depolarisation from muonium formation is not a concern.

Finally, there is depolarisation due to muon-electron scattering. This has been calculated for 29.6 MeV/c muons, and is at the level of 1×10^{-5}[31, 32].

[8]The "onium" suffix is usually reserved for bound-states of a particle and its antiparticle, such as positronium, pionium (or pion-onium) and quarkonium. Therefore the assignment of "muonium" to the state μ^+e^- is not strictly correct, and means that the bound-state $\mu^+\mu^-$ is referred to as the "bound muon-antimuon state". After more than two decades of widespread usage, the name "muonium" for μ^+e^- was eventually recommended by the International Union of Pure and Applied Chemistry (IUPAC)[29] in 2001.

1.6.3 Depolarisation after thermalisation

In the present experiment, muons were stopped in aluminium and silver foils of purity greater than 99.999%, while immersed in an external 2.0 T longitudinal magnetic field. After motional thermalisation, the muons can be depolarised by nuclear dipole moments, conduction electrons, and paramagnetic impurities. The form of the resulting depolarisation has been studied by the condensed matter community using the μ^+SR technique[33], but not to the precision required for this measurement of $P_\mu^\pi \xi$.

The behaviour of the muon after motional thermalisation and the possible depolarisation mechanisms will now be described.

Muon motion

The muon is positively charged, which limits its instantaneous position to interstitial sites (*i.e.* between nuclei), or substitutional sites, (*i.e.* "vacancies", where a nucleus is absent from the lattice). When nearly thermalised, a muon lowers its energy by attracting a screening charge of conduction electrons[34]. After motional thermalisation, room temperature ensures that a muon is not stationary, but instead it usually diffuses ("hops") between energetically allowed sites before decaying (there is also a zero-point motion that causes the muon to undergo spin exchange with its nearest neighbours). The conduction electrons in aluminium and silver efficiently screen the ionic potentials, allowing for high mobility, even over macroscopic distances of $\approx 1\,\mu$m[35, 36].

Muons can become trapped at defects, of which there are a wide variety. The most common defects originate from the manufacturing process, such as when an aluminium sheet is cold-rolled to produce a thin foil, and these defects can limit the muon mobility with a strong sample dependence[36]. The defects can be enhanced by quenching, where the sample is heated and then rapidly cooled. They can be removed by annealing, a process in which the metal is maintained close to the melting point for several hours and then slowly cooled. The current silver target was annealed in an inert argon atmosphere, after machining. High purity aluminium has been studied under annealing and quenching, over a temperature range of 19 K to 900 K: most defects were found to be absent after allowing the quenched sample to reach room temperature[37].

Thermalisation of the muon itself can result in lattice defects. The muon imparts recoil energy to the lattice on a time scale 10^{-17} s, and the lattice distributes this energy to neighbouring atoms in about $\sim 10^{-12}$ s[38]. A nucleus can be knocked out of its lattice position and into an interstitial site, leaving a vacancy (a "Frenkel pair"). However, these vacancies are unlikely to affect the muon's diffusion since they are eliminated in $\sim 10^{-11}\,$s[38], and the location of thermalisation is of order $1\,\mu$m from the last defect introduced[39].

Nuclear dipole moments

The magnetic field experienced by a muon due to nuclei and lattice impurities can be modelled as static, isotropic and Gaussian. At each hop, the muon sees a new, randomised field, and the depolarisation is then given by[40, 41],

$$P_\mu(t) = P_\mu(0)\exp\left\{-\frac{2\Delta^2}{\nu^2}\left[\exp(-\nu t) - 1 + \nu t\right]\right\}, \tag{1.41}$$

where Δ is a parameter describing the magnetic field distribution[9] and $1/\nu$ is the mean time between hops (γ_μ is the muon's gyromagnetic ratio). If an external field $\vec{B}_{\rm ext}$ is now applied in a direction *transverse* to the muon polarisation, the muon spins precess and are depolarised according to the Abragam formula[10],

$$P_\mu(t) = P_\mu(0)\exp\left\{-\frac{\Delta^2}{\nu^2}\left[\exp(-\nu t) - 1 + \nu t\right]\right\}\cos\left(\omega_\mu t\right), \tag{1.42}$$

where $\omega_\mu = \gamma_\mu B_{\rm ext}$. In the "motional narrowing" limit, the muons move quickly so that ν is large, $\exp\left(-\nu t\right) \to 0$, and the envelope of Eq. (1.42) limits to an exponential time dependence. In the static limit, the envelope limits instead to a Gaussian time dependence. If a *longitudinal* field B_0 is applied instead, the static relaxation function becomes[43],

$$\begin{aligned} P_\mu(t) &= 1 - \frac{2\Delta^2}{\omega_0^2}\left[1 - \exp\left(-\tfrac{1}{2}\Delta^2 t^2\right)\cos\omega_0 t\right] \\ &\quad + \frac{2\Delta^4}{\omega_0^3}\int_0^t \exp\left(-\frac{1}{2}\Delta^2\tau^2\right)\sin\omega_0\tau\, d\tau, \end{aligned} \tag{1.43}$$

where $\omega_0 = \gamma_\mu B_0$, and the longitudinal field is seen to suppress the depolarisation due to nuclear dipole moments. The largest observed field on a muon in a crystal cell is $\Delta/\gamma_\mu = 0.47\,$mT[41], and the field at the metal target in TWIST is longitudinal with $B_0 = 2\,$T, so that $(2\Delta^2/\omega_0^2) < 10^{-7}$. Depolarisation by nuclear dipole moments is therefore

[9] Each field component is independently modelled by $\mathcal{D}(B_{\rm local}) \sim \exp\left[-B_{\rm local}^2/(2\Delta^2/\gamma_\mu^2)\right]$, so that Δ/γ_μ is the standard deviation of a Gaussian distribution.

[10] See p.439 of Ref. [42]. Note that this is also called the Anderson form, or the Kubo-Tomita form.

negligible for TWIST in the static limit (no muon hopping)[11].

Depolarisation in aluminium and silver from nuclear dipole moments has been measured in μ^+SR experiments, using a transverse magnetic field arrangement[33]. There are more studies on aluminium since its nuclear dipole moment is about 35 times larger than silver. Even with its large dipole moment, high purity aluminium leads to almost negligible depolarisation down to 1 K[44, 45, 46, 47, 48]. As a result, experimenters have to dope the sample with impurities to actually measure depolarisation[34, 49, 50].

There is a contradictory measurement in Ref. [51], which used aluminium and silver targets of 99.99% purity in a *transverse* field arrangement, at room temperature, and observed a Gaussian form for the depolarisation in aluminium. They explained this anomalous result as muons trapping in defects, which originated from the cold-rolling during manufacture of the foil[52]. Note that a later publication by the same group used a *longitudinal* field, and found no clear evidence of depolarisation, as expected from Eq. (1.43).

Korringa relaxation

A hyperfine contact interaction between the muon spin and the conduction electron spins can lead to a depolarisation. The theory for such an interaction was originally treated in the context of NMR[12] by Korringa[53], and is therefore named "Korringa relaxation". In the context of muons, the conduction electrons hop on and off the muon, making the net hyperfine coupling experienced an average of the electron spin orientations[54, 55], A simultaneous flip of the electron and muon spins can take place, with the energy provided by a change in the electron's kinetic energy[42]. The participating electrons are within kT of the Fermi surface, and the exponential relaxation rate (λ) is proportional to $K_\mu^2 T$. The constant K_μ is the "muon Knight shift", which is a characteristic value associated with the host material; the signature of Korringa relaxation is a linear increase in rate with temperature, and insensitivity to applied magnetic field[54].

Korringa relaxation has been observed in several non-magnetic metals (lead, cadmium,

[11] If the muon hops at exactly the Larmor frequency, it "sees" a magnetic field that fluctuates at the resonant frequency to drive transitions between Zeeman states. Since the muon moves stochastically, this mechanism is only relevant for a small fraction of hops. Usually the higher the field, the smaller the fraction, and the slower the relaxation[36].

[12] In NMR spectroscopy, a substance is immersed in a static magnetic field and then exposed to electromagnetic (EM) radiation. For the nuclei in the substance that have intrinsic magnetic moments (those with an odd number of protons or neutrons), the static magnetic field creates an energy difference between the spin states. The frequency of the EM radiation is swept, and peak absorption will occur ("resonance") when the energy of the photons matches the energy difference between the spin states.

zinc, copper), where the relaxation rates increased with temperature, and were robust to field changes in the range 0.010 T to 0.200 T[56].

Impurities

The aluminium stopping target was purchased from Goodfellow, who gave the typical impurities as Cu 0.3 ppm[13], Fe 0.3 ppm, Mg 1.2 ppm, and Si 0.8 ppm. The silver stopping target was purchased from ESPI Metals, who gave the typical impurities as Fe 2 ppm, Bi $<$ 2 ppm, Cu 0.6 ppm, and Pd 0.6 ppm. The impurities take the place of an aluminium or silver atom. In aluminium, the muons are not trapped by impurities above $\approx 100\,\mathrm{K}$[48]. In silver, there is evidence that room temperature trapping at impurities can occur[57].

The non-paramagnetic ions can depolarise muons due to their nuclear dipole moments, but this is heavily suppressed by a longitudinal magnetic field (see Section 1.6.3). Paramagnetic ions (Fe in this case) are a concern since they can depolarise due to their electronic dipole moment, which is much larger than the nuclear dipole moment, producing fields up to 1 kG at a distance of one lattice spacing[58, 59]. The depolarisation form for a muon trapped at a paramagnetic ion is exponential[42, 57, 58].

Summary

The expected form of $P_\mu(t)$ in the TWIST targets is exponential. The arguments for this form will now be summarised.

Muon spin relaxation due to nuclear dipole moments, whether the source is metal nuclei or non-paramagnetic impurities, is heavily suppressed by the presence of a longitudinal field. Even if there were contributions from nuclear dipole moments, the muons are expected to be in rapid motion (the experiment is performed at room temperature), for which the appropriate form is exponential. If the muons become trapped at defects or vacancies, which is possible in our aluminium foil since it was not annealed, the longitudinal field holds the muon spin against depolarisation. Later it will be shown that the measured relaxation rates for silver and aluminium differ by roughly a factor of two, yet the nuclear dipole moments differ by a factor of 35, providing further evidence that the depolarisation is not from nuclear dipole moments.

In the silver foil, trapping at paramagnetic impurities cannot be excluded. The appropriate form is then exponential, as long as the muons diffuse sufficiently fast to find

[13]ppm = parts per million. A concentration of 1 ppm corresponds to one impurity for every 100 crystal cells.

the impurities promptly, which is a good assumption in an annealed sample such as the TWIST foil. For both foils, Korringa relaxation is expected to occur; this has been observed in other metals, and has an exponential form.

1.7 Previous $P_\mu^\pi \xi$ measurements

In 1956, Lee and Yang observed that parity is conserved in strong and electromagnetic interactions, but in weak interactions "is so far only an extrapolated hypothesis unsupported by experimental evidence"[60]. They suggested several experiments to investigate parity conservation in the weak interaction, including the asymmetry of muon decay.

The angular distribution of decay positrons will generally follow a distribution of the form $(1 + a\cos\theta)$, where θ is the angle between the muon polarisation and the decay positron. For $\eta = 0$, $a = P_\mu^D \xi/3$, where ξ is the intrinsic asymmetry parameter, and P_μ^D is the polarisation of the muon at the moment of decay. The polarisation at the time of muon production can only be inferred if all sources of depolarisation are evaluated. Note that $a = 0$ implies that parity is conserved, and $P_\mu \xi = \pm 1$ implies maximal violation of parity.

Two classes of asymmetry measurements will now be described: $P_\mu^\pi \xi$ and $P_\mu^K \xi$, where muons are produced from pion and kaon decay respectively. The results are interpreted differently in the context of LRS models (see Section 1.5.2). In addition, τ decay experiments have confirmed $P_\mu^\tau \xi$ that is consistent with 1.0 using the modes $\tau^\pm \to \mu^\pm \nu\bar{\nu}$ and $\tau^\pm \to e^\pm \nu\bar{\nu}$, but these measurements have statistical uncertainties greater than 10%[61, 62, 63]. The τ decay experiments are more useful as a check on lepton universality, rather than as a precision asymmetry measurement.

1.7.1 Measurements of $P_\mu^\pi \xi$

A year after Lee and Yang's publication, Garwin *et al.* stopped positive muons from pion decay in carbon, and the angular distribution of decay positrons was found to follow $(1 + a\cos\theta)$, and "$a = -\frac{1}{3}$ with an estimated error of 10%"[64]. Assuming no depolarisation took place while slowing down, stopping, and during the $1\,\mu$s the muon spent in the carbon target, this result suggested $P_\mu \xi = 1.0 \pm 0.1$, which was consistent with maximal parity violation.

Over the next three years (1957 to 1960) many similar experiments took place[65]. A popular technique was to stop a beam of pions in a nuclear emulsion[14]; this had the advantage of delivering muons with full polarisation, but the disadvantages of low statistics and poorly determined depolarisation within the emulsion itself due to muonium formation. This depolarisation was found to depend strongly on the applied magnetic field, yielding values of $P_\mu \xi$ in the range 0.33 ± 0.03 (zero field) to 0.97 ± 0.06[65].

Other techniques suffered from problems in producing a high polarisation muon beam. Before the surface muon beam was developed, a production target was simply placed inside a cyclotron, leading to a a muon beam with poorly determined polarisation.

The Bardon, Berley and Lederman experiment[66] instead used a π^+ beam that decayed in-flight, to produce a highly transverse polarised μ^+ beam at the Jacobian angle. The μ^+ were stopped in a bromoform target, which had small but unknown depolarising effects; in consequence, the experiment's result was a lower limit. Depolarisation from magnetic fields was eliminated by using Helmholtz coils to cancel the cyclotron's field. Decay positrons were detected with opposing counters. A solenoid surrounding the target rotated the muon spin by $\pm 90°$, allowing the e^+ distribution to be turned around in one hour cycles. The experiment found $|\xi| \geq 0.97 \pm 0.05$.

Muons were also stopped in liquid hydrogen in bubble chamber experiments. The particles left ionisation tracks that were curved by a magnetic field to determine their momentum. The most accurate experiment was carried out by Plano, who measured $P_\mu \xi$, δ and ρ, and found $|\xi| = 0.94 \pm 0.07$[67].

There was apparently no improvement in direct measurements until 1968, when a nuclear emulsion experiment with a 14 T pulsed magnetic field determined that $P_\mu \xi = 0.975 \pm 0.015$, which is 1.7σ from the (V-A) prediction[68].

Prior to the TWIST experiment, the most precise direct measurement of $P_\mu^\pi \xi$ was reported in Ref. [69]. A π^+ beam with a momentum of 150 MeV/c decayed in-flight, and a transversely polarised muon beam was extracted close to the Jacobian angle; this is the same approach taken by the Bardon experiment described above[66]. The muons were moderated by aluminium and stopped in a beryllium (Be) target. The μ^+SR technique was then used to measure the decay e^+ asymmetry, with the muon spins precessed by a 3 mT field. A pure exponential depolarisation form for $P_\mu(t)$ was assumed in the Be target. The experiment did not measure the decay e^+ momentum, so there was no magnetic field requirement, and hence no associated systematic uncertainty. The contribution from muons stopping in a trigger scintillator immediately before the Be target was measured

[14] A photographic plate is exposed to particles and later developed. The emulsion is predominantly silver bromide by weight, but by number of atoms, it's 25% silver bromide, 75% gelatin[65].

by removing the target, and found to be negligible (2×10^{-3} per event). The final result was $1.0027 \pm 0.0079\,(\text{stat.}) \pm 0.0030\,(\text{syst.})$.

The most precise direct measurement is $P_\mu^\pi\,\xi = 1.0003 \pm 0.0006\,(\text{stat.}) \pm 0.0038\,(\text{syst.})$, published by the TWIST collaboration in 2006[22, 59]. The largest contribution to the total systematic uncertainty was 34×10^{-4}, which was mainly due to limitations in the reproducibility of the muon beam's initial position and angle; the current measurement aims to reduce this uncertainty by improving the knowledge of the muon beam and solenoidal magnetic field through which the beam passes.

In Fig. 1.6, the *direct* measurements that have been described so far are summarised. All are seen to be consistent with $P_\mu\xi = 1$. Additional *indirect* measurements are possible using the result $P_\mu\xi\delta/\rho > 0.99682$ (90% C.L.)[31, 32], where muons were stopped in several targets, including high purity silver and aluminium foils. The confidence interval can be combined with published ρ and δ measurements[10],

$$\rho = 0.75014 \pm 0.00017\,(\text{stat.}) \pm 0.00044\,(\text{syst.}) \pm 0.00011\,(\eta), \tag{1.44}$$

$$\delta = 0.75067 \pm 0.00030\,(\text{stat.}) \pm 0.00067\,(\text{syst.}), \tag{1.45}$$

to determine

$$0.99524 < P_\mu^\pi\xi \leq \xi < 1.00091 \quad (90\%)\,\text{C.L.} \tag{1.46}$$

1.7.2 Measurements of $P_\mu^K\xi$

The first measurement of $P_\mu^K\xi$ with an uncertainty below 10% is described in Refs. [71, 73]. (There were three prior experiments that were consistent with $P_\mu^K\xi = 1$, but with uncertainties > 10%[74, 75, 76].) A 236 MeV/c muon beam was produced from stationary K^+ decays via the mode $K^+ \to \mu^+\nu$. The μ^+ were degraded by carbon, and stopped in an aluminium target of 99.99% purity. The muon trajectory immediately before the target was measured using a pair of orthogonal multiwire proportional chambers (MWPCs), which established the muon spin, and the decay positrons were measured with a different pair of MWPCs. A transverse magnetic field was applied to precess the muon spin, and the asymmetry's dependence on time was fit to determine $P_\mu^K\xi = 0.970 \pm 0.047\,(\text{stat.}) \pm 0.005\,(\text{syst.})$, under the assumption of no depolarisation while slowing down, and within the aluminium target; these are safe assumptions at the few percent level of precision. The measurement was limited by statistical uncertainty.

A more precise measurement using a similar technique is described in Ref. [72]. The

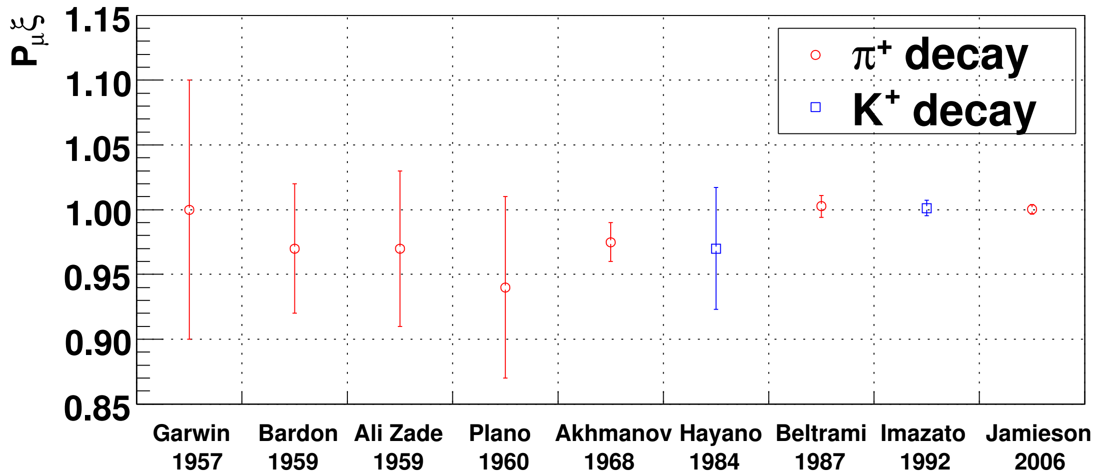

Figure 1.6: Previous direct measurements of $P_\mu\xi$ with uncertainties less than 10%. The statistical and systematic uncertainties have been added in quadrature. Measurements are from Refs [22, 64, 66, 67, 68, 69, 70, 71, 72].

muon source was the same, but with higher flux, and the muons were again degraded by carbon and stopped in high purity aluminium ($>$ 99.999% in this case). The incident muon beam was measured more carefully to determine its spin, using four x and y layers of drift chambers before the carbon degrader, and two MWPCs after the degrader and before the aluminium stopping target. The decay positrons were measured using two arms, each with four horizontal and four vertical drift chambers. A transverse magnetic field precessed the muon's spin. The polarisation was allowed to relax in the aluminium target, but the form for $P_\mu(t)$ was not stated. The final result was $P_\mu^K\xi = 1.0013 \pm 0.0030\,(\text{stat.}) \pm 0.0053\,(\text{syst.})$, where the leading systematic uncertainty was from a correction due to knock-on electron production in the stopping target.

Chapter 2

Apparatus

2.1 Overview

This chapter will describe the production and delivery of muons, the time expansion chambers that measured the muon beam, and the detector used to track the muons and reconstruct the positron trajectories. The alignment of the apparatus will also be described. There are publications that provide more detail on the detector[77] and the time expansion chambers[78]. The beam line has been described in an earlier publication[79], but the tune used for the current experiment is only available as an internal report[80].

2.2 Muon production and delivery

2.2.1 Location

The experiment was performed at the TRIUMF laboratory, located on the University of British Columbia's campus in Vancouver. The TRIUMF cyclotron delivered a 500 MeV proton beam to the T1 production target in the TRIUMF meson hall. The protons were produced in bunches of ≈ 4 ns width, separated by 43 ns. A 10 mm thick graphite production target was used at the T1 position. The particles from the graphite target were selected by a secondary beam line named M13. The experiment occupied the M13 beam line area from November 2001 to September 2007.

2.2.2 Muon production

The proton beam incident on the graphite target produced pions. The dominant pion decay mode is to a muon and a neutrino, a two body process that ensures the muon has momentum 29.79 MeV/c in the pion centre of mass frame. The experiment selected positive muons that decayed from stationary pions within microns of the production target surface. These are called "surface muons", and are highly polarised as described in Section 1.3.

The graphite production target was housed in a water cooled steel holder; it is shown schematically in Fig. 2.1(a). The targets typically lasted longer than a month, and were changed either due to destructive damage from the proton beam (see Fig. 2.1(b)) or after producing measurable levels of radioactive gas; this did not impair the data quality since the gas was mostly stopped by a 3 μm polyester (Polyethylene terephthalatate) window valve in the beam line[15].

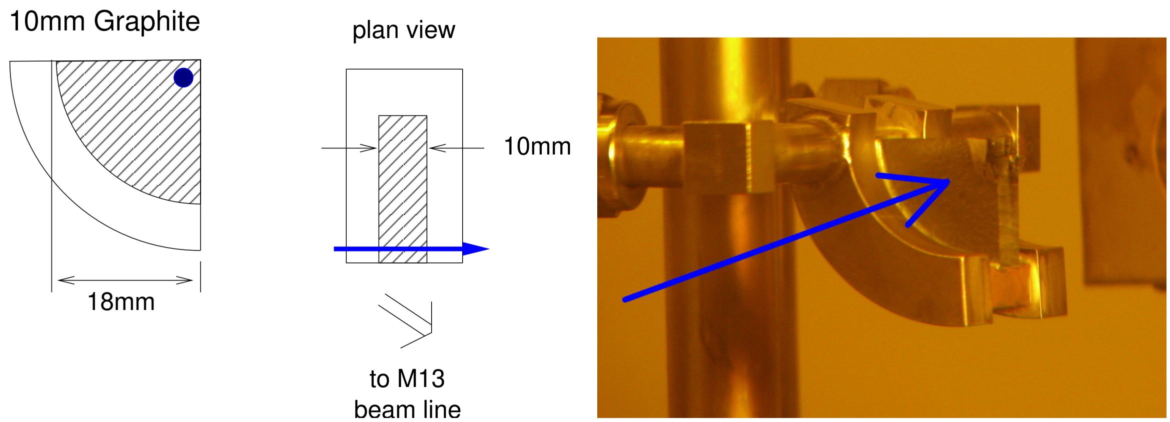

(a) Schematic of the graphite production target. The blue arrow indicates the proton beam.

(b) Photograph of a damaged graphite production target. The blue arrow indicates the proton beam direction.

Figure 2.1: Schematic and photograph of a graphite production target.

2.2.3 M13 beam line

The M13 beam line[79] that selected and transported the muons is shown schematically in Fig. 2.2. The graphite production target is labelled as 1AT1, and this was seen by M13 at 135°. The channel had two dipole magnets (B1 and B2) that selected the particle momentum by bending the beam through 60°, three vertically focusing quadrupole magnets (Q1, Q4, Q7), and four horizontally focusing quadrupole magnets (Q2, Q3, Q5, Q6). The M13 beam line was upgraded for the current measurement by adding current sources to certain quadrupole magnets, allowing additional muon beam steering; this technique will be described in Section 2.4.

[15]During the 2006 data acquisition, a small amount of gas was able to enter the detector via a bypass line. The particles making it to the detector were successfully identified using the off-line data analysis. For the 2007 data acquisition, a pump was installed to remove any radioactive gas in the bypass line.

The S1 label in Fig. 2.2 indicates the position of a horizontal slit and a vertical jaw, which controlled both the rate and the emittance of the beam. A compromise was reached that minimised emittance while maintaining an acceptable rate; this measurement used a muon rate of between $2000\,\mathrm{s}^{-1}$ and $5000\,\mathrm{s}^{-1}$. Even at the highest rate the muons were separated by nearly 100 lifetimes, ensuring that most events had just one muon in the detector at a time.

There are three possible focal points (F1, F2 and F3) indicated in Fig. 2.2. The tune for this measurement aimed for horizontal focii at F1, F2 and F3, and a vertical focus at F3. The beam's vertical divergence was limited at F2 by a slit. After B1 had dispersed the beam, the position of a horizontal slit at F1 selected the average momentum, and the width of this slit defined the momentum resolution of the channel.

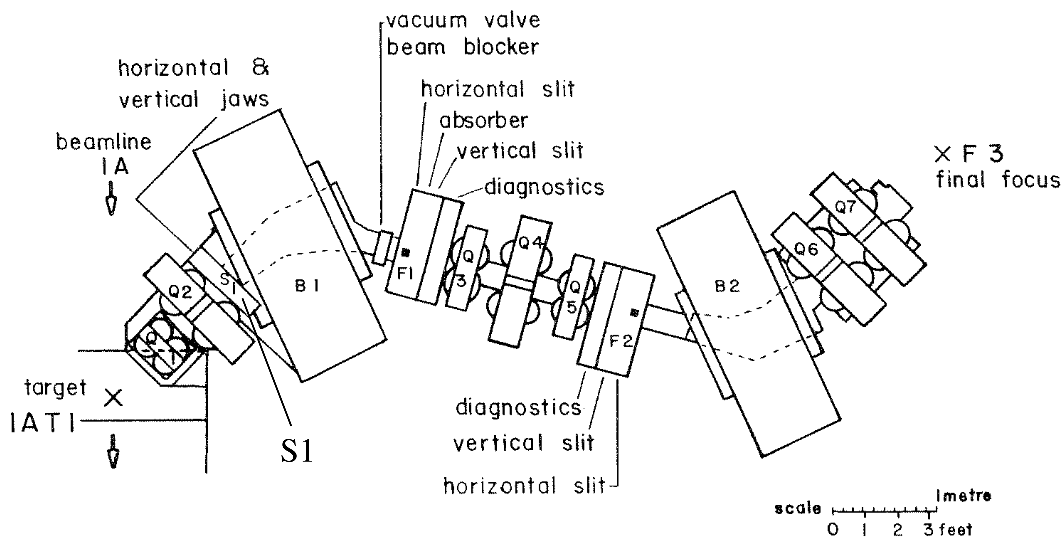

Figure 2.2: Schematic of the M13 beam line[79]. The distance from the production target to the centre of the detector (not shown) is approximately 10m.

2.2.4 Beam line momentum calibration

The initial muon momentum distribution was measured on a weekly basis using the dependence of the muon rate on channel momentum. An example is shown in Fig. 2.3, where the shape arises from a convolution of the initial muon momentum at the target and the approximately Gaussian momentum acceptance of the channel. The "surface muon edge" is a calibration feature, corresponding to muons from the surface of the target that have a momentum of exactly 29.79 MeV/c (kinetic energy 4.12 MeV). Note that for $p > 30$ MeV/c the rate is non-zero, since muons from pion decay-in-flight are accepted.

Operating at the surface muon momentum is undesirable since the rate is low. Instead the experiment used an average momentum of 29.6 MeV/c (kinetic energy 4.07 MeV), which corresponds to a thin layer (about 16 μm) close to the production target's surface. The selected muons underwent a small amount of multiple scattering inside the production target, which breaks the relationship between the momentum and spin, and later we make a systematic correction for this. The beam line was set to a momentum resolution of FWHM[16] 0.7%.

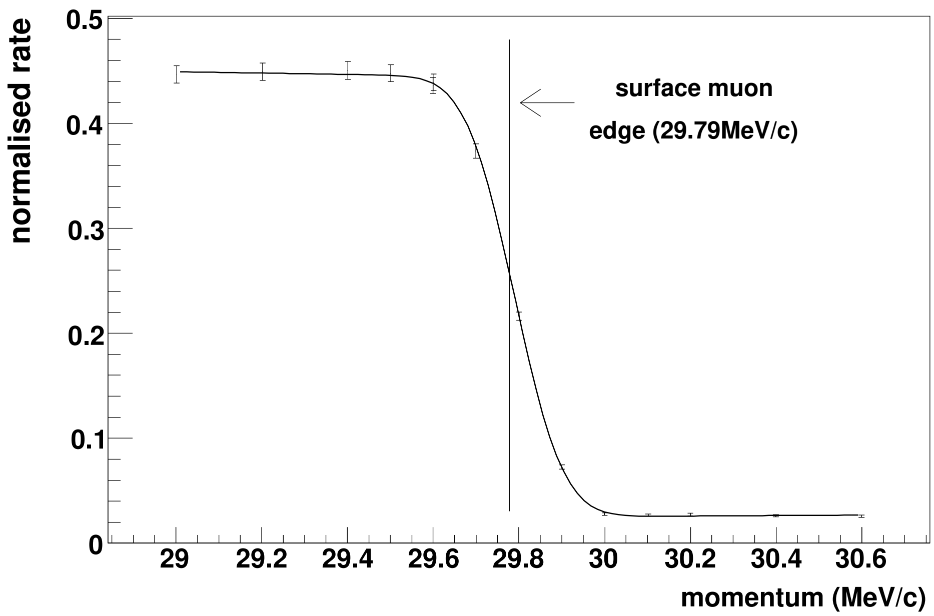

Figure 2.3: Rate dependence on channel momentum. The shape arises from a convolution of the initial muon momentum at the target and the beam line's approximately Gaussian acceptance. The "normalised rate" on the ordinate is proportional to the muon rate divided by the proton current.

[16]The momentum resolution, $\Delta p/p$, is sometimes referred to as the momentum "bite". In either case the meaning of Δp must be qualified, since there can be confusion between a Gaussian standard deviation and the full width at half maximum (FWHM).

2.2.5 Particles delivered

The muon beam had a contamination of "cloud muons", "beam positrons", pions and protons. The "cloud muons" originated from pions moving between the production target and B1 that decayed with a muon in the direction *opposite* to their motion. This decay must occur before the pion reaches the momentum selection at B1, and is therefore prompt. These muons had opposite polarisation to the surface muons, and were removed by a time-of-flight cut.

The contamination of "beam positrons" came from muon decays within the production target and its surrounding materials. There were also contributions from neutral pion decays ($\pi^0 \rightarrow \gamma\gamma$) since the γ's shower in the production target and beam pipe. The trigger scintillator was set to mostly ignore positrons, but they were high rate (1.9 to 3.6 times the muon rate, depending on the beam tune) and therefore still produced a significant background in muon-triggered events. The off-line analysis easily identified the beam positron trajectories since they passed through the entire detector.

The proton contamination was removed by a beam line window, and the pion contamination could be separated by time-of-flight due to their heavier mass[17]. The time-of-flight selection is described later in Section 3.3.1.

2.2.6 M13 configurations

The alignment and calibration runs used different M13 settings. For example, the drift chamber alignment used higher momentum pions, and the calibration of the time expansion chambers used a spread muon beam. Three smaller data sets were taken at lower momentum as a consistency check. The special configurations and data sets are described in more detail in Section 5.3.

2.3 Measuring the muon beam

The muon beam was measured at the end of the M13 channel, before the muons had encountered any significant material or magnetic field, so that the momentum and spin vectors were anti-parallel to an acceptable approximation. A simulation then propagated the muons to the stopping target to estimate the final polarisation of the beam. A pair

[17] The experiment has observed the decay $\pi^+ \rightarrow e^+\nu_\mu$, by selecting pions using their time-of-flight and pulse height in the trigger scintillator. However the statistics were too low to make a useful energy calibration feature.

of orthogonal time expansion chambers (TECs) measured x (y) and θ_x (θ_y). The mass of the system was kept as low as possible to minimise multiple scattering, which would otherwise degrade the resolution. Despite these efforts, the multiple scattering meant that the TECs could not remain in place during normal data acquisition. Instead the modules were inserted on a weekly basis for approximately one hour of beam measurements.

The positioning of the box containing the modules is shown in Fig. 2.4. The TEC modules were located close to the F3 focus (see Fig. 2.2), where the beam had a small extent and the magnetic field was weak. When the muon beam encountered the TECs, it had passed through no material except for a single thin (6 μm Mylar) beam line window. The weak magnetic field meant that trajectories were well approximated by straight lines, which simplified the analysis.

The TEC modules are labelled "X Module" and "Y Module" in Fig. 2.5. Each module was 8.0 cm long with an active area of 6.0 cm $\times$ 6.0 cm. They were inside a box that was filled with low pressure (8 kPa) dimethyl ether[18] gas that constantly flowed at a rate of 100 cm^3/min. An electric field was maintained over the drift volume by a graded voltage applied to the (drift) field wires, as shown in Fig. 2.6. Charged particles ionised the gas, and the ionisation electrons then drifted towards the sense plane. The electrons entered a high field region, separated by grid ("guard") wires. In this "multiplication region" the electrons were accelerated towards small diameter (25 μm) sense wires, and avalanches occurred within a few wire diameters. The movement of the electrons and ions from the avalanche induced a signal on the sense wire. During the analysis, the time of the signal relative to the muon trigger was converted into a distance using space-time relationships (STRs). Each TEC had 24 sense wires at 0.2 cm pitch. Thicker shield wires were placed between the sense wires to ensure that the avalanches remained localised.

The electric field in the drift volume was nominally uniform with a strength of 16 V/mm. However, there were non-uniformities of up to 10% due to the field at the sense planes leaking into the drift volume, and by $\sim 1\%$ due to modular interference[78]. Ultimately the non-uniformity in the electric field is addressed by iterating the STRs using real tracks; the analysis and calibration of the TECs is described in Appendix G.

[18]This gas was chosen since it was already in use in the drift chambers. See Section 2.8 for the advantages of dimethyl ether.

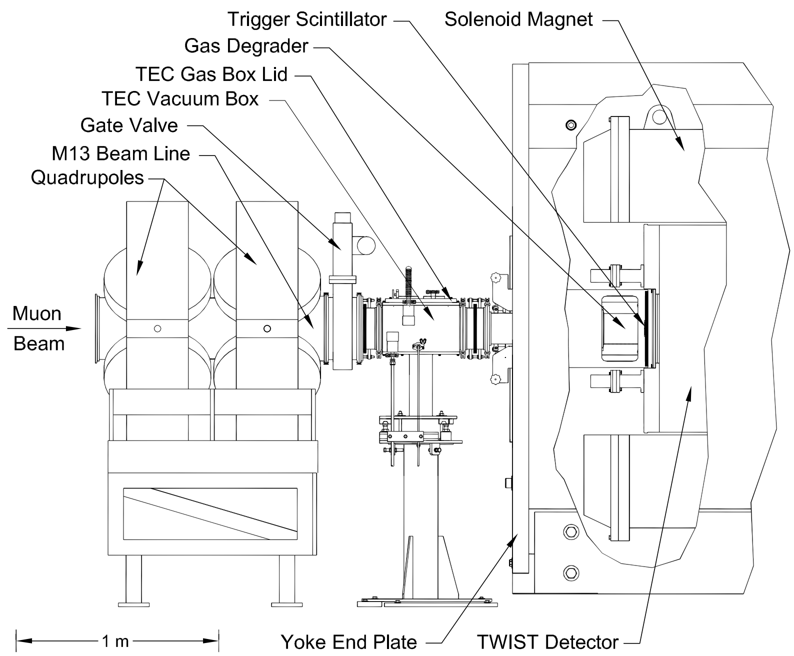

Figure 2.4: Positioning of the time expansion chambers relative to the M13 beam line and solenoid[78].

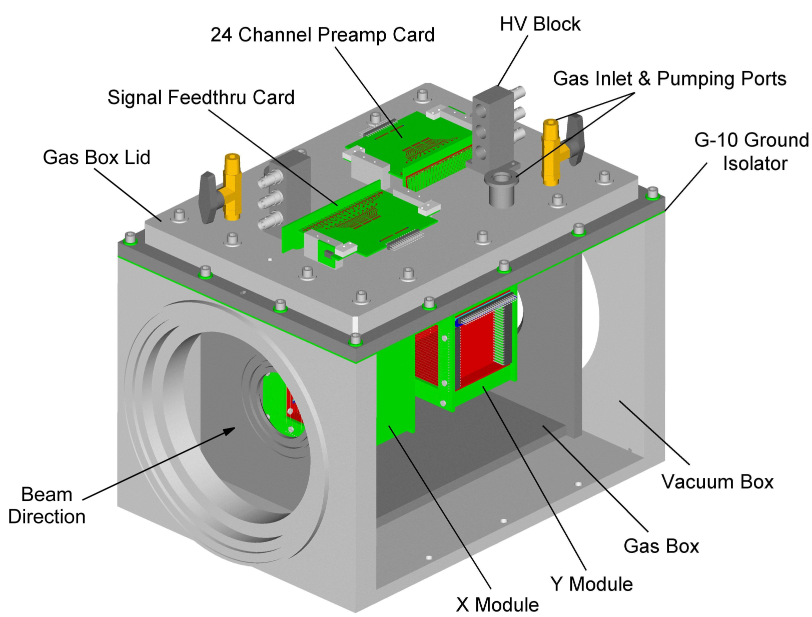

Figure 2.5: Location of the time expansion chamber modules and their electronics inside the gas box[78].

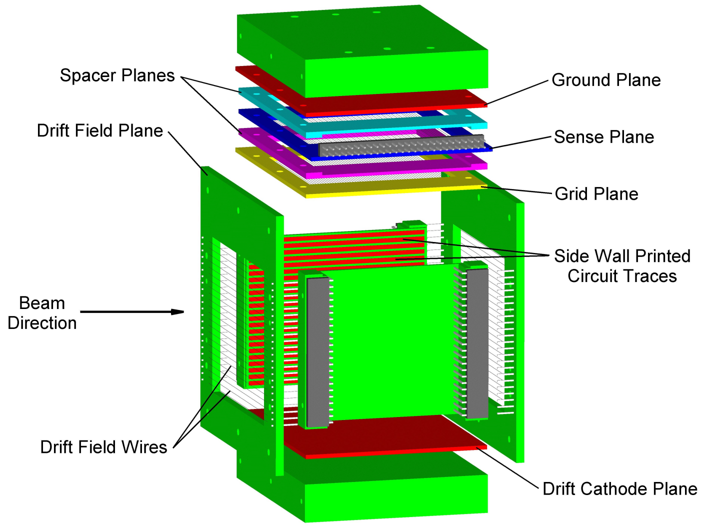

(a) Components of a single time expansion module[78].

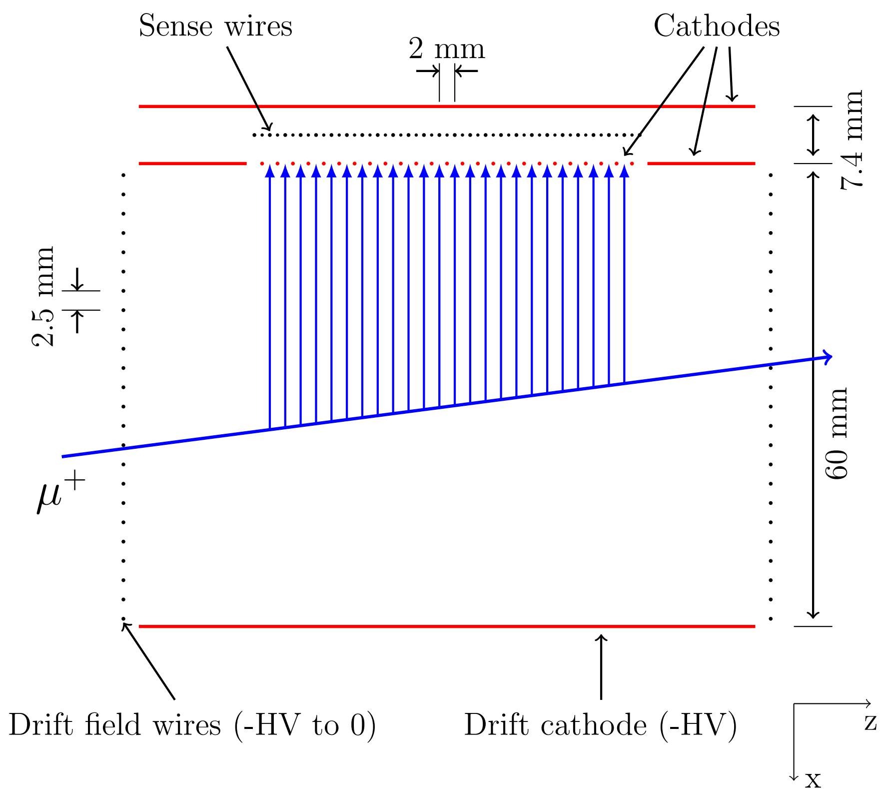

(b) Cross section through a single time expansion module. The initial ionisation is indicated by the arrow through the entire module. The other vertical blue arrows show the direction of the drift electrons[81].

Figure 2.6: An individual time expansion chamber module.

2.4 Quadrupole steering

The muon beam from M13 used to be approximately 1 cm too high at the TECs, with an average y-angle of about 20 mrad; this resulted in a larger depolarisation than the design specification. The deflection was blamed on interference between the quadrupoles Q3, Q4, Q5, and the solenoid's fringe field[82]. The dipole magnets only steered horizontally, so they were not able to correct for this deflection. A major improvement for this measurement was a correction of the deflected beam by applying additional currents to the poles of the M13 quadrupole magnets[19].

The left hand side of Fig. 2.7 illustrates the behaviour of an unmodified vertically focusing quadrupole magnet. In the vertical direction the charged particles are focussed, and in the orthogonal horizontal direction the particles are defocussed. If the poles of the quadrupole are asymmetrically excited using additional current sources then the field zero is shifted to a new location, and a charged beam is both focussed and steered[84]. For example, the right hand side of Fig. 2.7 shows that current sources added to two of the poles results in a horizontal steering force.

There are theoretical calculations of the steering effect[84]. However, since the experiment is able to precisely measure the muon beam after the quadrupoles and before the detector's fringe field, no theoretical predictions were made. The beam was simply observed with the TECs while increasing currents were applied to the pole(s) of the accessible quadrupole magnets.

Five power supplies were available to be placed over the poles of Q4, Q6 and Q7. Horizontal steering used Q6 and the dipole magnet B2. The magnets Q4 and Q7 were used to steer vertically, with Q4 providing most of the steering. The response of the beam's average position ($\langle x \rangle$,$\langle y \rangle$) and angle ($\langle \theta_x \rangle$, $\langle \theta_y \rangle$) were determined by steering with each element separately. While attempting to steer primarily in the x or y direction, all four parameters of the beam were changed. For example, Fig. 2.8 shows that when primarily steering horizontally, there is a second order beam movement in the vertical direction.

Relationships such as the example in Fig. 2.8 were determined for each quadrupole. Fortunately the response was almost always linear, and assuming the beam's response

[19]Using quadrupoles to deflect a charged particle beam is not a new technique. Prior to installation in M13, it was already in use in the TRIUMF proton beam line. Also new beam lines such as μE4 at PSI are being constructed with quadrupole steering as part of the design[83].

didn't change as the beam was steered, the results were combined according to

$$\begin{pmatrix} \langle x\rangle' \\ \langle y\rangle' \\ \langle \theta_x\rangle' \\ \langle \theta_y\rangle' \end{pmatrix} = \begin{pmatrix} \langle x\rangle \\ \langle y\rangle \\ \langle \theta_x\rangle \\ \langle \theta_y\rangle \end{pmatrix} + \begin{pmatrix} r_x^{B2} & r_x^{Q6} & r_x^{Q4} & r_x^{Q7} \\ r_y^{B2} & r_y^{Q6} & r_y^{Q4} & r_y^{Q7} \\ r_{\theta_x}^{B2} & r_{\theta_x}^{Q6} & r_{\theta_x}^{Q4} & r_{\theta_x}^{Q7} \\ r_{\theta_y}^{B2} & r_{\theta_y}^{Q6} & r_{\theta_y}^{Q4} & r_{\theta_y}^{Q7} \end{pmatrix} \begin{pmatrix} \Delta B2 \\ I_6 \\ I_4 \\ I_7 \end{pmatrix},$$

where the primed quantities are the final beam parameters, $\Delta B2$ is the change in B2's magnetic field, and the I terms refer to the additional currents applied to the quadrupoles. The r quantities are the gradients of relationships such as Fig. 2.8, where the subscript refers to the beam parameter and the superscript refers to the steering element. For example, r_x^{Q6} and r_y^{Q6} are the gradients in Fig. 2.8. The beam tuning criteria will be described later.

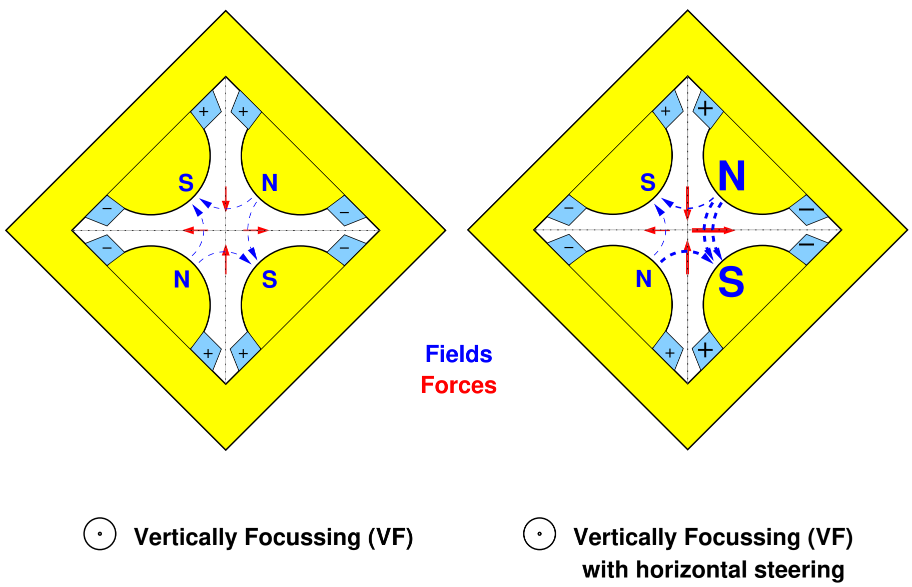

Figure 2.7: An unmodified vertically focusing quadrupole is shown on the left. The same quadrupole with asymmetrically excited poles is shown on the right. The modified quadrupole both focusses and deflects the beam.

2.5 The detector

Prior to this measurement, the TWIST detector had been in operation for four years, and has been described in detail elsewhere[59, 77, 85]. A schematic of the apparatus is shown

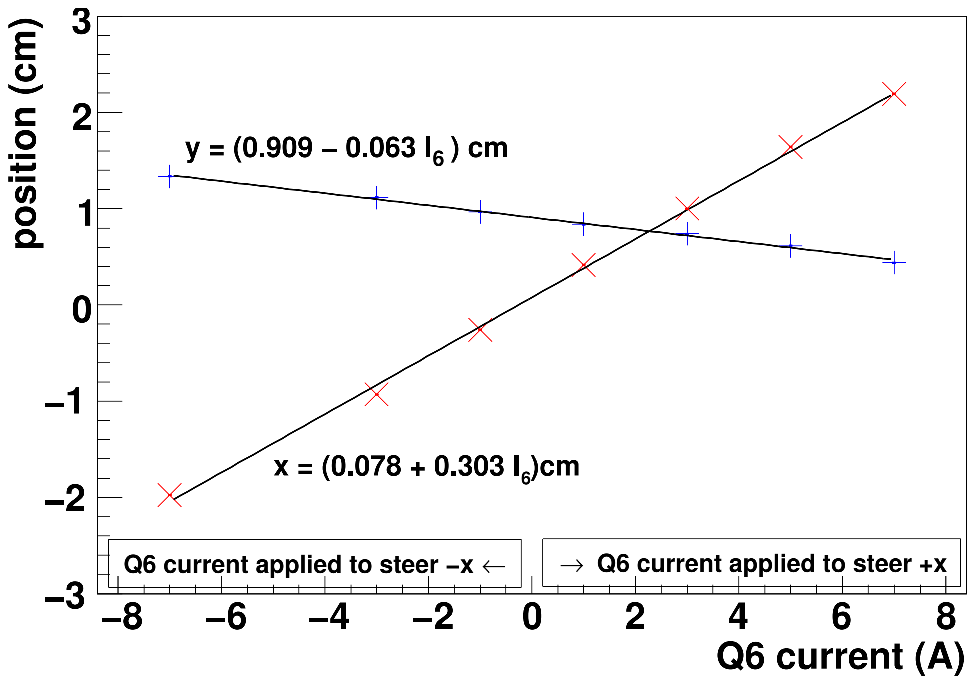

Figure 2.8: The response of the muon beam to changes in the horizontal steering of Q6.

in Fig. 2.9. The essential features are a superconducting magnet with an iron yoke, a stack of proportional and drift chambers, and a metal foil at the centre. For simplicity, the figure does not show the material immediately upstream and downstream of the chamber stack. Muons entered the detector and were stopped in a high purity metal target. The decay positrons spiralled in the magnetic field, leaving ionisation in the chambers. This allowed the positron's position to be measured and its trajectory to be reconstructed, from which the initial energy and angle were inferred. The spacing of the chambers has been changed for this measurement, which has helped to resolve wavelength degeneracies.

2.6 Coordinate system

The longitudinal direction, z, is along the beam line axis shown in Fig. 2.9. The coordinate z is positive in the direction of the beam ("downstream"), and $z = 0$ is defined to be half way between the wire planes of the central proportional chambers. The y-direction is vertical. As will be described in Section 2.8, the wire chambers are rotated by 45° with respect to the vertical, and this rotated system is described by u and v coordinates.

In terms of decay positrons, θ is measured with respect to the z-axis, so that $(0 <$

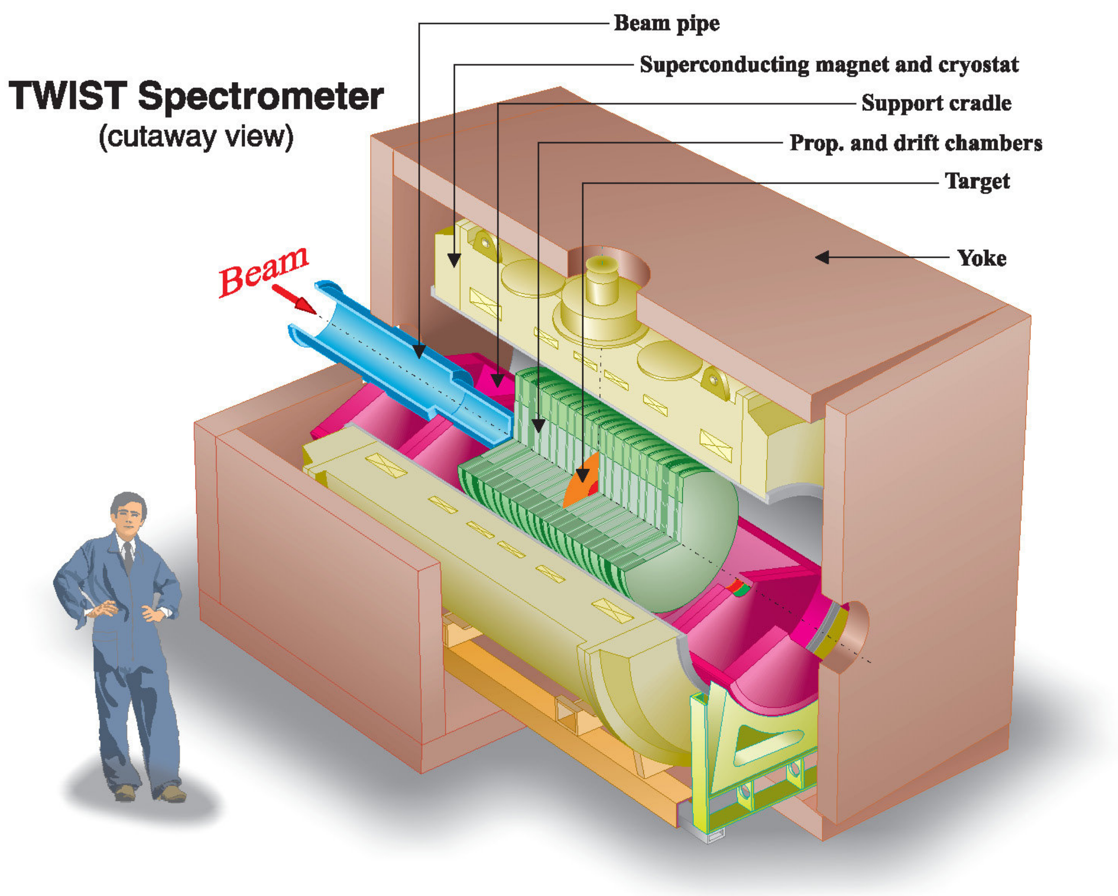

Figure 2.9: Schematic of the TWIST spectrometer.

$\cos\theta < 1$) corresponds to downstream decays, and ($-1 < \cos\theta < 0$) to upstream decays.

2.7 Solenoid magnet

The solenoid magnet and surrounding iron yoke are shown in Fig. 2.9. The magnet is a liquid helium cooled superconductor that is used in persistent mode; it was manufactured by Oxford Magnet Technologies Limited (UK), and was originally part of medical apparatus used for magnetic resonance imaging[20]. The magnetic field was necessary for the following reasons:

1. Inside the solenoid, where the field was strong (nominally 2 T), uniform, and lon-

[20]The magnet is no longer at TRIUMF. In September 2008 it was transported to Michigan State University, home of the National Superconducting Cyclotron Laboratory.

gitudinal (z), the decay positrons spiralled, with the radius and pitch angle related to the positron momentum and decay angle.

2. While muons were inside the target, the strong field held their spin against depolarisation due to nuclear dipole moments. The field also suppressed depolarisation due to muonium formation in the gas immediately before the target, although muons that stopped in the gas were removed in the analysis.

3. In the fringe field region upstream of the detector's chambers, the field radially focussed the muon beam, which maintained a high polarisation with respect to the z-axis.

The solenoid had a bore of diameter 1.05 m, and a length of only 2.23 m, resulting in significant departures from uniformity. For this reason the magnet had doors of high permeability steel at both ends (visible in Fig. 2.9), at which the field lines must be close to perpendicular, resulting in increased uniformity through the central region. The thickness of the doors (8 cm) was carefully chosen to maximise uniformity. Each door had a 40 cm diameter circular hole that detracted from the field uniformity[21]. The doors were part of a custom built iron "yoke" with square cross-section of side 3 m and thickness 0.20 m, which provided an easy return path for the magnetic flux. The yoke limited the solenoid's fringe field in the z-direction, which minimised interference with the beam line, and allowed the experiment to measure the muons after the beam line but before the magnetic field.

The strength of the field dictated the radius and angle that could be reconstructed. The nominal field strength at the centre was $|B| = 2\,\mathrm{T}$, but data were also taken at $B = 1.96\,\mathrm{T}$ and $B = 2.04\,\mathrm{T}$ as a consistency check. The field was continuously monitored using an NMR probe just outside the tracking region, and it was found to decrease by 0.02 mT over a period of approximately 3.5 months[86]. The analysis corrected for the field decrease by scaling the field map for each data set. Note that the solenoid required about 8 hours to ramp from 0 T to 2 T, and from 2 T back down to 0 T.

The magnetic field was measured with a specially constructed Hall probe apparatus. A simulated field was produced using the `Opera` software package[87], and this simulated field was used for the analysis. Further detail can be found in Appendix D. The accuracy of the `Opera` field map will contribute to the largest systematic uncertainty for $P_\mu^\pi \xi$.

[21] The extent of the muon beam is $\sim 5\,\mathrm{cm}$, yet the holes in the door are 40 cm. There are two reasons for this: the trigger and light guides required a significant amount of space, and the muon beam will be less sensitive to imperfections in the doors' steel if it is further from the hole's boundary.

2.8 Wire chambers

The 120 cm detector stack included 44 multiwire drift chambers and 12 multiwire proportional chambers, which are shown in Fig. 2.10, and will now be referred to as DCs and PCs. The construction of the DCs and PCs was very similar, and a chamber is shown schematically in Fig. 2.11. Wires of diameter 15 μm were placed at pitches of 0.4 cm (DCs) and 0.2 cm (PCs), with a total of 80 wires per plane for the DCs, and 160 wires per plane for the PCs. All of the DC wires were individually instrumented, but the outer wires of the PCs were joined together in groups of four and then read out[22]. The cathodes (which served as gas windows) for each chamber were separated by 0.4 cm, and constructed from 6.35 μm thick Mylar foil that was aluminised on both sides to allow conduction[23]. A pair of proportional chambers was just 1×10^{-4} radiation lengths in thickness[88]; the low mass was important to minimise multiple scattering and energy loss, which ultimately improved the reconstruction of the positron trajectories.

The DCs used dimethyl ether (DME) gas, and the PCs used a mixture of CF_4 and isobutane. The DME had the advantage of a high number of clusters per cm ($30\,\mathrm{cm}^{-1}$), small Lorentz angle between the electric field and electron drift velocity (5°), and a low drift velocity (2.5 cm/ μs) that resulted in DC response times up to 1 μs[89]. These properties allowed a high precision position measurement. The PC gas was chosen to minimise response time ($\lesssim 20$ ns was achieved), allowing separation of particles in time, and muon/positron identification based on pulse width. Proportional chambers were chosen over scintillators to minimise material, cover a large area, and operate efficiently in the experiment's strong magnetic field. Unfortunately a proportional chamber could not be used as the trigger since the time resolution was insufficient to separate cloud muons from surface muons; instead a dedicated trigger scintillator was used, which is described in Section 2.9.

A mixture of helium and nitrogen in the ratio 97:03 was continuously flowed at 1 litre/min into the volume between the chambers. The continuous flow was necessary to prevent contamination of the He/N_2 mixture by diffusion of air or chamber gasses into the volume. The helium minimised material and the nitrogen prevented sparking[24]. A

[22]The availability of TDCs (time-to-digital converters) limited the instrumentation of the PCs. The chosen instrumentation was adequate since the spatial distribution at the PCs was only used to select muons close to the target, and these target PC wires *were* individually instrumented.

[23]The Mylar foils were doubly aluminised regardless of whether they were exposed to chamber gas or the helium/nitrogen mixture. This reduced static charge build-up.

[24]The cross section for electrons in pure helium is very small, allowing a long mean free path. Hence the electron can accelerate and ionise, eventually causing a spark. The nitrogen significantly reduces the

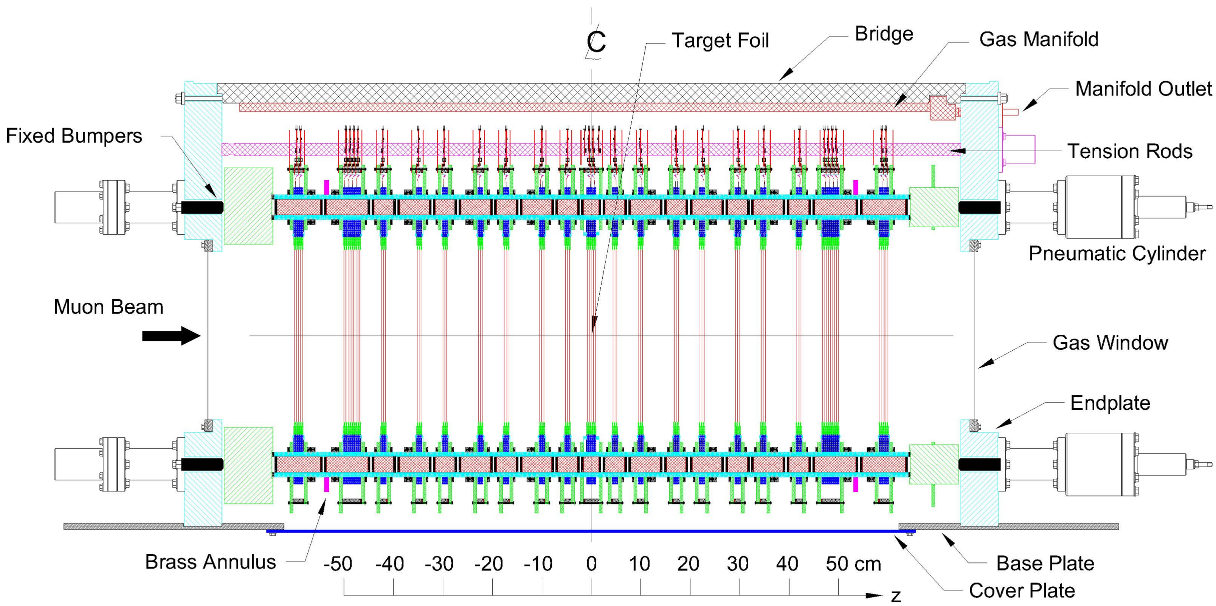

Figure 2.10: Side view of the detector. A planar array of drift and proportional chambers are symmetrically placed about a thin metal target foil.

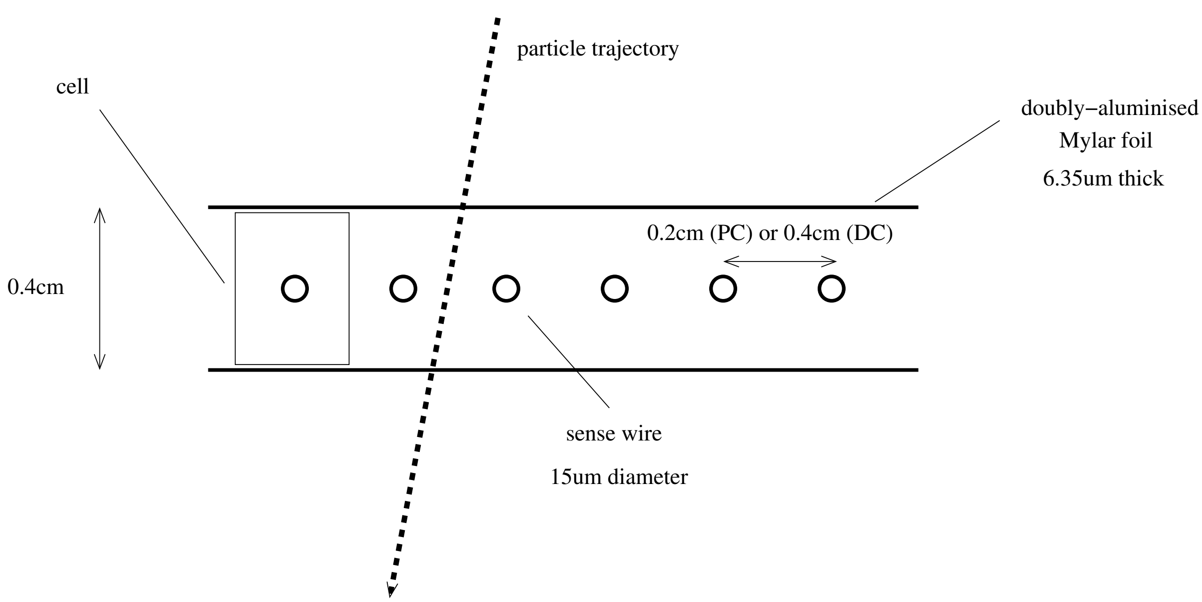

Figure 2.11: Schematic of a wire chamber. Both the DCs and PCs have cathodes of Mylar coated with aluminium; these also act as gas windows. The cathode-to-cathode distance was always 0.4 cm.

continuous gas flow of 20 cc/min into each chamber was provided to dilute contaminants due to He/N_2 diffusion through the cathode foils and outgassing from chamber materials. A pressure control system adjusted the output from the chambers to maintain the differential pressure across the cathode foils to less than $\pm 2\,$mTorr[90].

The chambers were assembled into modules with shared cathodes. In the simplest case two planes were orthogonal with a single shared cathode. These two plane modules were used in the region $-45.0\,\mathrm{cm} < z < 45.0\,\mathrm{cm}$ in Fig. 2.10, and made up the "sparse stack". There were also two "dense stacks" of eight DCs, where the seven internal cathodes were shared. At the beginning and end of the detector there were four plane modules of proportional chambers. The module containing the stopping target used the target foil as a shared cathode. For all modules the wire planes were inclined at 45° with respect to the vertical, and the rotated system was described by u and v coordinates, rather than x and y. The rotation allowed the electronics to be placed at the same position on each plane.

For this measurement the voltages on the two PCs immediately before the stopping target were reduced. The resulting pulse widths were sensitive to the muon's energy loss, allowing muons that stopped before the metal target to be rejected; see Section 3.3.3.

The detector was symmetric in z about the central muon stopping target (see Fig. 2.10). The modules were positioned in z using ceramic spacers provided by the experiment's Russian collaborators. The material was "Sitall CO-115M", which had a coefficient of linear thermal expansion of just $\sim 1 \times 10^{-7}$(dL/L)/°C[77]. For comparison, the coefficients of diamond and carbon steel are 1×10^{-6}(dL/L)°C and 11×10^{-6}(dL/L)°C respectively. The Sitall material was strong and had surfaces that were flat and parallel to $< 0.5\,\mu$m[77]. The spacers were 0.4 cm for the chambers, and 2.0 cm or 4.0 cm for spacing between the modules. The z position was maintained by four pneumatic cylinders, which are indicated in Fig. 2.10, and these exerted a force of 1470 N on the chamber stack. The length of the detector assembly was measured "with a precision considerably better than $50\,\mu$m" so that the "relative position of each 4 mm Sitall in the stack is known with a precision of a few microns"[77].

mean free path, and therefore minimises sparking.

2.9 The upstream "beam package"

The beam line vacuum extended into the iron yoke (see Fig. 2.12), and the first material traversed by the muons was the upstream window of the gas degrader. The other components in Fig. 2.12 made up an upstream "beam package". The 20.4 cm gas degrader was filled with a helium / carbon dioxide mixture; the gas ratio was automatically adjusted to stabilise the muon stopping distribution, in response to ambient pressure changes. The original design specification estimated the helium and carbon dioxide should have a 50:50 volume ratio. In the winter of 2006 the degrader typically contained $\approx 35\%$ CO_2 by volume, and in the summer of 2007 the fraction was $\approx 43\%$.

Figure 2.12 shows a central muon scintillator, which provided the trigger for the experiment, and an annular positron scintillator, which was used as part of the wire time offset calibration. The muon scintillator was a circular plastic scintillator of radius 3 cm, thickness 240 μm, wrapped in aluminised Mylar. The muon scintillator was read out with two light guides (M1 and M2), and the coincidence of M1, M2 and M1+M2 was used as the trigger[25].

The annular scintillator was expected to establish the decay positron track time. However, the helix reconstruction software was able to establish this time itself. Therefore the only use of the annular scintillator was in the calibration of the wire time offsets. The scintillator also shielded the upstream half of the detector from backscatters.

Lastly, Fig. 2.12 shows a "filmstrip degrader". This allowed up to 0.10 cm of Mylar to be put in the path of particles. For nominal data acquisition, this degrader was set to the hole position (no material). Special data were taken with the muons stopped far upstream in the detector, and this required the use of the filmstrip degrader.

In nominal operation there is no corresponding downstream package, and this is expected to cause an asymmetry in backscatters[26]. In order to study this asymmetry, a copy of the beam package was specially constructed and placed downstream of the chamber stack for a single data set. This set will later be shown to produce a consistent $P_\mu^\pi \xi$ result.

[25]The coincidence of M1 and M2 is used to suppress false triggers due to photomultiplier noise. The sum of M1 and M2 provides a larger pulse for improved time resolution.

[26]The primary backscatter process in the experiment is a positron rebounding, following a Coulomb interaction with a nucleus. Therefore the number of measured backscatters in the experiment depends on the thickness and positioning of materials outside of the tracking volume. Secondary processes such as showering are possible, but for low energy positrons this is less of a problem.

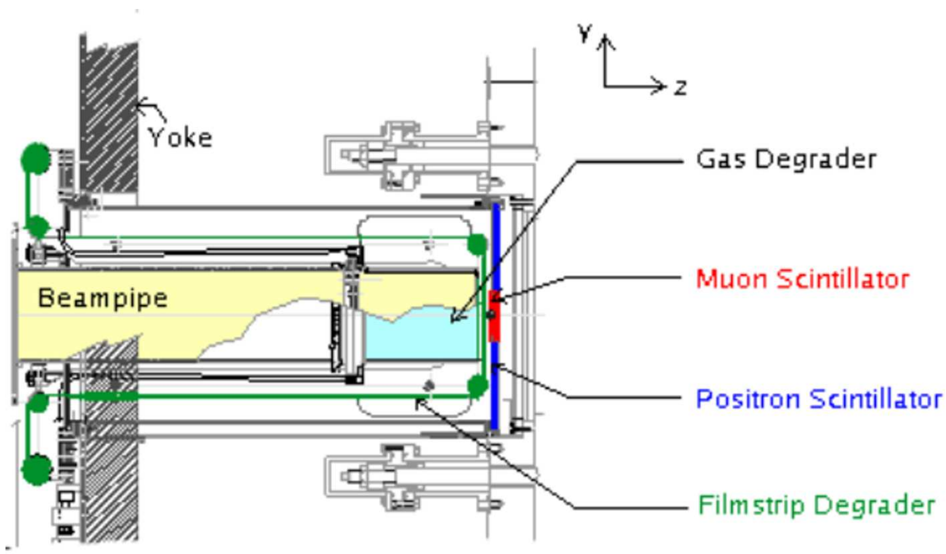

Figure 2.12: Beam package for the upstream end of the detector[59]. The beam did not encounter material until the upstream window of the gas degrader ($z = -103.2\,\mathrm{cm}$). At this location the magnetic field strength was 1.5 T.

2.10 Muon stopping target

Muons were stopped in high purity foils of silver (2006) and aluminium (2007), which both had purity $> 99.999\%$. The aluminium target was the same foil used for the previous TWIST $P_\mu^\pi \xi$ measurement[22]. The nominal thicknesses were $29.5\,\mu$m (31 mg/cm^2) for silver, and $71.0\,\mu$m (19 mg/cm^2) for aluminium. Both targets were expected to depolarise the muons by a small amount, with a simple exponential form for $P_\mu(t)$, as described in Section 1.6.3

The target region is shown in Fig. 2.13, where the metal target is seen to be part of the cathode between PC6 and PC7. The figure shows Kapton masks that avoid problems with the high electric fields where the metal target is attached to the aluminised Mylar foil.

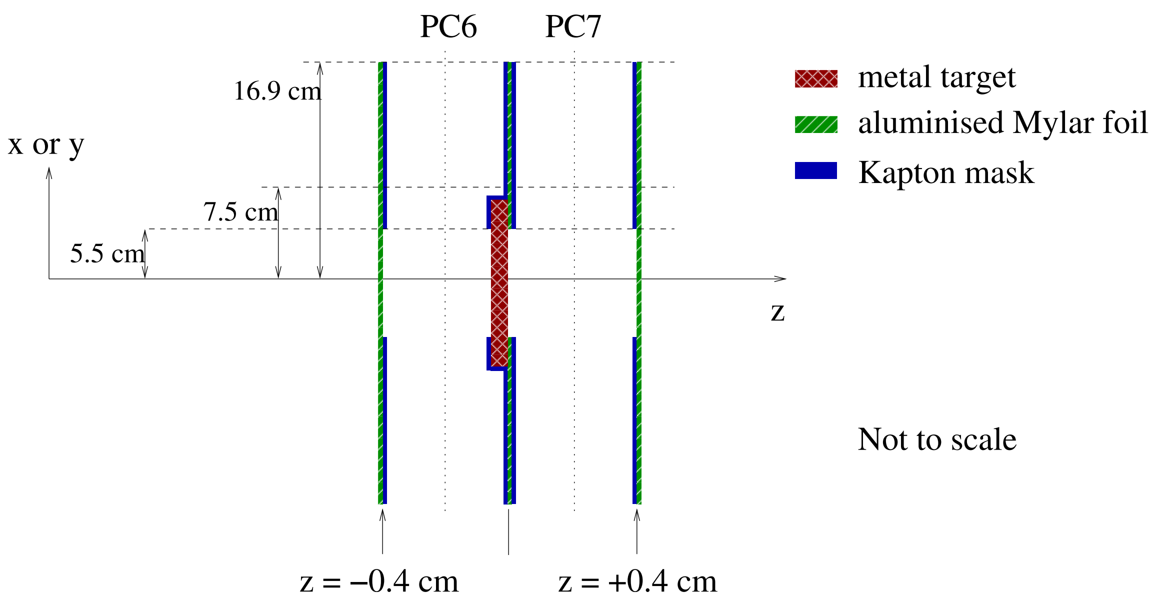

Figure 2.13: Schematic of the muon stopping target, which is part of the cathode foil for PC6 and PC7.

2.11 Muon ranging

The total range of the muons in this experiment was about 0.14 cm of water. The materials traversed by muons stopping in the target are summarised in Table 2.1. A previous TWIST analysis found that if the adjustable gas degrader in the simulation was set to the real value, then an extra 12.8 mg/cm^2 of material was needed to match the stopping distribution from data[59]. The apparatus has now been disassembled and the materials were reweighed; the upstream gas degrader window was found to have a thickness of 93.0 μm instead of 79.3 μm, and the muon scintillator had a thickness of 240 μm instead of 195 μm[27]. The simulation now only required an extra 1.9 mg/cm^2 of material, which is within the uncertainties from matching the stopping distribution (see Section 6.2.12), the unsimulated bulge in the vacuum window ($\approx$ 1.0 mg/cm^2 of extra gas), uncertainties in other material thicknesses (stopping target, cathode foil, cathode foil bulges), and uncertainties in modelling the energy loss for highly ionising particles in thin foils.

2.12 Electronics

The electronics and data acquisition system have been described elsewhere[59, 86, 88]. In summary, time-to-digital converters (TDCs) recorded the start- and stop-times for pulses on each wire. The TDCs recorded in the interval starting 6 μs before, to 10 μs after the trigger. If the TDCs were busy, or an additional trigger particle arrived, a new event was *not* started. Instead the hits were recorded and the events were later identified as having multiple trigger particles.

An individual TDC channel could suffer from nonlinearities. The internal oscillator may not have been exactly at the manufacturer's frequency, leading to a so-called "integral nonlinearity". The TDCs used by the experiment (LeCroy model number 1877 and 1877s) had an integral nonlinearity of < 25 ppm for full range[91], which is beyond the required accuracy. The TDCs had a least significant bit of 0.5 ns[92], and the differential nonlinearity between channels in one TDC was < 0.1 ns. The experiment spread signals through multiple TDCs, further reducing the effects of nonlinearities.

[27] The muon scintillator was not properly weighed before construction, and was assumed to be 8 thou instead of its correct thickness of 10 thou.

Table 2.1: Materials traversed by a muon reaching the centre of an aluminium target. The gas densities assume a pressure of 750 torr, temperature 31.0°C.

Description	Thickness (mg/cm^2)
Graphite production target	3.8[a]
Polyester window valve	0.4
Upstream gas degrader foil (“vacuum window”)	12.9
Downstream gas degrader foil	0.9
Adjustable degrader gas (He/CO_2)	3.6 (pure He) to 40.3 (pure CO_2)
Air	5.2
Muon scintillator	24.8
Scintillator wrapping (Aluminised Mylar)	8.9
Cradle window	0.9
Detector gas (He/N mixture)	12.6
PC module (4 chambers)	9.2
Dense stack (8 chambers)	13.4
Sparse stack (7 chambers)	27.9
Target PCs	4.2
Half of Al target	9.6
	138.3 to 175.0

[a] This is the average amount of material seen. The momentum resolution of the channel allows muons with a range of production target depths to be accepted.

2.13 Alignments

All components were ultimately aligned to the steel yoke, using the techniques summarised in Fig. 2.14, which will now be briefly described. Further detail can be found elsewhere[81].

The sense planes were positioned in each TEC module using alignments pins, allowing a precision of $< 200\,\mu$m. The angle of the sense planes relative to the TEC gas box was part of the wire time offset calibration, and was determined to better than 1 mrad (see Section G.4.1). To determine the positional alignment of the TEC gas box relative to the yoke, collimators were screwed onto each end of the box, and these were aligned to the

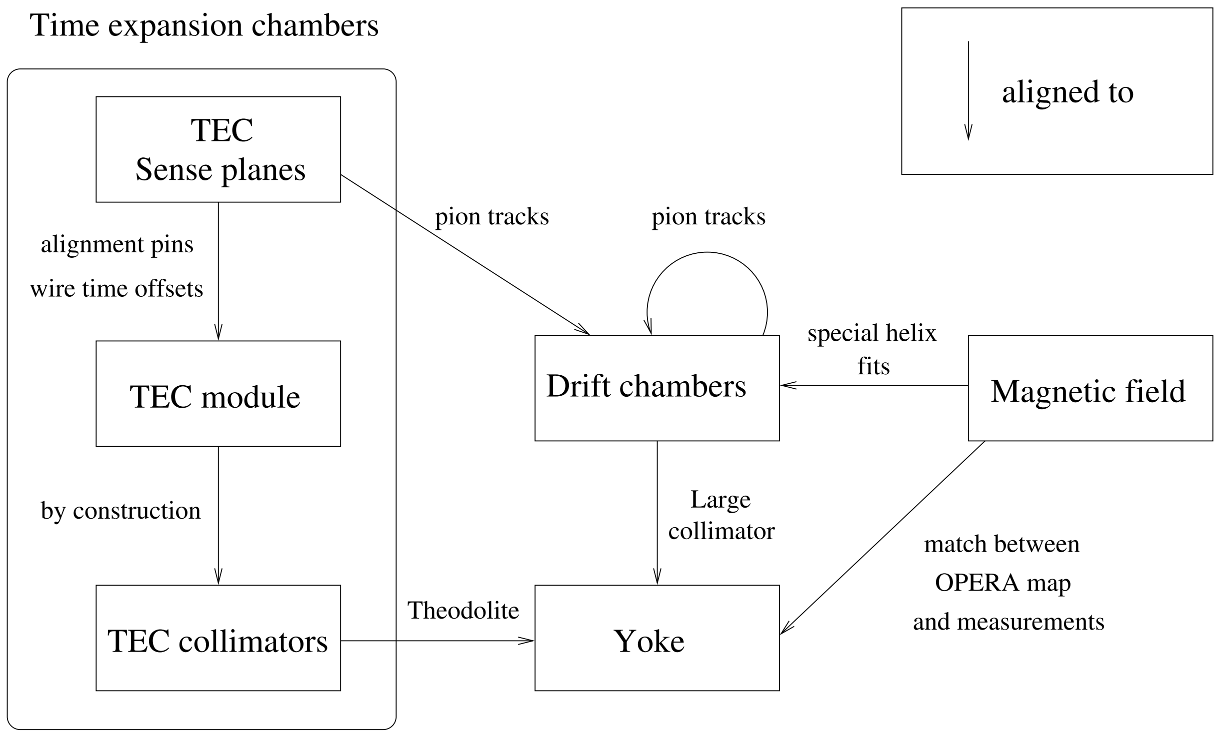

Figure 2.14: Summary of alignments relative to the steel yoke.

yoke using a theodolite; the precision was limited by the positioning of the collimators, which was better than $< 500\,\mu$m. The angle of the sense planes relative to the drift chambers was determined by turning off the magnetic field, tuning the M13 beam line to accept 55 MeV/c pions, and selecting trajectories that passed through both the TEC modules and the detector. These trajectories were well approximated by straight lines, which allowed the angle of the sense planes relative to the detector to be determined to better than about 2 mrad. Later it will be shown that these alignments were all smaller than the limitations of reproducibility when the TECs were removed/inserted into the beam line.

A new technique aligned the drift chambers to the yoke. Identical collimators with 14 holes of 2 cm diameter were placed over the upstream and downstream yoke windows. Pions with momentum 120 MeV/c were put through the collimators while the magnetic field was off, and their straight line trajectories were reconstructed with the drift chambers. The drift chambers were consistent with being at $(x, y) = (0, 0)$, to a precision of $400\,\mu$m. The angle of the stack relative to the yoke was found to be -1.4 mrad in θ_x, and -0.4 mrad in θ_y, with a precision of 0.1 mrad.

The drift chambers were aligned relative to each other using using 120 MeV/c straight line pion trajectories. The drift chamber residuals were minimised to determine the translational (u, v) and rotational alignment (about the $z-$axis) simultaneously. The precision was better than 10 μm in translation, and 0.2 mrad in rotation. The z alignment of the chambers was known to a few μm, as discussed in Section 2.8.

The magnetic field was rotationally aligned with the drift chambers by applying a special analysis to the nominal muon decay data. Simple helices were fit to the decay positron, with no energy loss or scattering, but the helix axis was allowed to have a non-zero angle with respect to the z-axis. This technique found that the magnetic field was misaligned with the drift chambers by 1.15 mrad in θ_x, and -0.33 mrad in θ_y, with precision < 0.03 mrad. The translational alignment of the magnetic field is less certain; this was determined from matching the `Opera` field to the Hall probe measurements, and had a precision of $\lesssim 0.2$ cm.

Long term stability measurements were also taken. For example, there were concerns that the box containing the TEC modules moved by a small amount. This might happen when inserting the box and pumping down the beam line to vacuum, and while turning on the 2 T solenoid. There was also the question of reproducibility over a period of several weeks. The position was therefore monitored in three ways: total station[28] measurements at intervals of several months, a laser crosshair directed onto the box while changing the magnetic field, and the same theodolite used for alignment was placed on a catwalk overlooking the beam line. The only conclusive measurement came from the total station, which was accurate to ~ 0.1 cm, and found movements of $\lesssim 0.1$ cm in x and y, but up to 0.4 cm in z.

[28]A total station measures both position and angle of a target. For this experiment a laser based total station was used. The device could only be used every few months since it was required elsewhere on site, and needed an expert to setup.

Chapter 3

Analysis

3.1 Overview

The analysis procedure is summarised in Fig. 3.1. The time and position of signals from the detector's wires ("hits") were used to identify the different particles, and then to characterise the events. Wherever possible, the particle trajectories were reconstructed. Information for every event and track were written into a ROOT tree structure database[93]. Track selection and cuts were then applied to construct a clean and unbiased decay positron spectrum.

At this point the spectrum from data could have been fit with a function that allowed for non-standard model muon decay parameters. However, this would have required a careful correction for detector acceptance and inefficiency. Instead, the analysis used a `GEANT3` simulation[94] that included the full detector response (see Chapter 4). The simulation was analysed with the same software and cuts as the real data to produce a simulated spectrum, *except* it was generated with hidden values of the muon decay parameters. The difference between the data and simulation spectra was then fit to obtain the difference in $P_\mu^\pi \xi$, ρ and δ. All systematic uncertainties were determined on the difference in the decay parameters from their hidden values. The hidden simulation parameters were only revealed after all the systematic uncertainties had been evaluated. This "blind" analysis procedure was adopted to reduce the influence of human bias.

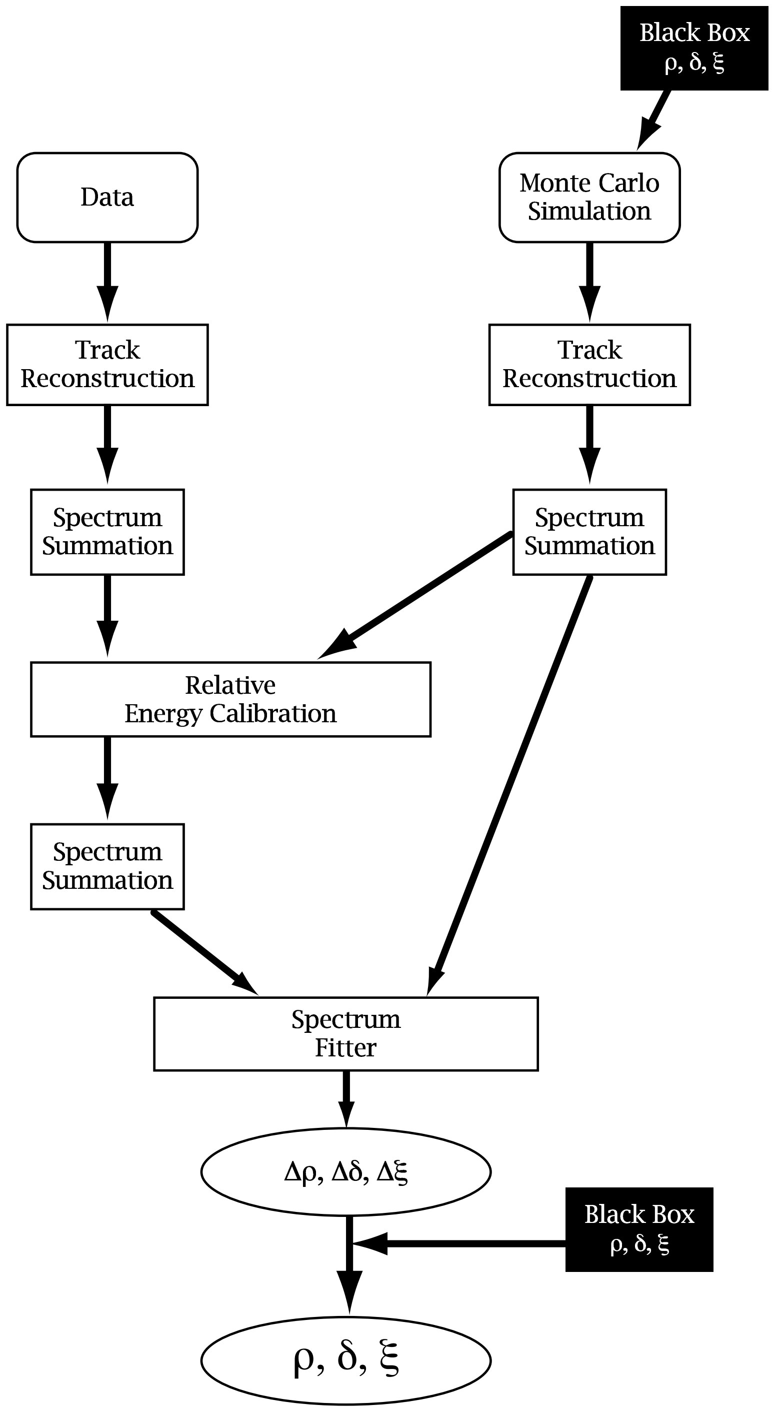

Figure 3.1: Flow diagram of the analysis procedure[85]. The "Black Box" corresponds to the hidden values of the muon decay parameters.

3.2 Track reconstruction

The collaboration used internally authored software to identify particles, classify events, and reconstruct trajectories. For this measurement almost every code module was reviewed and improved where necessary. The software will now be described, except for the time expansion chamber analysis, which appears separately as Appendix G.

3.2.1 "Unpacking"

The data acquisition system and the simulation created data files with the leading and trailing edge times of the signals from the scintillators and wires. These were first "unpacked" according to the following steps:

- Pulses where the time-to-digital converter reported an error were recovered or discarded as appropriate.
- Wire time offsets (see below) were applied, and the leading/trailing edge times were converted into a time and width.
- The trigger scintillator used a "pulse amplitude charge-to-time converter" (PACT), which converted the integrated charge to a leading/trailing edge time. This signal was converted back to an energy deposit during unpacking. This is the only place in the electronics where the pulse amplitude is used.

A new technique determined the wire time offsets, which varied from wire-to-wire due to variations in cable lengths (up to $\sim$ cms), electronics (primarily the TDCs, up to 1 ns), and discriminator amplitude time walk. A selection was made on decay positrons that triggered the annular scintillator upstream (see Section 2.9), or a newly constructed downstream scintillator, which is described in Appendix E. The time of each wire's hit relative to the trigger scintillator was histogrammed, after correcting for the time-of-flight using a simple helix fit; these histograms were then fit to determine the zero-time.

An additional analysis selected beam positrons that passed through the entire detector, including both the upstream annular scintillator and the downstream scintillator. These positrons illuminated the central wires and allowed the determination of a very precise relationship between the two detector halves.

The technique was validated using the simulation, where the wire time offsets are known to be zero. The results are shown in Fig. 3.2(a), where a wire-dependent bias is evident. The wires at the edge of the planes had poor statistics, resulting in anomalous offsets. Figure 3.2(b) shows a histogram of the wire time offsets in each detector half; the bias changes both halves by about 0.4 ns on average.

The wire time offsets for one set of the data are shown in Figs. 3.2(c) and 3.2(d). As expected, there is rich structure, and the scale of variations is much greater than the bias shown in Fig. 3.2(a).

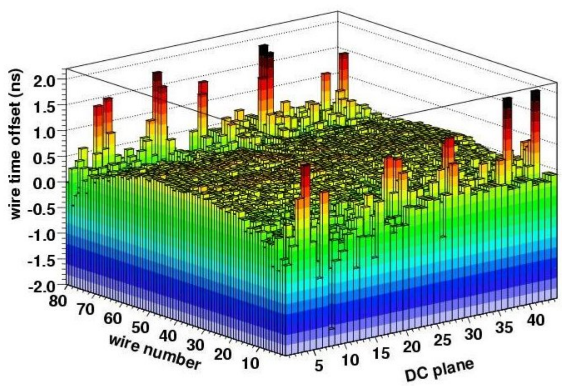

(a) Wire time offset technique applied to simulation. This shows the bias in the technique, since ideally all wires would be at zero.

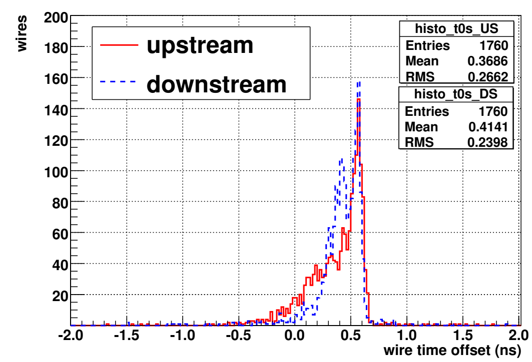

(b) Simulation wire time offsets.

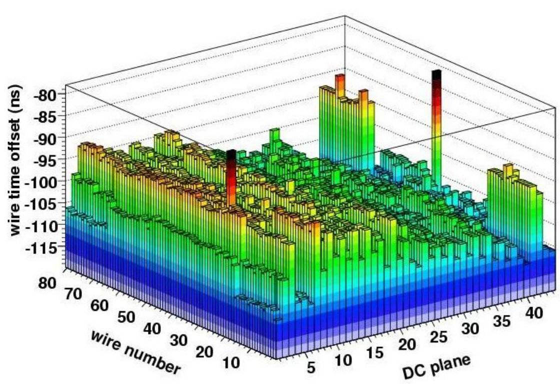

(c) Wire time offsets for the data, showing rich structure. The scale differs significantly from Fig 3.2(a).

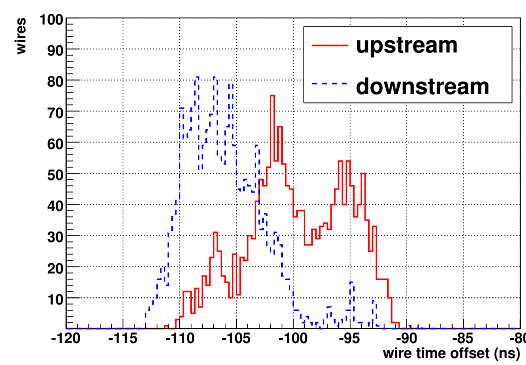

(d) Wire time offsets for the data.

Figure 3.2: Results of the wire time offset analysis, applied to simulation and data.

3.2.2 Crosstalk removal

Electronic crosstalk occurred when a large pulse induced a false smaller pulse nearby. This could occur inside a drift chamber, where an avalanche at a wire created photons, and these underwent a secondary process resulting in ionisation at a nearby wire. Crosstalk could also occur within the pre-amplifier cards, where a large pulse (*e.g.* from a muon) could induce a signal in a nearby conductive path on the pre-amplifier board. (In the time expansion chambers, described in Section 2.3, the finer spacing of the sense wires required guard wires to be placed in between, to minimise the crosstalk.)

The crosstalk pulses were identified and removed by the analysis software: if a narrow pulse (less than 50 ns wide in a DC, and 60 ns in a PC) occurred within 50 ns of a good pulse, then the narrow pulse was removed. In previous TWIST analyses, only the 10 (32) wires closest to the good DC (PC) hit were checked for crosstalk. For this measurement, it was found safe to expand the check to all wires in the same plane.

The simulation did not reproduce the crosstalk mechanisms. Therefore the crosstalk removal code was disabled when analysing the simulation. This was the only place in the track reconstruction code where the data and simulation were treated differently.

3.2.3 Windowing

A "time window" was started by a hit in a proportional chamber (PC). All subsequent hits within 1.05 μs were put into this time window. Historically the window was chosen to be 1 μs, which is sufficiently long for the ionisation in the drift chamber's slow gas to register at the wires; an extra 50 ns was later added to account for the PC rise time. The time window was intended to contain all hits of a track. If additional PC hits occurred while filling the window, a new window was not started. The window code was reviewed for this analysis; see Ref. [81].

3.2.4 Classification

The identification of the particle in a window used several pieces of information. The muons were identified from their pulse widths in the upstream PCs. Positrons from the beam line were distinguished from decay positrons since they passed through all the chambers, rather than just one half of the detector. Delta electrons were identified, as well as broken trajectories due to a large angle multiple scatter from the detector material, and backscatters from material outside of the DC region. Events were then classified according to the particles observed and their time separation. The classification code was reviewed for this analysis; see Ref. [81].

3.2.5 Pattern recognition

A helix with centre $(x, y) = (\Delta_x, \Delta_y)$ can be written in the notation

$$\begin{pmatrix} x \\ y \end{pmatrix} = r \begin{pmatrix} \cos(f(z)) \\ \sin(f(z)) \end{pmatrix} + \begin{pmatrix} \Delta_x \\ \Delta_y \end{pmatrix}, \tag{3.1}$$

where (x, y, z) is the position in space, r is the radius, and

$$f(z) = \frac{2\pi z}{\lambda} + \phi, \tag{3.2}$$

where λ is the wavelength and ϕ is the phase. The pattern recognition algorithm made a first estimate of the parameters in Eqs. (3.1) and (3.2).

A helix with $\lambda = 1$ and $\phi = \pi/4$ is shown in Fig. 3.3(a). Projecting this helix onto the $x - y$ plane allows both the radius, r, and the centre, (Δ_x, Δ_y), to be determined. If the angle between each point and the line $x = 0$ is then plotted against z, the wavelength (λ) and phase (ϕ) can be determined, as shown in Fig. 3.3(b).

A real measurement samples the helix at discrete z locations, which allows the centre and radius to be readily estimated, but results in a wavelength ambiguity. This is illustrated in Fig. 3.4(a), where the circles show the helix phase if it's assumed to be in the range $0 < f(z) < 2\pi$. The other markers are mathematically allowed solutions, of which there are an infinite number.

The projection of coordinates from the drift chambers were used in combinations of three to estimate the radius. The combination with the smallest χ^2 was kept, and the phase for each combination was determined by the procedure already described. The wavelength degeneracy was resolved since higher angle tracks registered in more than one drift cell[29]. At this stage there was no estimate made of the positron energy loss, or changes in angle due to multiple scattering.

The modules in the sparse stacks (see Section 2.8) were re-arranged prior to this measurement in order to minimise "magic wavelengths". This concept is illustrated in Fig. 3.4(b), which shows the worst case scenario where the wavelength is ambiguous.

[29]Further detail can be found in Section 4.2.1 of Ref. [95].

Wavelengths close to this situation were poorly reconstructed.

For this measurement, the existing quality cuts were upgraded to depend on track angle and radius, and the range of potential helix parameters was reduced. The changes were tuned to help situations where multiple tracks overlapped, and tracks that had smaller angles.

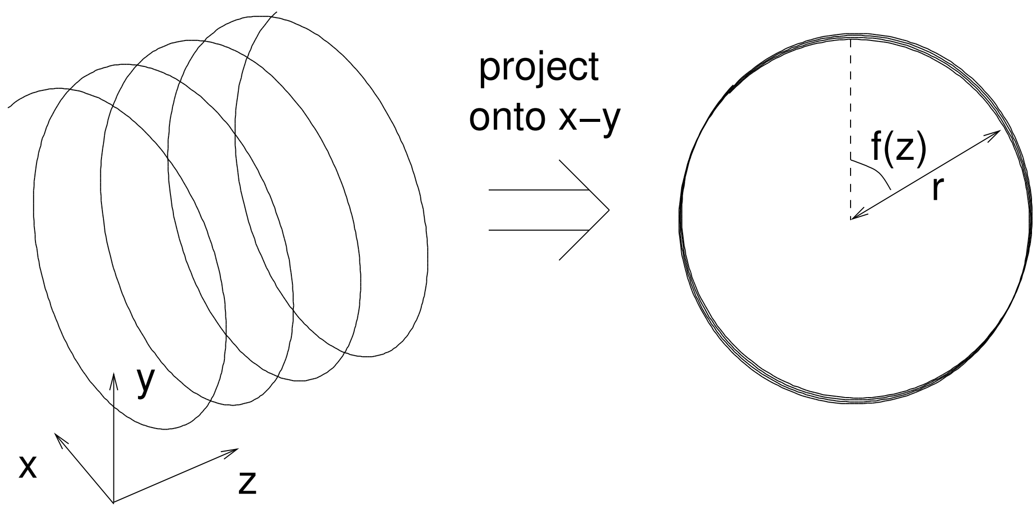

(a) Projecting onto the transverse plane to determine r and (Δ_x, Δ_y).

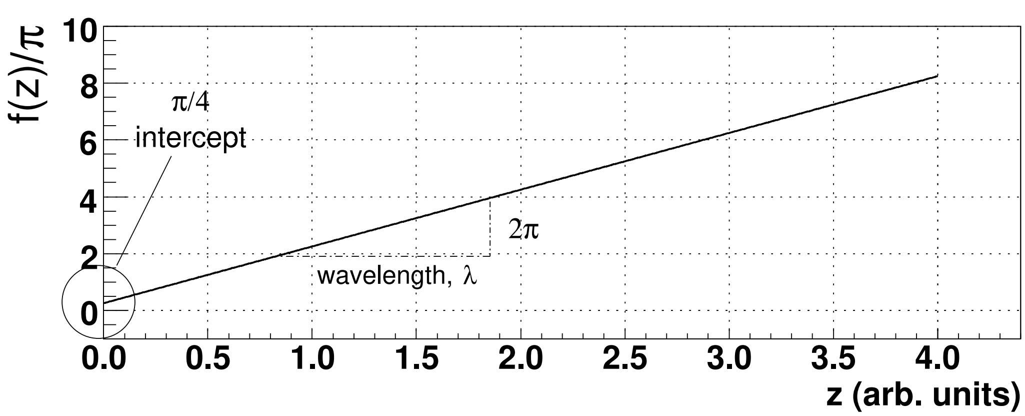

(b) Angle with respect to $x = 0$ can determine λ and ϕ.

Figure 3.3: Determining helix fit parameters.

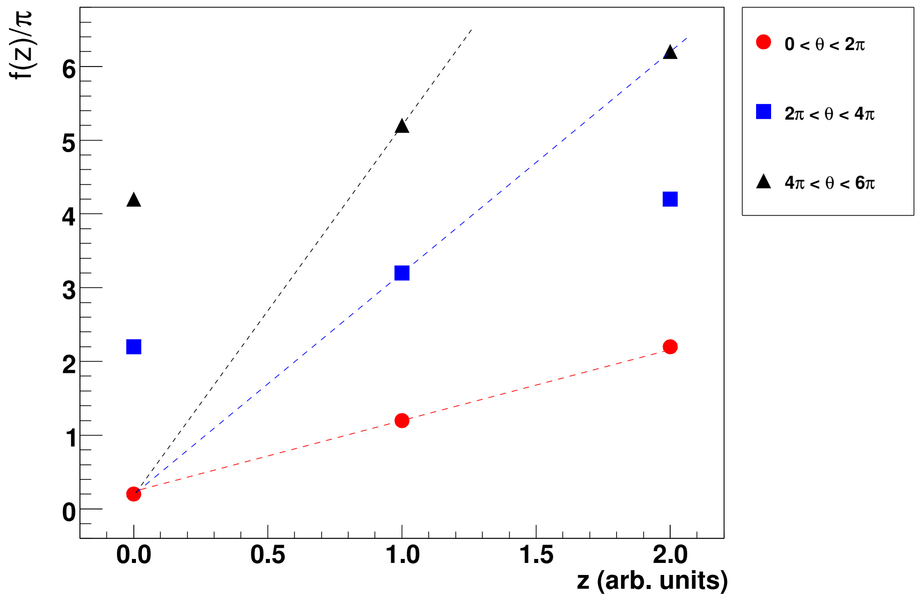

(a) There is a 2π ambiguity in determining the angle, which leads to an infinite number of possible wavelengths.

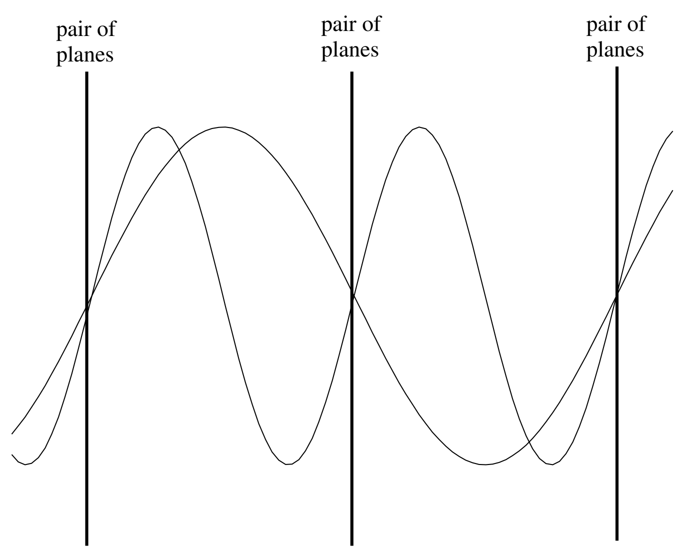

(b) The concept of a "magic wavelength" is illustrated. Wavelengths close to the spacing of the chambers in the sparse stack will have poor longitudinal momentum reconstruction.

Figure 3.4: Demonstration of the difficulties in estimating the wavelength of helical trajectories.

3.2.6 Fitting the helical trajectory of the decay positron

The pattern recognition described in the previous section gave initial estimates to the helix fitting algorithm. The fitter then "swam" helices with parameters close to those determined by the pattern recognition, and iterated using a least squares approach. For the swimming, the trajectory was divided up into segments where the magnetic field changed by less than 0.075 mT. The field was assumed to be uniform over these segments, so that the trajectory was a perfect helix arc. This approach was validated using the simulation, by comparing the true and reconstructed trajectories.

Initially a fit was made to the positions of the wires, ignoring any drift time information from the chambers. The final iterations then used the drift times[30]. There was a left-right ambiguity that was resolved iteratively, by choosing the side of the wire that best fit the rest of the trajectory at each iteration. The final fit to drift times includes energy loss and allowed for a scatter in the trajectory at pairs of planes.

The helix fitter minimised

$$\chi^2 = \sum_{\text{hits}} \frac{(d_f - d_m)^2}{\sigma_d^2} + \sum_{\text{scatters}} \frac{\theta_s^2}{\sigma_\theta^2}, \tag{3.3}$$

where d_f is the fitted position, d_m is the position from the space-time relationship, σ_d is the resolution, θ_s is the scatter angle and σ_θ is the theoretical scattering distribution width, which is set to the approximate expression for multiple scattering from the Particle Data Group[3]. The fitting algorithm established its own track time, and reported the position and momentum at the first plane that sees the positron, along with a code that described the algorithm's success. This success code is later used to select good tracks. The helix fitting algorithm is described in more detail elsewhere[85]. An example of an event with a successful fit is shown in Fig. 3.5.

Sections of the fitting code were reviewed for the current analysis. A small improvement was made to the algorithm that calculated the scatter angle. The energy loss model used by the fitter was reviewed with the conclusion that bremsstrahlung can be safely neglected, instead assuming that all the energy loss is from ionisation. There were significant improvements made to the space-time relationship and drift cell resolution, and these will described later (Sections 3.2.7 and 3.2.8).

Note that muons and beam positrons were not reconstructed since their transverse momenta (and hence their radii) were too small.

[30] The PCs were not used in the final helix fit since their resolution was limited to the wire spacing.

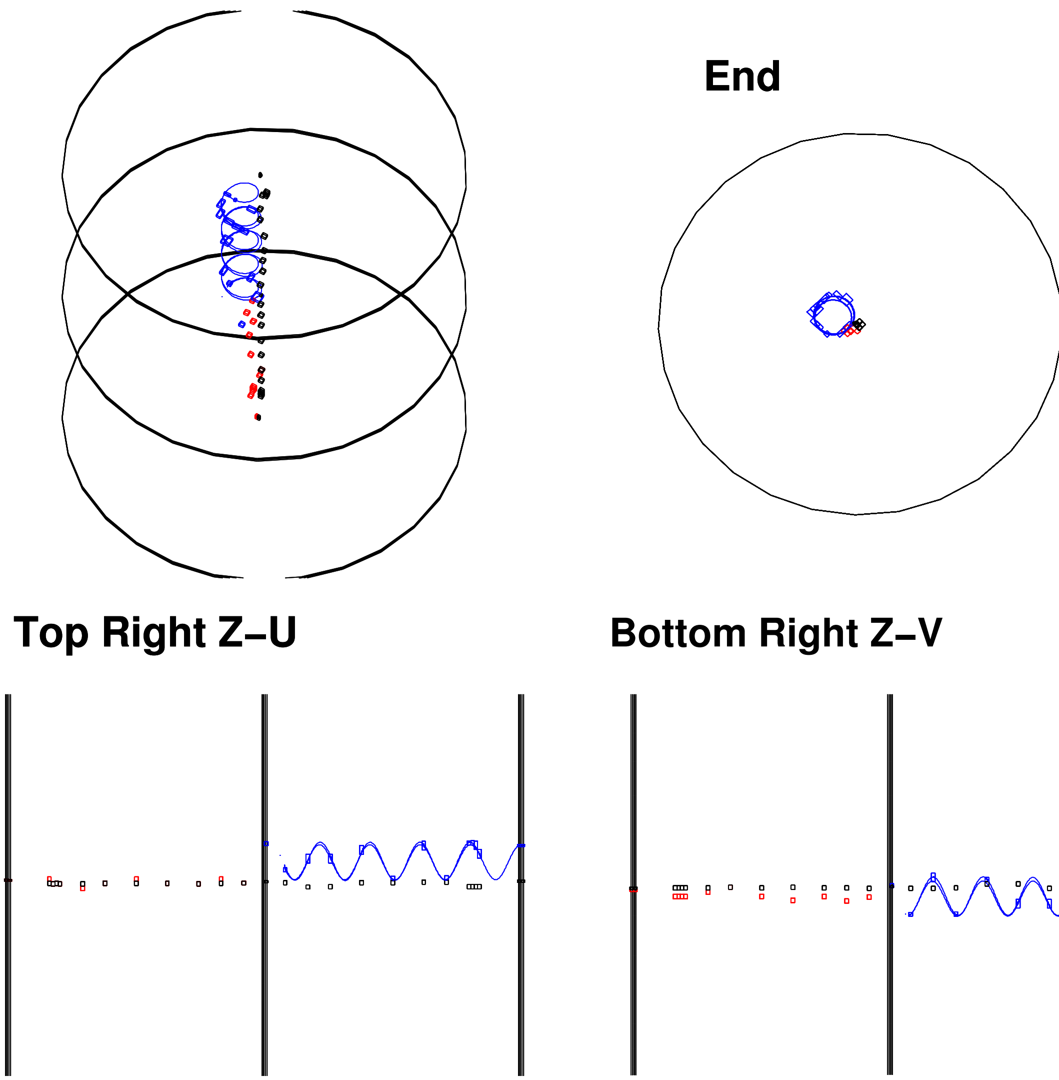

Figure 3.5: A muon (red) stops in the metal foil (centre of the Z-U and Z-V displays) and decays to a positron (blue). The black hits that span the length of the detector are a beam positron.

3.2.7 Improved drift chamber space-time-relationship

The drift chambers measured the arrival time of the ionisation at a wire, and this was converted into a distance using a drift time map, also known as a space-time-relationship (STR). The STR map for a single drift cell was generated using `GARFIELD`[96] at intervals of 20 μm. For the current analysis the STRs were refined by analysing events using the experiment's helix fitting code, and adjusting the STR at each space point to minimise the average drift time residual. This was performed iteratively, and allowed the STRs to correct for biases in the helix fitting software, and for uncertainties in the `GARFIELD` inputs such as voltage, gas densities and geometry[31] (see Section 2.8).

The refined drift cell isochrones are shown in Fig. 3.6(a), along with the `GARFIELD` STRs. The significant differences are close to the wire, and in the corners of the cells. The refined STRs reduced the average χ^2 of the positron helices, and improved the momentum bias and resolution. The method was extended to make refined STRs that were dependent on drift plane, thereby correcting for small variations in plane assembly such as wire placement and cathode-to-cathode differences. Although temperature gradients exist in the detector, the STRs for each plane were dominated by mechanical variations, not by temperature. For each drift plane there was just one cell specified; STRs depending on the region of the plane were investigated and found to be unwarranted.

Refined STRs were also produced from the simulation, and these are shown in Fig. 3.6(b). This was done independently from the data, and in this case any differences from `GARFIELD` must be due to biases in the helix reconstruction software. The simulation was then analysed using these refined STRs, allowing the data and simulation to be treated in the same way; in other words, for both cases the reconstruction biases were absorbed into the STRs. For convenience, the analysis chain is described schematically in Fig. 3.7.

[31] The previous $P_\mu^\pi \xi$ analysis[22] used `GARFIELD` STRs. The most recent ρ and δ measurements did not use refined STRs for the analysis since they were not initially available, but *did* make a correction for them at the end of the analysis.

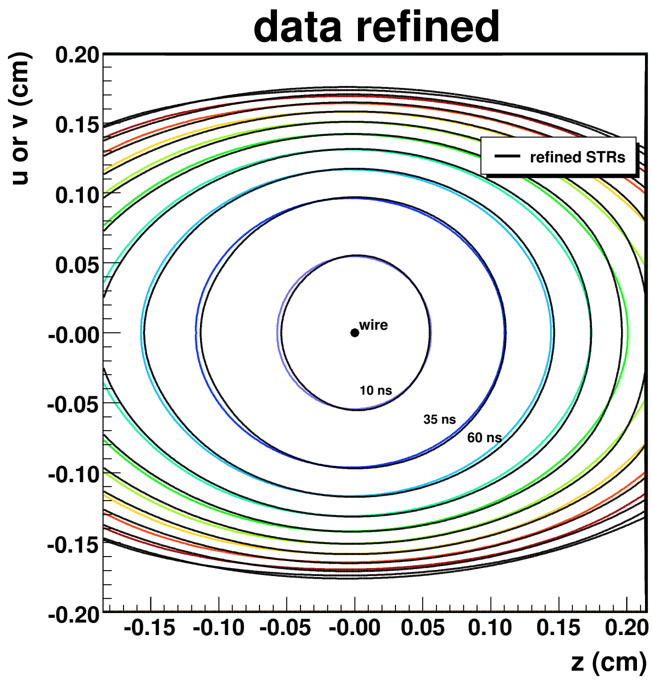

(a) STRs refined from helix fits to data.

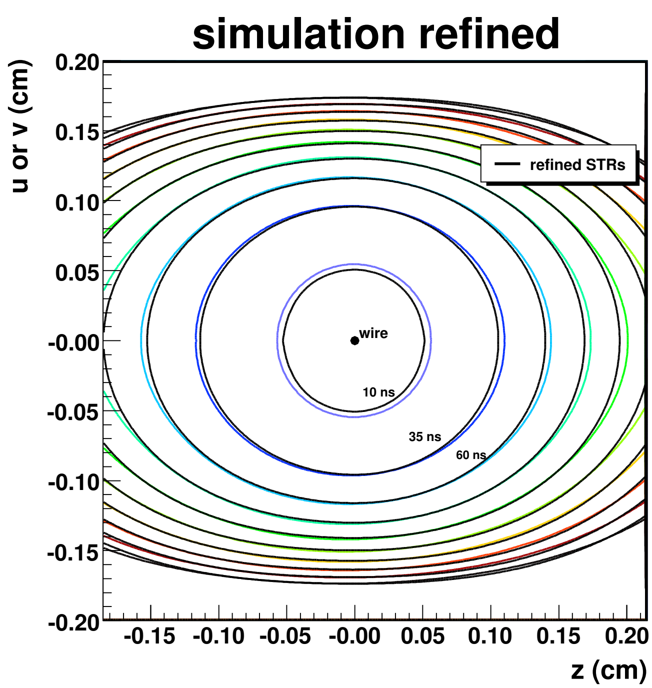

(b) STRs refined from helix fits to simulation.

Figure 3.6: Comparison of drift cell space-time-relationships from the `GARFIELD` software, and after refinement. The elliptical contours arise due to the geometry of the drift planes (see Section 2.8).

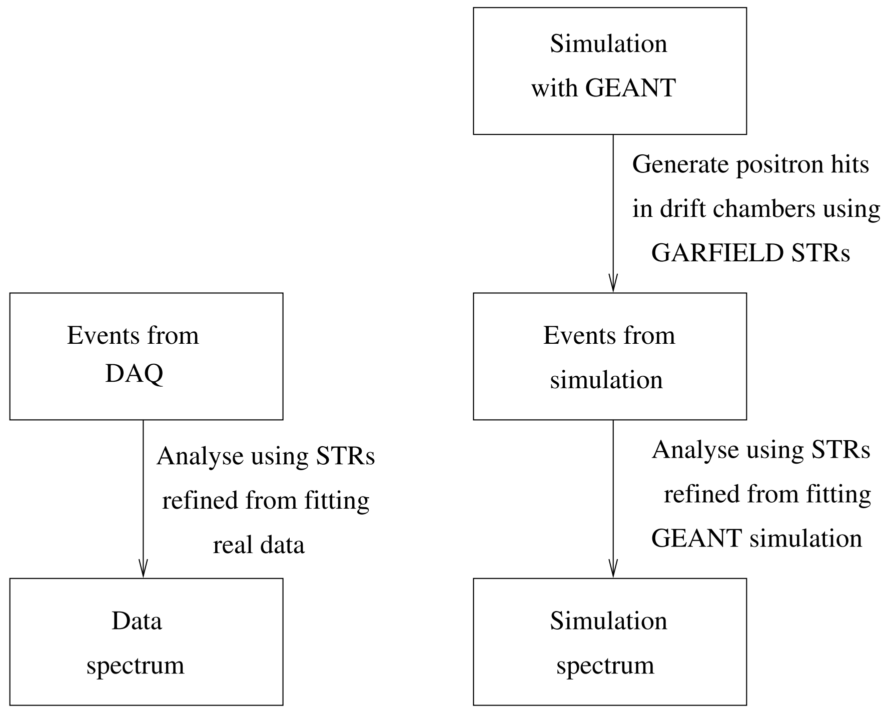

Figure 3.7: Analysis chain for using refined STRs.

3.2.8 Improved resolution function

The helix fitting routine must assign a weight to each point (σ_d in Eq. (3.3)). In the past a constant has been used, $\sigma_d = 100\,\mu$m. A study using the simulation tested many reasonable functions, and concluded that the best choice was the one shown in Fig. 3.8. This was an "effective resolution", which incorporated the effects of bias in the fitting procedure.

Figure 3.8 shows a degradation of resolution at longer times; this was due to diffusion spreading out the ionisation clusters, resulting in several smaller pulses. The baseline in the figure was expected from the timing resolution. The resolution from ionisation statistics degraded closer to the wire, but Fig. 3.8 shows that the *effective* resolution was nearly constant. This was because the "left-right ambiguity" dominated the resolution from ionisation statistics; specifically, the wires only recorded time, and this alone could not determine which side of the wire the particle went past. The left-right ambiguity was resolved by iterating the fit, keeping the side that was compatible with the rest of the trajectory. In the first iteration, ionisation with distance (x) closer to the wire than 0.1 cm was assigned $\sigma_d = 2x + \sigma(x)$, where σ is taken from Fig. 3.8. In subsequent iterations, as the ambiguity was resolved, this weighting only occurred for hits closer than 0.05 cm, 0.025 cm *etc.* The dominance of the left-right ambiguity therefore resulted in almost no sensitivity to the choice of $\sigma(x)$ below 0.1 cm.

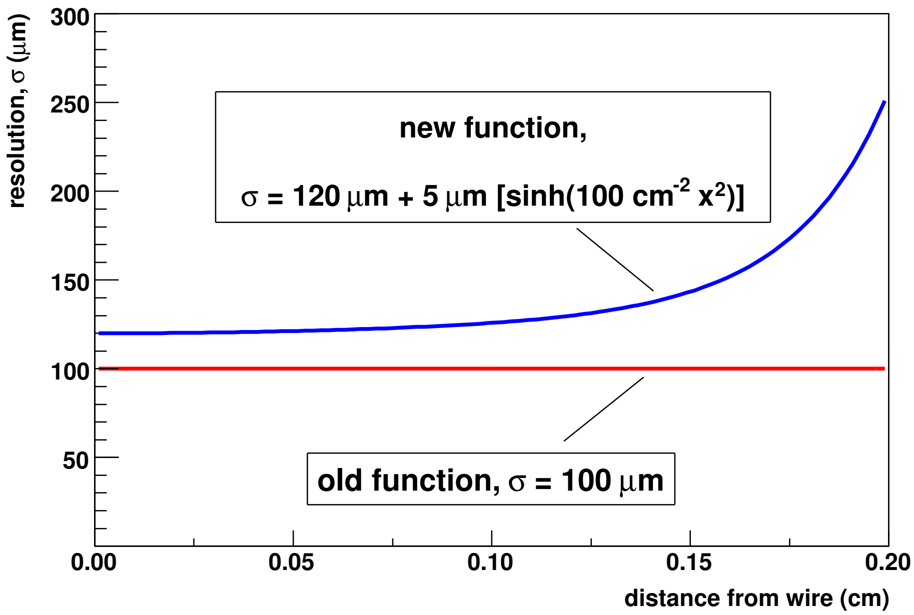

Figure 3.8: Resolution function used during the helix fitting. The resolution is worse at longer distances from the wire due to diffusion. The new resolution function is ad hoc.

3.3 Track selection and cuts

Cuts and selections were applied to produce a clean, unbiased decay positron spectrum. The analysis was the same for the data and simulation, with the exception of a time-of-flight cut that was applied in data to remove cloud muons and pions, which are not simulated. The order of these selections and their effect on statistics is shown in Fig. 3.9. The fraction of events removed at each stage was well matched in data and simulation, with about 13% of surface muon events entering the region used to extract muon decay parameters (the “kinematic fiducial”). Each of these selections and cuts will now be briefly described.

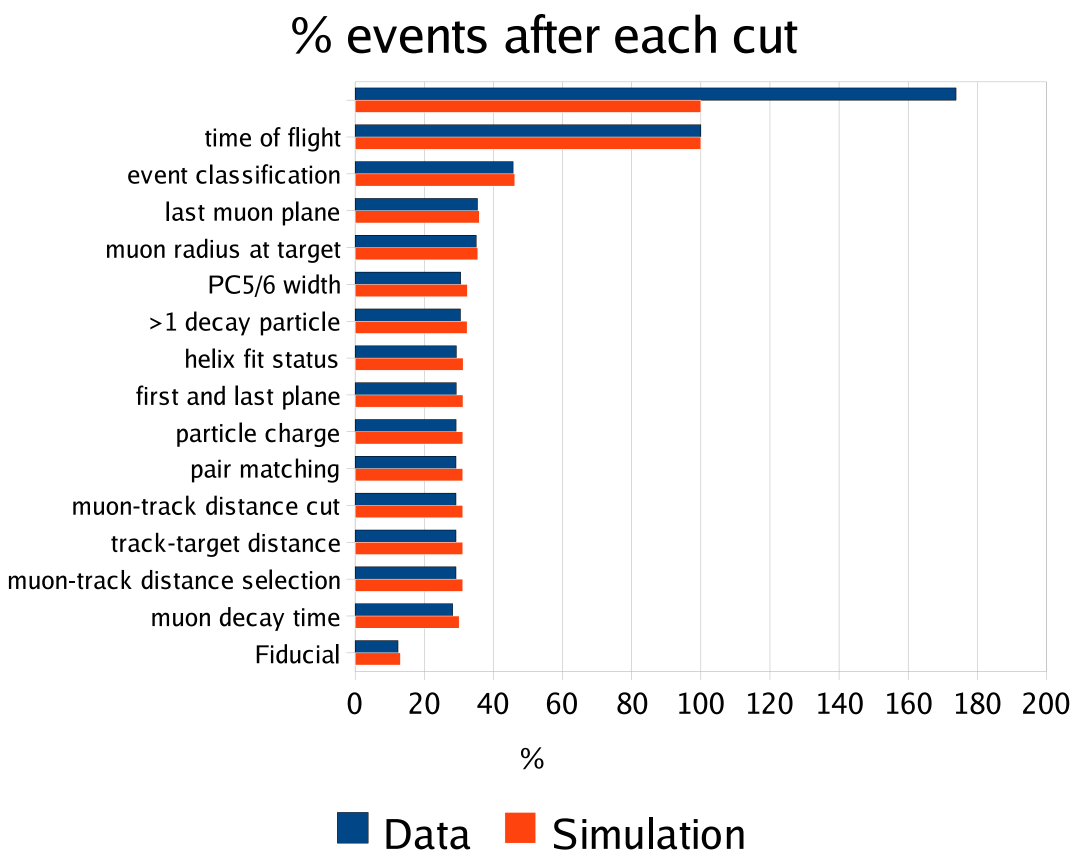

Figure 3.9: Fraction of events after each selection or cut is applied. The data has an additional time-of-flight cut to remove cloud muons and pions; for this reason, 100% in the figure corresponds to surface muons. The figure shows that after applying the fiducial cut, less than 15% of the original surface muons remain.

3.3.1 Trigger particle time-of-flight

A cut on the time-of-flight through the M13 beam line allowed heavier pions and prompt cloud muons to be eliminated (see Section 2.2.5). The cut was tuned by observing the asymmetry in the data, as shown in Fig. 3.10. This cut was not applied to the simulation since only surface muons and positrons were generated[32].

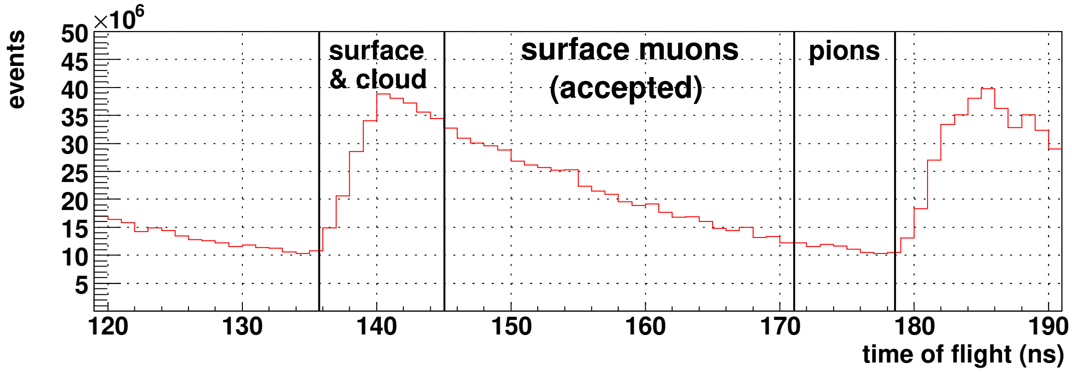

(a) Number of events for each time-of-flight.

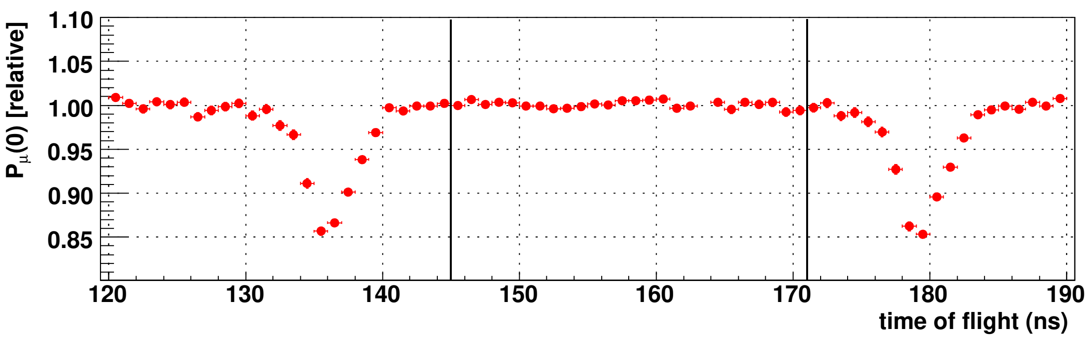

(b) Relative asymmetry for each time-of-flight.

Figure 3.10: Tuning the time-of-flight cut that selects surface muons. The surface muons are accepted, and the cloud muons and pions are rejected. These backgrounds were described in Section 2.2.5.

3.3.2 Event classification cut

The classes of events selected at this stage are shown in Table 3.1. Nearly 90% of events that subsequently pass the remaining cuts are of the simplest type. The remaining ≈ 10%

[32]The capability to simulate proton bunches on the target and the M13 beam line existed, but there was little motivation since the experiment had a high precision muon beam measurement at the end of the channel.

have a beam positron that doesn't interfere with the decay positron's reconstruction. Events with multiple muons have been removed. Note that the event classification cut includes an implicit time cut at 1.05 μs, which removes 38% of events due to the muon's lifetime.

Table 3.1: Fraction of event types in the fiducial.

Event type[a]	Description	% of fiducial events data	% of fiducial events simulation
1	μ^+ and decay-e^+, separated by $> 1.05\,\mu$s	88	87
2	μ^+, decay-e^+, beam-e^+(s), all separated by $1.05\,\mu$s	11	12
6	As (1), with delta-e^- removed	0.6	0.5
7	As (2), with delta-e^- removed	< 0.1	< 0.1
10	As (1), decay-e^+ scattered	< 0.1	< 0.1
11	As (2), decay-e^+ scattered	< 0.1	< 0.1
21	As (1), but beam-e^+ overlapping decay-e^+ in PCs[b]	< 0.1	0.2
22	As (2), but beam-e^+ overlapping decay-e^+ in PCs	< 0.1	< 0.1

[a] This is an identifier used internally by the collaboration. See the appendix of Ref. [85] for more detail.
[b] Events where the beam-e^+ overlap in the DCs are removed, since DC hits are used to reconstruct the decay-e^+.

3.3.3 Muon cuts

Last muon plane

The muon must have a hit in the PC immediately before the stopping target, but not in the PC following the target.

Muon radius at target

A radial cut was needed on the muon's final position to ensure that the decay occurred in the metal target rather than the surrounding foils (see Fig. 2.13). A cut was also needed to keep decay positrons with high transverse momentum inside the tracking region; otherwise they could strike external material (*e.g.* glass frames) and scatter. This cut demanded

$$r = \sqrt{u_{\mathrm{PC5}}^2 + v_{\mathrm{PC6}}^2} < 2.5\,\mathrm{cm}, \tag{3.4}$$

where $(u_{\mathrm{PC5}}, v_{\mathrm{PC6}})$ were the positions of the wire centres in the PCs immediately before the target.

Pulse width at target

This analysis used a new cut on the pulse widths in the PCs immediately before the target (PC5 and PC6). The voltages on these PCs were deliberately lowered while acquiring data, which increased the sensitivity to the muon energy deposited but lowered the positron efficiency. The pulse widths were used to reject muons stopping in the PC gas rather than in the metal target, since these muons will deposit more energy in both PC5 and PC6.

Figure 3.11 shows the two cut lines that selected "zone 1", which contained a sample of muons that primarily stopped in the metal target. There is a gas "band" on the right hand side that was almost completely removed by the cut-B line; its position in data was tuned so that no more than 0.5% of the gas distribution leaked across the cut line. The cut-A line removed muons that stopped in the PC6 wires. The sensitivity of the result to the cut position in data is described in Section 6.2.10.

The simulation's PC response was not tuned, resulting in the differences between Fig. 3.11(a) and 3.11(b). In order to avoid a systematic uncertainty from a mismatch in stopping distributions between data and simulation, the cut-B was tuned in simulation so that the fraction of muons surviving the cut was the same as in the data.

3.3.4 Decay particle trajectories

Events can have multiple decay positron tracks due to a hard scatter (*i.e.* a large angle multiple scatter that breaks a trajectory to the point where it is reconstructed as two tracks), delta-electrons, decay positrons that scatter off material and cause a second track, or a real beam positron that has survived the earlier selections. This section describes the selections and cuts that result in a single decay trajectory.

Number of tracks

There had to be at least one decay trajectory in the event.

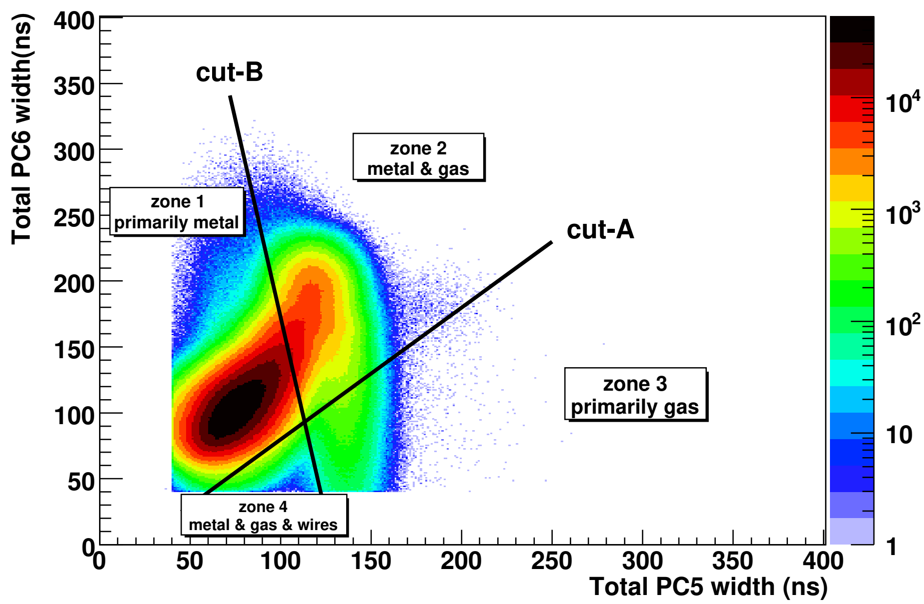

(a) PC5 and PC6 widths in **data**.

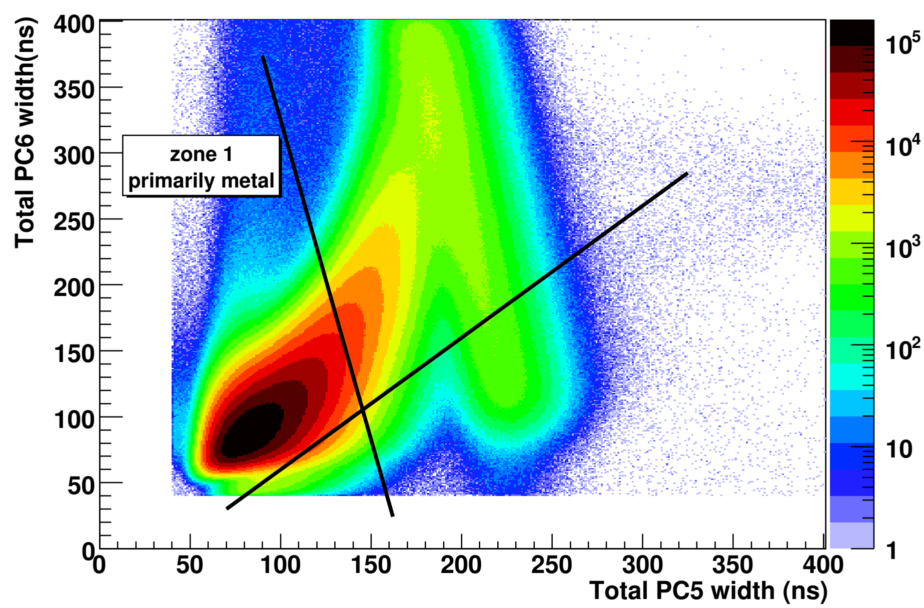

(b) PC5 and PC6 widths in **simulation**.

Figure 3.11: Pulse widths in the PCs immediately before the target. Muons that stop in the target will deposit less energy in PC5 and PC6, allowing them to be selected. The cut positions differ in data and simulation since the simulation's PC response has not been tuned.

Helix fit status

The helix fitting algorithm reported a code to describe how successful it was in fitting a track. Only completely successful fits were used.

First and last plane

A track had to be fully contained in either the upstream or downstream half of the detector. Tracks could not cross the target module. There was no cut on the number of planes crossed, or the first plane crossed.

Particle charge

Only particles that were reconstructed with positive charge were accepted (this was determined by their direction of winding). There were genuinely negative particles in the detector; principally these were delta-electrons, but there were also small contributions from bremsstrahlung (the γ could undergo pair production), and the decay process $\mu^+ \to e^-\bar{\nu}_e\nu_\mu e^+e^-$. The analysis could incorrectly identify particles as negative; for example, a beam positron that only appears in the upstream detector half due to a scatter in the stopping target could look like an upstream-going decay particle with negative charge.

Pair matching

Pairs of track candidates were extrapolated to find their closest distance-of-approach. If both tracks were on the same side of the target, and they appeared to overlap (came within 2 cm), the track furthest from the target was removed since it probably corresponded to a large angle scatter. If there was a track in each detector half, and they appeared to overlap at the target (came within 0.5 cm), then both tracks were removed since they were probably one beam positron traversing the entire detector.

Muon-track distance cut

Tracks were extrapolated to the z-position of the stopping target. If the distance between the extrapolated position and the muon's last position exceeded $(1 + 1/|\cos\theta|)$ cm, then the track was removed. The $\cos\theta$ dependence of this cut was a result of the reconstruction quality depending on angle.

Track-target distance

The track with a starting plane closest in z to the target was selected.

Muon-track distance selection

If there was still more than one decay positron candidate, the extrapolation of the track to the z-position of the target was used to select the track closest to the muon's last position.

Muon decay time

If a muon decayed in less than $1.05\,\mu$s, and the decay positron was upstream, then the drift chambers could still receive late ionisation from the muon, reducing the efficiency of upstream positron reconstruction. The event classification selection has already made an implicit time cut on the decay positron at $1.05\,\mu$s. However, this was determined from the PC times, which had a resolution of $\approx 20\,$ns. An additional cut was now made at $1.05\,\mu$s, using the time determined while fitting the helix, which has better resolution. There was also a cut made at $9.0\,\mu$s, since this allowed $1\,\mu$s of decay positron information to be analysed up to the data acquisition limit of $10\,\mu$s.

3.3.5 Kinematic fiducial region

The helix fitting routine determined the radius and wavelength of each track, from which the transverse momentum was approximately given by

$$p_t\,[\,\mathrm{MeV/c}] = 300 \times \mathrm{B}\,[\,\mathrm{T}] \times \mathrm{r}\,[\mathrm{m}], \tag{3.5}$$

and the longitudinal momentum was approximately given by

$$p_z\,[\,\mathrm{MeV/c}] = \frac{300}{2\pi} \times \mathrm{B}\,[\,\mathrm{T}] \times \lambda\,[\mathrm{m}]. \tag{3.6}$$

After all the selections and cuts from the previous section, the decay spectrum was reconstructed; this is shown alongside the theoretical spectrum in Fig. 3.12.

The region of $(p, \cos\theta)$ for spectrum fitting (the "fiducial") was then chosen to minimise bias and inefficiency, maximise resolution and sensitivity, and ensure that any deficiencies were well matched between data and simulation. This selection is shown in Fig. 3.12(c), where the region *within* the dashed white lines is accepted. The fiducial boundaries will now be described.

Total momentum, $p < 52.0\,\mathrm{MeV/c}$

An upper momentum cut was placed at $p = 52.0\,\mathrm{MeV/c}$ for several reasons. First, the spectrum was most affected by radiative corrections at higher momenta; the cut therefore reduced the systematic uncertainty due to radiative corrections. Second, the momentum calibration of the spectrum used the higher momentum region; as a result, the region was excluded from the decay parameter extraction to be conservative. Third, the momentum resolution was degraded at higher momenta.

Longitudinal momentum, $|p_z| > 14.0\,\mathrm{MeV/c}$

The cut was necessary to avoid a wavelength ambiguity (see Section 3.2.5). Although the chambers were specifically repositioned to reduce this ambiguity, it was still present and had to be avoided. The statistical gain from adjusting this cut was insignificant for $P_\mu^\pi\,\xi$.

Angle, $0.54 < |\cos\theta| < 0.96$

At small angles (large $|\cos\theta|$) the wavelength was poorly resolved, and the helix fitting algorithm began to fail. At large angles (small $|\cos\theta|$) the reconstruction became unreliable due to increased multiple scattering. The measurement of $P_\mu^\pi\,\xi$ benefits the most by including larger angles, but the improvement is still only marginal.

Transverse momentum, $(10.0 < p_t < 38.0)\,\mathrm{MeV/c}$

The maximum transverse momentum cut was coupled to the target radius cut of 2.5 cm; together these cuts kept the positron trajectory within the instrumented region of the chambers.

There was a minimum transverse momentum cut to avoid small angle difficulties where the track radius became comparable to the wire spacing.

3.4 Extraction of muon decay parameters

The spectrum from data was compared to a simulation that had been generated with hidden values of ρ, δ and $P_\mu^\pi\,\xi$; the hidden range was limited to $\pm1\%$ of the standard

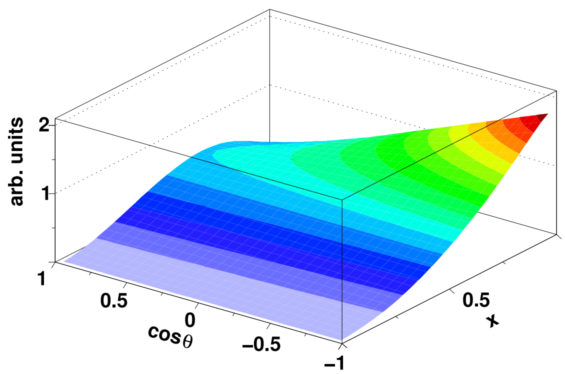

(a) Theoretical spectrum.

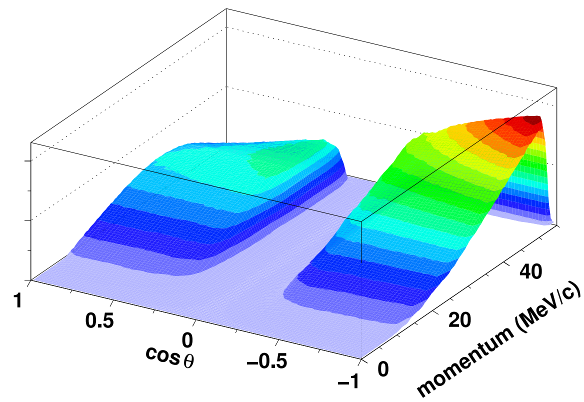

(b) Reconstructed spectrum.

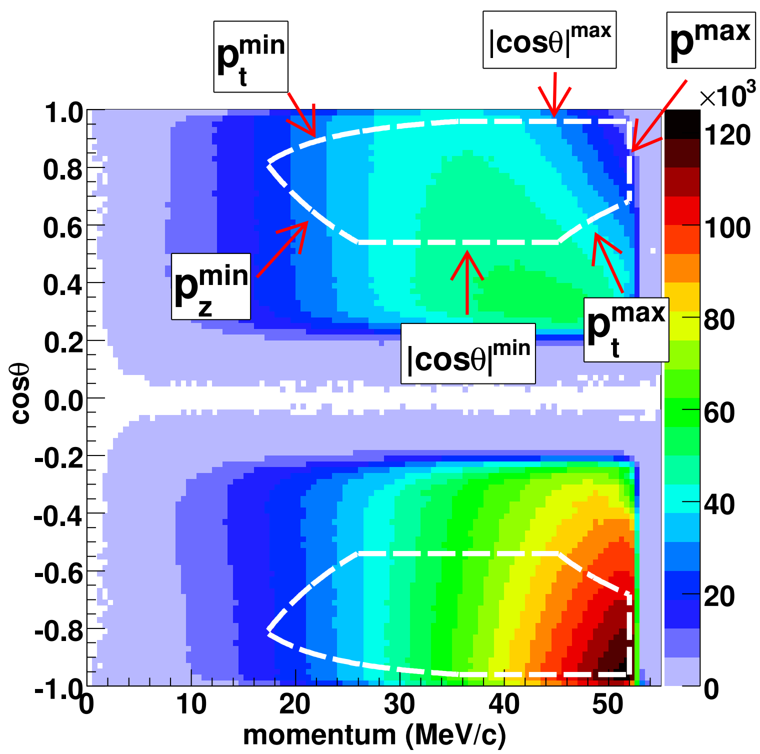

(c) Reconstructed spectrum with fiducial boundaries.

Figure 3.12: Comparison of theoretical and reconstructed muon decay spectra. The lower plot shows the fiducial region, which is *within* the dashed white lines.

model values. The parameter η was fixed to the world average, for the reasons already described in Section 1.4.5. The same analysis was applied to both data and simulation, so that biases and inefficiencies cancelled to first order.

The hidden values were only revealed after all the systematic uncertainties were evaluated on the *difference* between the data and simulation spectra. Since the result from the data was unknown until the end, this removed potential biases from tuning cuts, rejecting anomalous data, and evaluation of systematic uncertainties.

The method of comparing the data and simulation spectra will now be described. Recall the expression from Section 1.4.4 for the differential decay rate,

$$\frac{d^2\Gamma}{dx\,d\cos\theta} = k(x)\left\{F_{IS}(x) + P_\mu \cos\theta F_{AS}(x)\right\}, \tag{3.7}$$

where x is the reduced energy, and

$$k(x) = \frac{m_\mu}{4\pi^3} W_{e\mu}^4 G_F^2 \sqrt{x^2 - x_0^2} \tag{3.8}$$

$$F_{IS}(x) = x(1-x) + \frac{2}{9}\rho\left(4x^2 - 3x - x_0^2\right) + \eta x_0(1-x) + F_{IS}^{RC}(x), \tag{3.9}$$

$$\begin{aligned} F_{AS}(x) = {} & \frac{1}{3}\xi\sqrt{x^2 - x_0^2}\left[1 - x + \frac{2}{3}\delta\left(4x - 3 + \left(\sqrt{1 - x_0^2} - 1\right)\right)\right] \\ & + F_{AS}^{RC}(x), \end{aligned} \tag{3.10}$$

and the terms were defined in Section 1.4. The expression is linear in $P_\mu^\pi\,\xi$, $P_\mu^\pi\,\xi\delta$, ρ and η. Assuming standard model radiative corrections, this allows exact derivatives of the spectrum to be constructed,

$$S = \frac{d^2\Gamma}{dx\,d\cos\theta}, \tag{3.11}$$

$$\left.\frac{\partial S}{\partial \rho}\right|_{P_\mu^\pi\,\xi, P_\mu^\pi\,\xi\delta, \eta} = k(x)\cdot \tfrac{2}{9}(4x^2 - 3x - x_0^2), \tag{3.12}$$

$$\left.\frac{\partial S}{\partial \eta}\right|_{P_\mu^\pi\,\xi, P_\mu^\pi\,\xi\delta, \rho} = k(x)\cdot x_0(1-x), \tag{3.13}$$

$$\left.\frac{\partial S}{\partial P_\mu^\pi\,\xi}\right|_{P_\mu^\pi\,\xi\delta, \rho, \eta} = k(x)\cdot \tfrac{1}{3}\cos\theta\sqrt{x^2 - x_0^2}(1-x), \tag{3.14}$$

$$\left.\frac{\partial S}{\partial P_\mu^\pi\,\xi\delta}\right|_{P_\mu^\pi\,\xi, \rho, \eta} = k(x)\cdot \tfrac{2}{9}\cos\theta\sqrt{x^2 - x_0^2}(4x - 4 + \sqrt{1 - x_0^2}). \tag{3.15}$$

Each derivative is independent of the muon decay parameters. The shape of each derivative is shown in Fig. 3.13. The *difference* between the data and simulation spectrum can be used with the derivatives to determine how much of each muon decay parameter is needed to make the spectra agree. The final result ($P_\mu^{\pi,data}\xi^{data}$) is then given by

$$P_\mu^{\pi,data}\xi^{data} = P_\mu^{\pi,sim}\xi^{sim} + \Delta\left[P_\mu\xi\right], \tag{3.16}$$

where $\Delta\left[P_\mu\xi\right]$ is difference between the data and the simulation with a hidden ξ^{sim} value. The validation of this method is described elsewhere[88]. For $P_\mu^\pi\,\xi$, the method can introduce a bias at a level that is negligible for this measurement, but may be of interest to future experimenters; this is described in Appendix F.

The technique required five spectra to be reconstructed with the same analysis software and cuts: data, simulation and one spectrum for three of the four derivatives (Eqs. (3.12), (3.14), (3.15)). Since the derivative spectra had negative regions (they were, after all, just a spectrum shape), a positive value was used to generate the $(p, \cos\theta)$ for a decay positron, but the sign of the spectrum was passed to the analysis software, which applied it when reconstructing the derivative spectra.

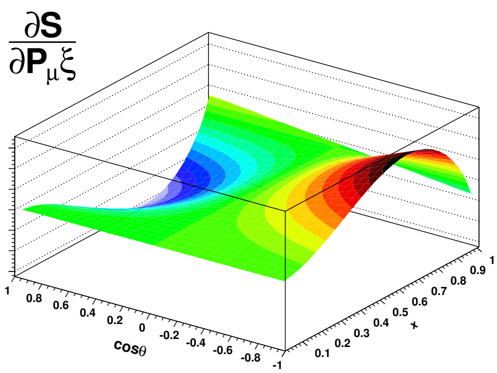

(a) Derivative with respect to $P_\mu^\pi \xi$.

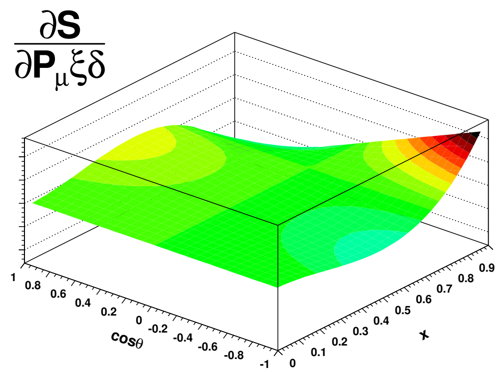

(b) Derivative with respect to $P_\mu^\pi \xi\delta$.

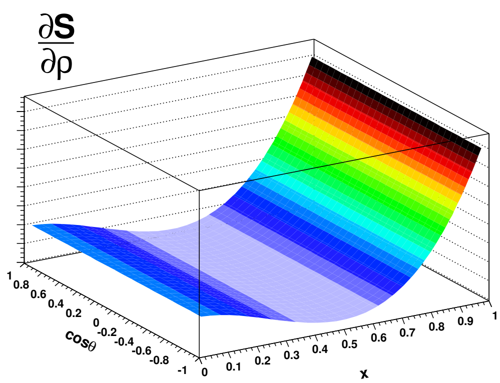

(c) Derivative with respect to ρ.

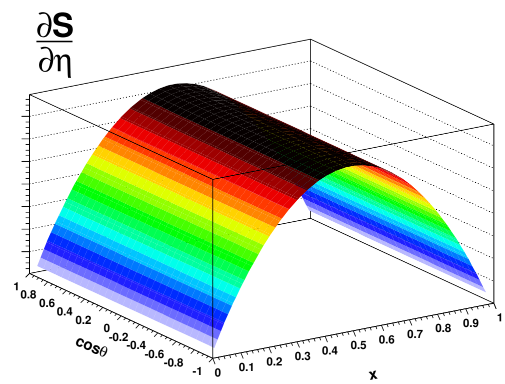

(d) Derivative with respect to η.

Figure 3.13: The shape of each muon decay parameter's contribution to the total spectrum. The vertical scales are inconsistent between the figures since only the *shape* is important here.

3.5 Energy calibration

The maximum positron energy from muon decay is approximately 52.83 MeV. This provided a spectrum feature that was used to calibrate the energy scale of the reconstruction. The region close to the kinematic endpoint is shown for a limited angular range in Fig. 3.14, where the sharp edge is smeared by the reconstruction resolution and radiative corrections. The figure includes a fit, where a step function has been convoluted with a Gaussian to model the resolution. The position of the step function had a linear dependence on $1/|\cos\theta|$ due to the planar construction of the detector, and this is shown for data and simulation in Fig. 3.15. They are both displaced from the kinematic endpoint due to positron energy loss in the metal target and reconstruction bias.

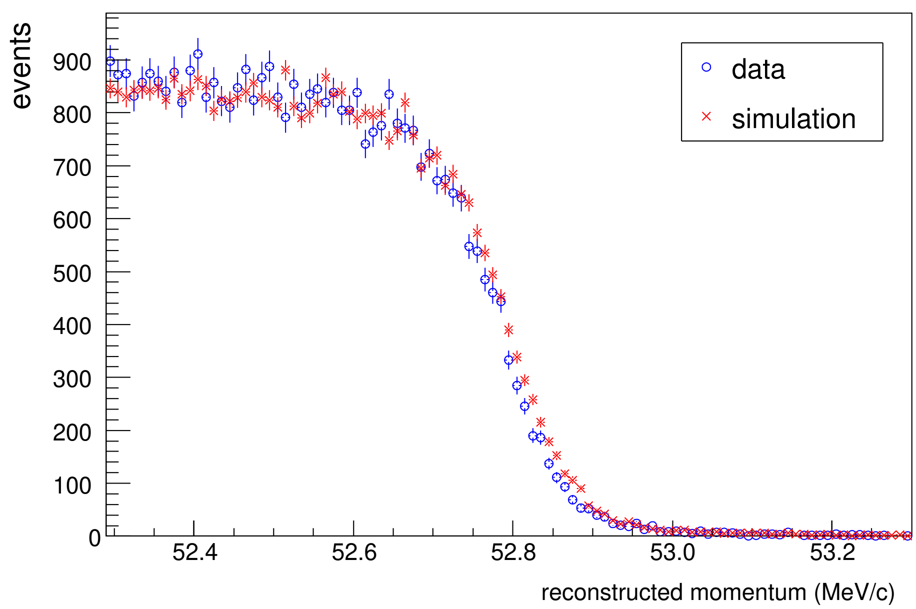

Figure 3.14: Example of the kinematic endpoint in data and simulation, over a limited angular range. The endpoint positions for all angles are shown in Fig. 3.15.

There is a discrepancy of about 10 keV/c between data and simulation, with a weak dependence on angle. The cause of this discrepancy was investigated, with the following explanations considered: a mismatch in the muon stopping distribution between data and simulation, an error in the stopping target thickness in the simulation, an error in the scale of the magnetic field map in the data, and an error in the positron energy loss physics in the simulation. These causes were ruled out[33], with the exception of an error in the

[33]Either they were too small, or implied the discrepancy between data and simulation should depend

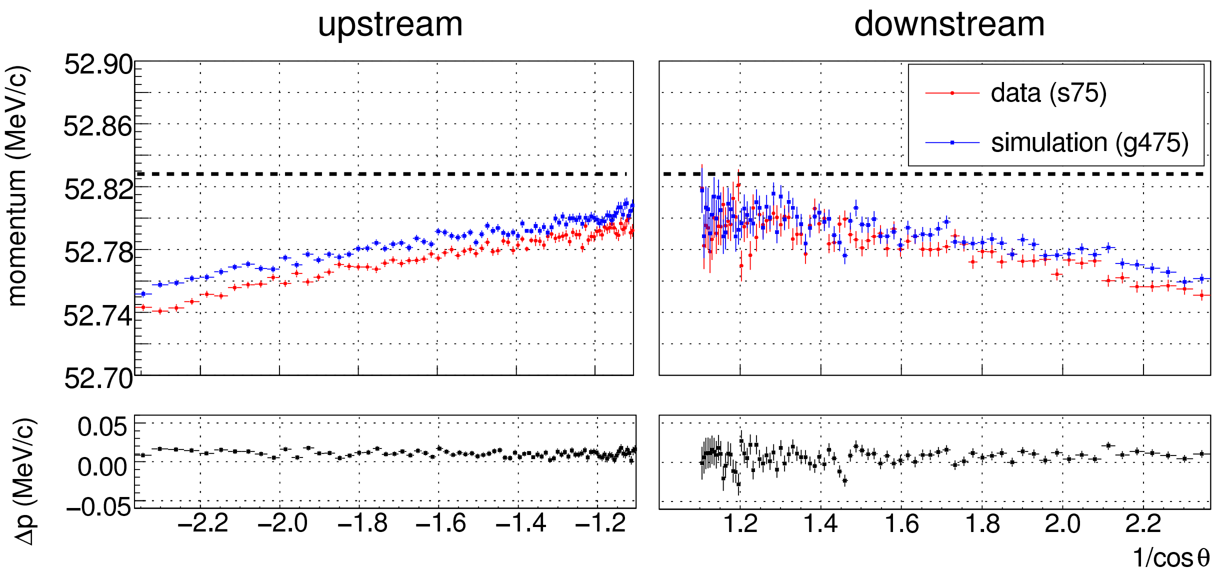

Figure 3.15: Dependence of endpoint position on $|1/\cos\theta|$. The black dashed line is the expected endpoint from kinematic considerations.

positron energy loss physics in the simulation, whose dependence on $\cos\theta$ is complicated by the use of refined STRs (see Section 3.2.7). Specifically, an energy loss error in the simulation should result in a data-simulation difference that is strongly angle-dependent. However, we apply an STR calibration procedure, and this appears to convert the data-simulation difference into an angle-independent offset of about 10 keV/c at the spectrum endpoint.

The analysis procedure corrected the data so that the endpoint overlapped the simulation. This change can be propagated to the rest of the spectrum assuming a constant shift, a constant scale, or a combination of shift and scale. The choice of propagation model will be later treated as a systematic uncertainty.

3.6 Time dependence of depolarisation

The previous $P_\mu^\pi \xi$ analysis[22] determined the time-dependent relaxation of the polarisation by fitting

$$G(t) = \frac{N_U(t) - N_D(t)}{N_U(t) + N_D(t)} = KP_\mu(t) \; , \tag{3.17}$$

on $\cos\theta$, which is not observed.

where N_U and N_D are the number of the upstream and downstream counts within the standard fiducial region, and K is a constant that depends on this region. Equation (3.17) can be explicitly written

$$G(t) = \frac{\int_p^q \int_a^b n(x, \cos\theta, t) dx\, d\cos\theta - \int_p^q \int_{-b}^{-a} n(x, \cos\theta, t) dx\, d\cos\theta}{\int_p^q \int_a^b n(x, \cos\theta, t) dx\, d\cos\theta + \int_p^q \int_a^b n(x, \cos\theta, t) dx\, d\cos\theta}. \tag{3.18}$$

If only the time dependence of G is sought, then a new quantity with the same time dependence can be defined as

$$G'(t) = \frac{\int_p^q \int_a^b w_1 n dx\, d\cos\theta - \int_p^q \int_{-b}^{-a} w_1 n dx\, d\cos\theta}{\int_p^q \int_a^b w_2 n dx\, d\cos\theta + \int_p^q \int_{-b}^{-a} w_2 n dx\, d\cos\theta}, \tag{3.19}$$

where $w_1(x, \cos\theta)$ and $w_2(x, \cos\theta)$ are weighting terms that can be freely chosen, and are time independent. Since $\cos\theta > 0$ corresponds to downstream decays, and $\cos\theta < 0$ to upstream decays, Eq. (3.19) can be simplified by making $w_1 \propto \cos\theta$ and $w_2 \propto |\cos\theta|$; the sum over the bins, both upstream (US) and downstream (DS), is then

$$G'(t) = \frac{\sum_{\mathrm{US,DS}} w_1 N}{\sum_{\mathrm{US,DS}} w_2 N}, \tag{3.20}$$

where N is the number of integrated counts in the bin.

For this measurement, w_1 and w_2 were related to the theoretical asymmetry itself[19]. The full expression for the differential decay rate was given in Eq. (3.7). Neglecting radiative corrections and the positron mass, and assuming standard model values for ρ and δ, the differential decay rate is

$$\frac{d^2\Gamma}{dx d\cos\theta} = x^2 \left[(3 - 2x) + P_\mu \xi \cos\theta (2x - 1)\right], \tag{3.21}$$

which produces a theoretical asymmetry of

$$A(x, \cos\theta) = P_\mu \xi \cos\theta \left(\frac{x - \frac{1}{2}}{\frac{3}{2} - x} \right). \tag{3.22}$$

After integration over $\cos\theta$, Eq. (3.22) is shown in Fig. 3.16. The greatest contribution to the integral asymmetry comes from the highest energy positrons; the positrons with $x \approx 0.5$ contribute little to the total asymmetry yet add statistical noise. Positrons with

$x < 0.5$ actually detract from the total asymmetry, and add more statistical noise. The optimum sensitivity is achieved by weighting each positron according to how much it contributes to the asymmetry, by relating w_1 and w_2 to Eq. (3.22).

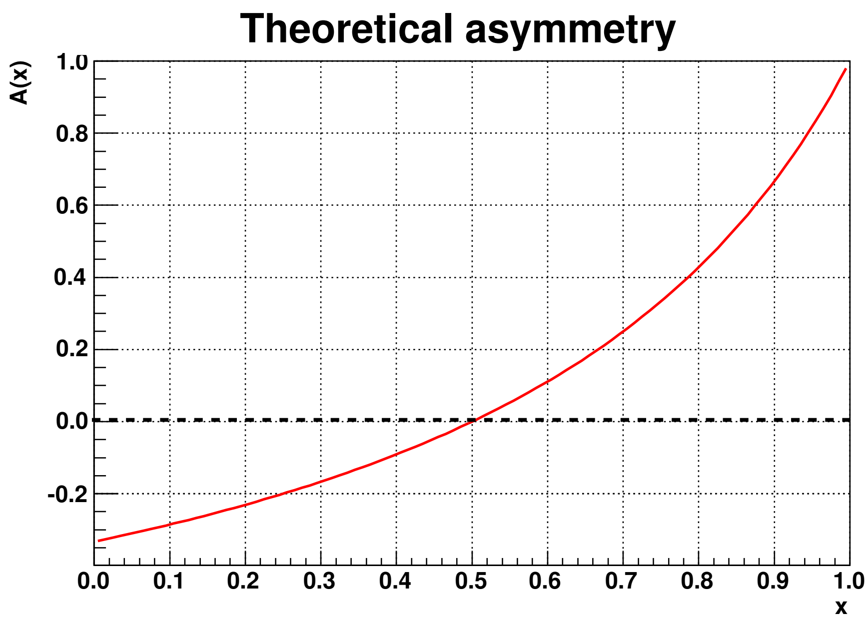

Figure 3.16: Theoretical prediction for the asymmetry, assuming standard model values for the muon decay parameters (Eq. (3.22)). Positrons with $x < 0.59$ are excluded from the analysis.

The weighting factors used in the current analysis were

$$w_1 = A(x, \cos\theta)\,|A(x, \cos\theta)|^n\,, \tag{3.23}$$

$$w_2 = |A(x, \cos\theta)|\,|A(x, \cos\theta)|^n\,, \tag{3.24}$$

$$n = 1. \tag{3.25}$$

Different powers of n were examined, but $n = 1$ combined with rejection of positrons with $x < 0.59$ was found to be close to optimum[34]. The combination of weighting the counts, and removing low energy positrons significantly reduced the uncertainty on the depolarisation rate. An example of a fit to a single set is shown in Fig. 3.17, where $P_\mu(t) = P_\mu(0)\exp(-\lambda t)$ has been used.

[34]The lower momentum positrons could be included, and their sign could be reversed for $x < 0.5$, but this doesn't improve the statistical precision by a significant amount.

Note that the weighting technique is only applied to the asymmetry analysis that determines the time dependence of P_μ. It is *not* applied to the normal spectrum analysis.

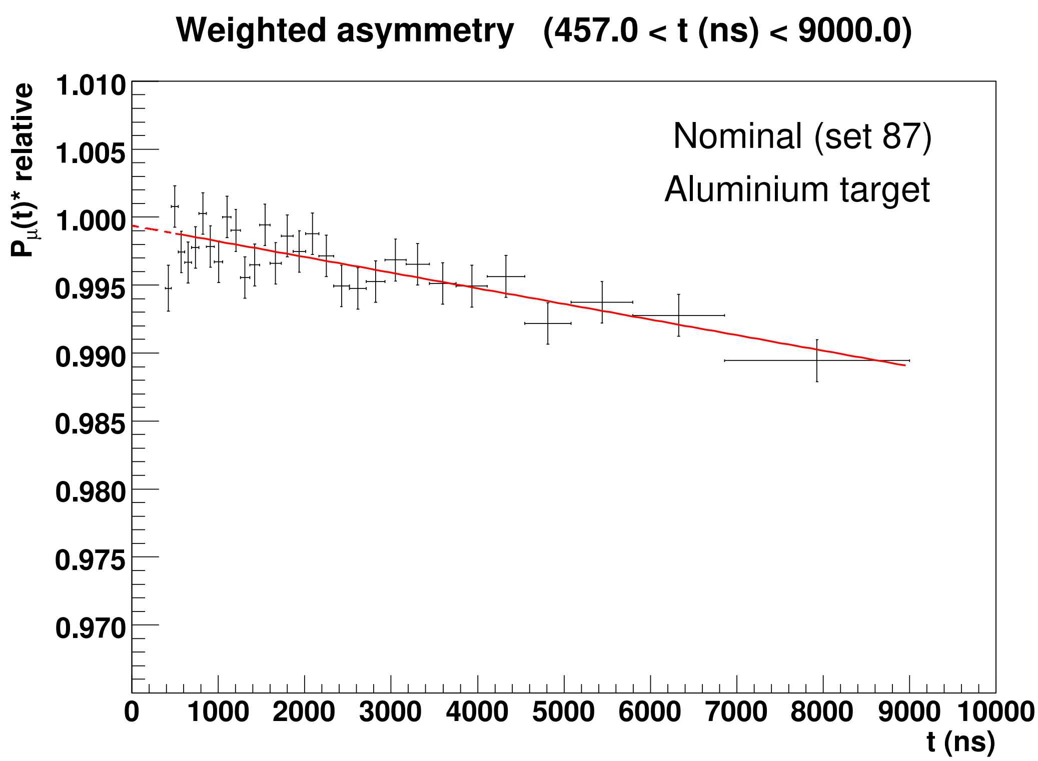

Figure 3.17: Example of an exponential fit to the asymmetry. The vertical scaling is arbitrary.

3.7 Muon beam spots

Individual muon trajectories could not be reconstructed due to their low transverse momentum, significant energy loss, and the muons suffering from more electronic crosstalk than the positrons. However, the average position of the muon beam within the chamber region was found to be well described by a helix with decaying amplitude and wavelength; this is demonstrated in Fig. 3.18, which shows the trajectory of a simulated beam with large average transverse momentum. The solid markers indicate a pair of orthogonal detector planes, and the positions of these were fit with

$$\begin{pmatrix} \langle x \rangle \\ \langle y \rangle \end{pmatrix} = (A - A_d \cdot z') \begin{pmatrix} \cos\left(f\left(z'\right)\right) \\ \sin\left(f\left(z'\right)\right) \end{pmatrix} + \begin{pmatrix} \Delta_x \\ \Delta_y \end{pmatrix}, \tag{3.26}$$

where $z' = z + 59.59\,\mathrm{cm}$, $\langle x \rangle$ and $\langle y \rangle$ are the mean muon positions, Δ_x and Δ_y are the positions of the helix centre, A is the amplitude of the helix at $z' = 0$, A_d is a linear decay parameter, and

$$f(z') = 2\pi \frac{z'}{\lambda - \lambda_d z'} + \phi, \tag{3.27}$$

where λ is the helix wavelength at $z' = 0$, ϕ is the phase at $z' = 0$ and λ_d is a linear decay parameter. The parameters A and λ approximately corresponded to average transverse and longitudinal momentum.

The fit used the mean positions at all chambers except the final proportional chambers and the last two pairs of drift chambers. This gave a total of 11 points, from which 7 fit parameters were determined. The parameters of the average muon beam were used to tune the M13 beam line, and monitor the muon beam's stability (see Sections 5.2 and 5.4.7 respectively).

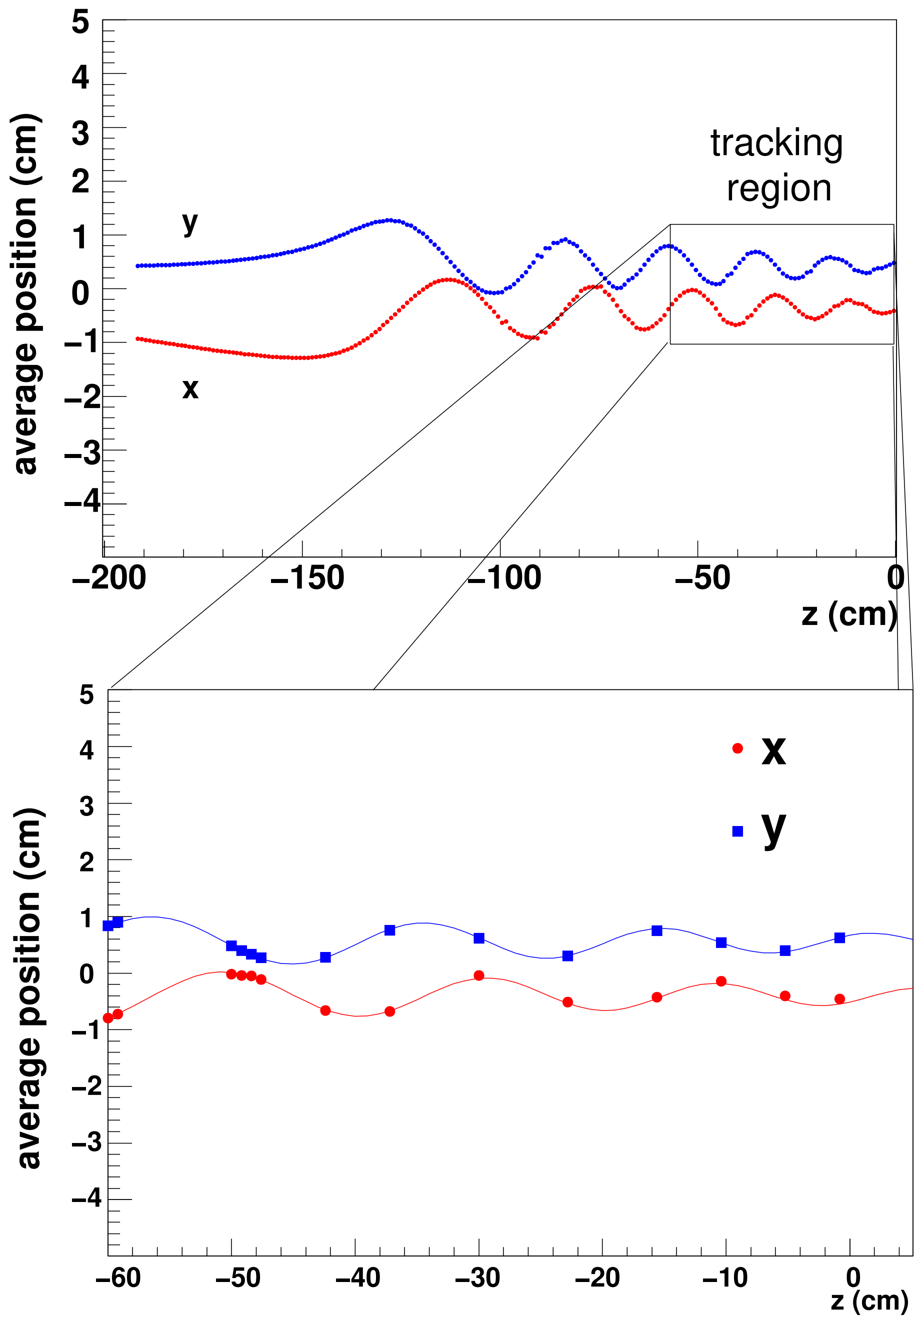

Figure 3.18: A simulated beam with large average transverse momentum. The upper figure shows the mean trajectory from the end of the M13 beam line up to the stopping target. The lower figure shows the beam inside the chambers, where the solid markers indicate a pair of planes.

Chapter 4

Simulation

4.1 Introduction

A detailed simulation of the detector was implemented in `GEANT3.21`[94] for several purposes. First, the performance of the reconstruction was studied and optimised. Second, the muon decay parameters could be measured while remaining blind to their data values (see Section 3.4). Third, the final polarisation of the muons was determined by using the simulation to transport their momentum and spin to the stopping location. An outline will be given of the `GEANT3` features that were important to the TWIST experiment, and the custom code modules most relevant to $P_\mu^\pi \xi$ will be described in more detail.

4.2 Overview of the simulation

Simulation events started with the generation of muons and beam line positrons at the end of the M13 channel, with rates that matched the data[35]. Pions, protons, cloud muons and radioactive beam line particles were not simulated, since they were removed from the real data by physical windows and a time-of-flight cut on the trigger particle (see Section 3.3.1). The simulated particles and their secondaries, including photons, e^+ and e^-, were transported through the detector materials. The response of the chambers to the subsequent ionisation was tuned to match the data.

When a muon decayed, the simulation produced a positron with energy and angle following a pre-defined spectrum that used a hidden value of $P_\mu^\pi \xi$. Transportation of the particles through the magnetic field used a fourth order Runge-Kutta numerical method to propagate the momentum. The muon's spin was transported with equivalent precision. The simulation demanded that at least one transportation step was taken in every material; this was important for the energy loss in thin materials such as the cathode foils, which each had a thickness of just $0.9\,\mathrm{mg/cm^2}$.

[35]The muon rate matched the average trigger rate from the DAQ. The beam positron rate was tuned so that the probability of observing a positron within an event was matched in data and simulation.

The physics processes in the simulation included ionisation energy loss, δ-electron production, Bremsstrahlung, multiple Coulomb scattering, positron annihilation, Compton scattering, pair production, the photo-electric effect and Bhabha scattering; these processes were *not* tuned from their `GEANT` defaults. Materials external to the tracking volume were included, such as the beam pipe and the upstream beam package (see Section 2.9); the positrons could backscatter from these materials and re-enter the tracking volume, introducing extra hits and a resulting inefficiency that the simulation had to reproduce.

The output of the simulated wire chambers and scintillators were written to disk through a digitisation routine that mimicked the real DAQ. This allowed the data and simulation to undergo an almost identical analysis, with the exception of the crosstalk removal routine, since the simulation did not contain this level of electronics detail.

For each data set accumulated, a simulation with at least twice the statistics was generated; this was adequate since the $P_\mu^\pi \xi$ precision is limited by systematic uncertainties, not statistical uncertainties.

The customised code modules that are important for $P_\mu^\pi \xi$ will now be described. Specifically these are the routines that generated the initial position/angle/momentum of the beam line particles, the muon depolarisation mechanisms, the generator of decay positrons, and the simulation of the wire chamber response. The performance of the simulation will be compared to the data when the systematic uncertainties are evaluated in Chapter 6.

4.3 Primary particle generation

The simulated beam (muons and positrons) is generated based on measurements of the real beam. For the muons this was achieved with a pair of time expansion chambers (TECs) that were inserted at the beginning and end of each data set, at a z-location upstream of the solenoidal fringe field. The TECs had low positron efficiency, so that the beam line positrons were instead measured using the upstream drift chambers, while the magnetic field was off. The precision was inferior to the muon beam measurement, but it was adequate since $P_\mu^\pi \xi$ is relatively insensitive to the quality of the simulated positron beam.

The initial position and angle of the muons and beam line positrons in the simulation were generated with the same algorithm: the position was selected by dividing up the $x - y$ distribution of the number of muons into $0.1\,\mathrm{cm} \times 0.1\,\mathrm{cm}$ bins, and using this as a probability distribution. The θ_x and θ_y angles were drawn from independent Gaussian distributions with means and widths that matched the data measurement for the particular (x, y) bin; the widths of the Gaussian angle distributions were multiplied by $c_x = 0.64$ in the $x-$module and $c_y = 0.48$ in the $y-$module in order to correct for multiple scattering (the muons were multiple scattered by the entrance window and the gas within the TECs while they were being measured, which increased the apparent emittance of the beam).

The initial z-position of the muons was always at the centre of the TEC modules. The initial momentum (p) depended linearly on the $x-$position at the TECs according to

$$p(x) = p_0 - p_1 x. \tag{4.1}$$

The p_1 parameter was tuned using comparisons such as Fig. 4.1, and $p_1 = 0.17\,\mathrm{MeV/c}$ per cm was found to be optimal. The p_0 parameter was then determined by setting the average momentum to 29.6 MeV/c so that

$$\begin{aligned} p_0 &= \langle p(x) \rangle + p_1 \cdot \langle x \rangle \\ &= 29.6\,\mathrm{MeV/c} + (0.17\,\mathrm{MeV/c}) \cdot \langle x \rangle\,, \end{aligned} \tag{4.2}$$

and $\langle x \rangle$ was different for each data set. Muons were not generated above the kinematic maximum momentum of 29.79 MeV/c.

The materials traversed by the muons were described earlier (Section 2.11), where it was noted that the simulation required an extra 1.9 mg/cm^2 of material to match the data distribution of where the muon was last seen. However, the small amount of extra material was only included in half of the simulations; see Section 6.2.12 for further detail.

4.4 Depolarisation

The simulation must transport the muon's spin through the solenoidal magnetic field, and reproduce the time dependence of the spin once inside the stopping material. These routines did not exist in `GEANT3`, and were specially written for the TWIST experiment.

4.4.1 Initial polarisation

The simulation started the muons with anti-parallel spin and momentum vectors. This is an approximation with two limitations; first the muons were selected from the graphite

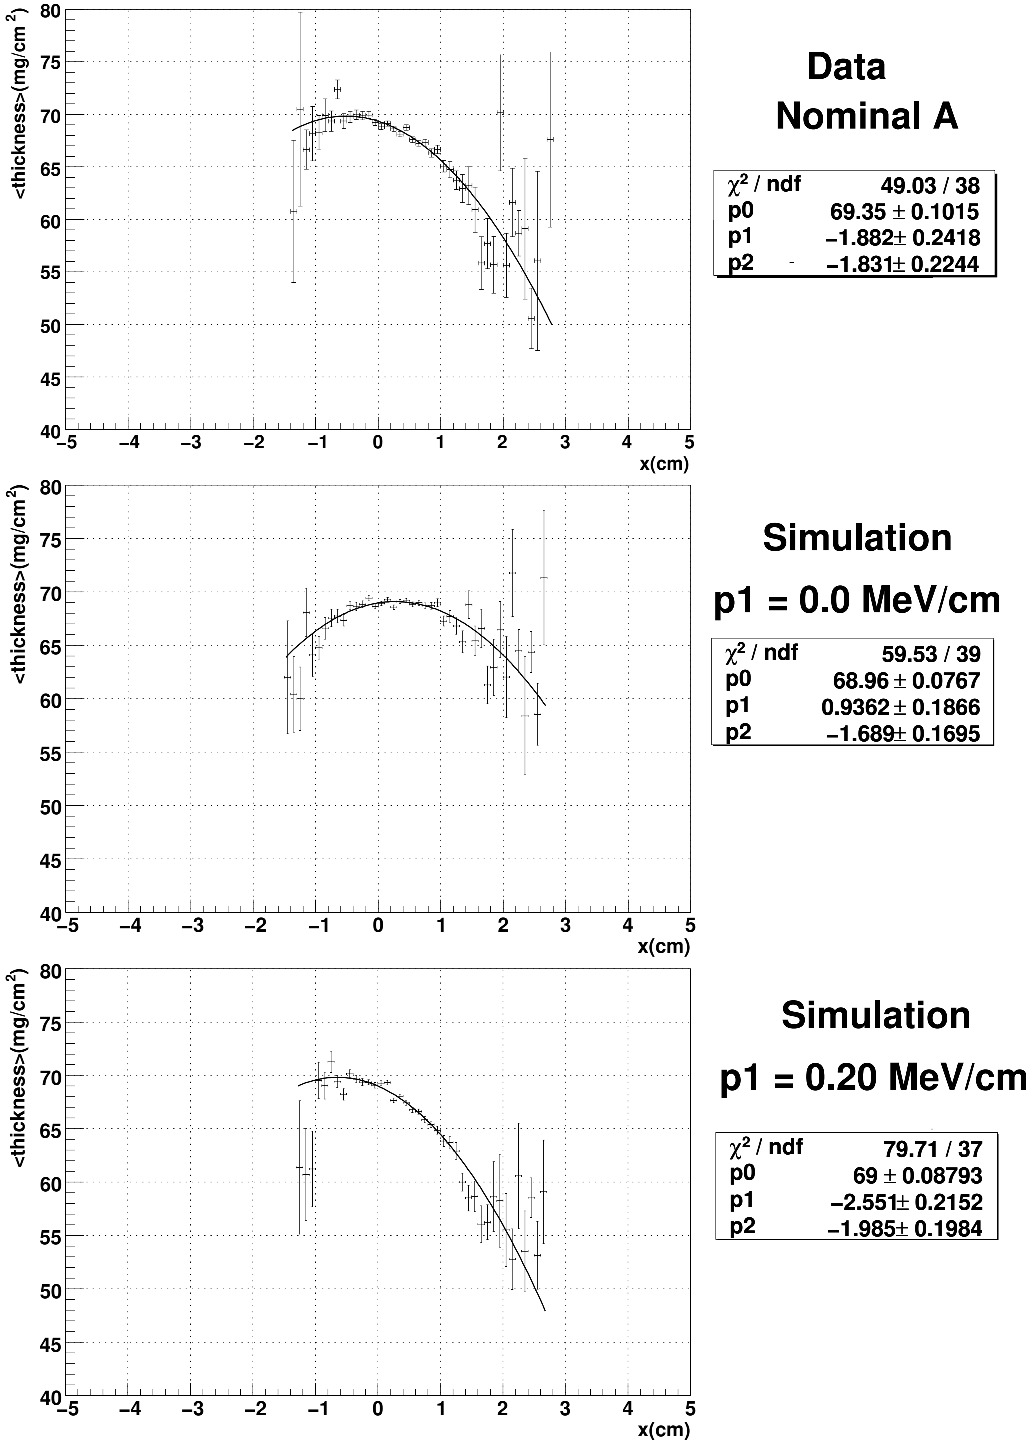

Figure 4.1: Average muon range (mg/cm^2) depends on the $x-$position of the muon at the TECs. The relationship is well approximated by a quadratic function of the form range $= p_0 + p_1 x + p_2^2 x^2$. The simulation required $p_1 = \frac{dp}{dx} = 0.17\,\mathrm{MeV/cm}$ in order to match the data.

production target at an average depth of 16 μm. While exiting the target, multiple scattering changed the momentum vector but not the spin vector, causing them to no longer be exactly anti-parallel. This effect is later treated as a correction. Second, the muons have already passed through the magnetic fields in the M13 beam line, and the difference in precession frequencies between the momentum and spin vectors (Eq. (1.38)) can change the angle between them. This second effect is negligible, changing P_μ at the level $\sim 10^{-8}$; see Appendix I.

4.4.2 Electromagnetic field

Recall the Thomas equation from Section 1.6.1, which described the propagation of the spin vector ($\vec{s}$) in an electromagnetic field,

$$\frac{d\vec{s}}{dt} = \frac{e}{mc}\vec{s}\times\left[\left(\frac{g}{2} - 1 + \frac{1}{\gamma}\right)\vec{B} - \left(\frac{g}{2} - 1\right)\frac{\gamma}{\gamma+1}(\vec{\beta}\cdot\vec{B})\vec{\beta} - \left(\frac{g}{2} - \frac{\gamma}{\gamma+1}\right)\vec{\beta}\times\vec{E}\right], \quad (1.40)'$$

where t is time, (e/m) is the charge to mass ratio, c is the speed of light, g is the Landé g factor, γ is the Lorentz factor, B is the magnetic field, $\vec{\beta}$ is the velocity divided by c, and E is the electric field. The E field term was not included in the simulation; the drift and proportional chambers have significant electric fields, but a field of $(+E)$ between the entrance cathode foil and wire was followed by a symmetric $(-E)$ field between the wire and exit cathode foil, which effectively cancels the change in spin from this term. For programming convenience, the Thomas equation was manipulated into this form:

$$\frac{d\vec{s}}{dx} = \frac{d\vec{s}}{dt}\frac{1}{|\beta|} = \frac{e}{mc}\left[U_1(\vec{s}\times\vec{B}) - U_2(\vec{s}\times\hat{v})\right], \quad (4.3)$$

where

$$U_1 = \left[\frac{1}{|\beta|}\left(a + \frac{1}{\gamma}\right)\right], U_2 = \left[a\frac{\gamma}{\gamma+1}|\beta|\left(\hat{v}\cdot\vec{B}\right)\right],$$

and $a = (g-2)/2$. The spin was then propagated using a Taylor expansion,

$$\vec{s}\,' = \vec{s} + \frac{d\vec{s}}{dx}\Delta x, \quad (4.4)$$

where $\vec{s}\,'$ is the new spin vector and $\vec{s}$ is the spin at the previous step. In order to correct for numerical precision errors, the components of the spin vector were renormalised at each step. For example,

$$s'_z = \frac{s_z}{\sqrt{s_x^2 + s_y^2 + s_z^2}}, \tag{4.5}$$

and similarly for s_x and s_y. For each GEANT particle transportation step, smaller sub-steps were taken that depended on the magnitude of the magnetic field and muon velocity.

The Thomas equation can be solved exactly in a uniform and stationary magnetic field[97], and this solution has been used to validate the numerical method[59]. However, a precision problem was discovered and fixed for this analysis: not enough sub-steps were being taken at low velocities, resulting in a P_μ error of 4×10^{-4} for a nominal beam. This was fixed by increasing the number of sub-steps by a factor of ten. A new validation used the simulation with energy loss and multiple scattering disabled, to show that momentum and spin remained anti-parallel to $< 2 \times 10^{-5}$.

4.4.3 Time-dependent depolarisation

The muon's spin is precessed in the time interval between thermalisation and decay. The spin is also randomly flipped during this period ($\vec{s} \to -\vec{s}$), with a rate that matches the observed relaxation in the real data. The simulation does not implement the random walk of the muon in the metal stopping target, since all the necessary physics is included in the random spin flips.

4.4.4 Other depolarisation

The simulation did not include depolarisation due to the field between the nucleus and atomic electrons, muonium formation at low velocities, and muon-electron scattering, since they were all negligible (see Section 1.6.2). Also the spin of the positron was not transported since the experiment could not detect this quantity.

4.5 Muon decay

The default GEANT software does not even simulate muon decays assuming the standard model ($V - A$) interaction; it simply simulates a three body decay phase space into a positron and two massless particles[98]. Therefore a special program was used that allowed for non-standard model values of the decay parameters, and included radiative corrections with the required accuracy[88]. The software ran on a dedicated server at TRIUMF, and each time a simulation began, the list of decay positron $(p, \cos\theta)$ pairs was retrieved from the server. The custom program also produced the derivative spectra that were described in Section 3.4. In this mode of operation, a $(p, \cos\theta)$ pair *and* the relevant sign ($\pm$) were supplied to the simulation. The simulation passed the sign to the analysis, allowing a derivative spectrum with positive and negative regions to be produced.

4.6 Wire chamber response

The response of the drift and proportional chambers were simulated in detail, with the aim of accurately reproducing inefficiencies, bias and resolution. As a particle passed through the chamber gas, ion clusters were randomly produced along the trajectory, with a mean cluster spacing that was tuned to match the data. The drift distance to the wire was converted into a time using space-time-relationships from the `GARFIELD` software[96] (see Section 3.2.7). The time was then smeared to include electronics effects and diffusion. Signals from wires that were dead in the data (typically one or two out of more than 3500) were deactivated in the simulation.

The muons produced more ionisation than the positrons, which deadened the wire within 0.06 cm, with a mean recovery time of 3.0 μs (both of these parameters were measured using data). This effect is included in the simulation since it caused a small upstream inefficiency.

Chapter 5

Data

5.1 Overview

The commissioning run of the TWIST experiment was in 2001. The first sets for muon decay parameter extraction were acquired in 2002, and results for ρ and δ were published in 2005[86, 88]. An engineering run took place in 2003, followed by the second set of physics data in 2004, which have been analysed twice: the first publication in 2006 was for $P_\mu^\pi \xi$[59], and the second publication in 2008 was for ρ and δ[85].

The author commenced his studies in September 2005, and the experiment soon undertook an engineering run from October 2005 to December 2005. The data analysed for this measurement were acquired in two periods: muons were stopped in a silver target from October 2006 to December 2006, and in an aluminium target from May 2007 to August 2007. From June 2006 to September 2006 data were accumulated on the same aluminium target, but were not analysed[36].

5.2 Muon beam tuning

The muon beam was tuned using the TECs, which measured individual muon trajectories before the solenoidal field, and the internal beam that was described in Section 3.7.

A muon beam directed along the magnetic field axis with minimised transverse momentum will have the highest polarisation. This beam is claimed to have the smallest polarisation uncertainty since it experiences minimal transverse magnetic field components, and is more robust to angle and position misalignments of the original beam. At the time of tuning the TECs were not calibrated or aligned, allowing only an approximate measure of the beam's position and angle. The muon beam was steered to be roughly on-axis ($x \approx y \approx 0$), with small average angle, using the quadrupole steering that was

[36] Two metal spanners were accidentally left inside the detector over this period, affecting the tracking field in a way that could not be easily corrected. In other words, it was less effort to re-take the data in 2007.

described in Section 2.4. The same quadrupole steering was then used to minimise the transverse momentum of the beam, as measured by the chambers. (Even though the drift and proportional chambers were not aligned at this stage, their relative placement was known with higher precision than the removable TECs.) An example of the beam spots inside the detector after tuning is shown in Fig. 5.1.

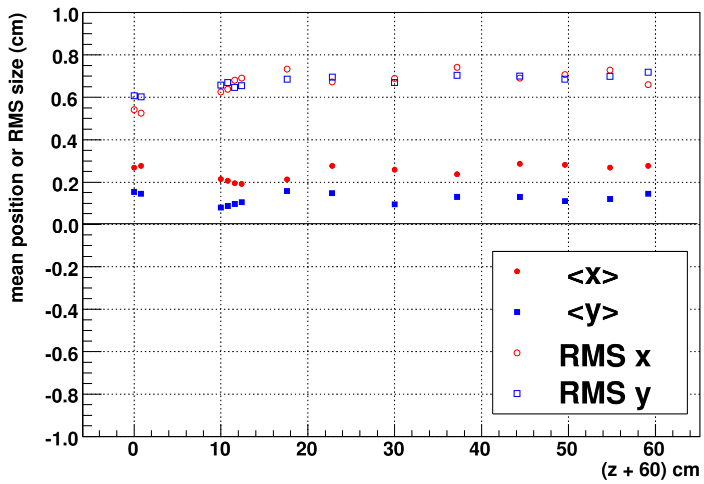

Figure 5.1: Muon beam inside the detector, after tuning.

5.3 Data sets

Data were accumulated in "runs" with a nominal file size of 2 GB, which typically took between 5 and 10 minutes. The runs were grouped into "sets", which are listed in chronological order in Table 5.1. A set contained about 900 runs, which required nearly six days of continuous acquisition; this amounted to about 0.9×10^9 triggers, which is a factor of three larger than the previous TWIST $P_\mu^\pi \xi$ measurement[22].

Table 5.1 shows there were four nominal sets (74, 75, 84, 87), two for each stopping target. Nominal was defined as an initial muon momentum of 29.6 MeV/c, a stopping distribution peaked at the centre of the target, a uniform magnetic field with a central value of 2.0 T, and no downstream beam package in place (see Section 2.9).

Set 68 had an incorrect stopping distribution, and was therefore used as a consistency check of the energy scale. The B= 1.96 T and B=2.04 T sets were tests of the magnetic field's effect on reconstruction. The set with the downstream beam package was used to validate the simulation of back-scattered positrons.

The sets with the steered beam (76 and 86) will be used later to evaluate the largest $P_\mu^\pi \xi$ systematic uncertainty. In these sets, the transverse momentum of the beam was maximised to deliberately lower the final polarisation of the muons. For the silver target (set 76), an average angle of $\theta_y \approx 30\,\mathrm{mrad}$ was introduced, and for the aluminium target (set 86), the beam was steered so that $\langle x \rangle \approx -1\,\mathrm{cm}$, and $\langle \theta_x \rangle \approx -10\,\mathrm{mrad}$.

In normal operation, the TECs were inserted at the start and end of each data set, but they were removed during the data set since they introduced additional depolarisation. Two special sets (72 and 82) were acquired with the TECs in place, to measure the long term stability of the muon beam, and the aging of the TEC sense planes.

Three smaller sets were taken at lower momenta: one at 28.75 MeV/c, and two at 28.85 MeV/c. These will later be used to validate the correction due to multiple scattering within the graphite production target.

Sets 73, 80 and 89 had the muons stopped immediately after the trigger scintillator, so that each decay positron was reconstructed independently in each half of the detector. These data were taken using the standard silver and aluminium targets, and a special large radius aluminium target that allowed the relaxation of certain analysis cuts. These sets will be used to measure the systematic uncertainties due to upstream-downstream inefficiencies, and from a resolution difference between data and simulation.

Table 5.1: Data sets accumulated in 2006/2007, in chronological order. The sets numbered 68-76 (80-93) used a silver (aluminium) stopping target of purity greater than 99.999%. Set 89 used a special larger radius aluminium target that had lower purity.

Set num.	Description	Number of events ($\times 10^6$) Before quality checks	Before cuts/ selections	Final spectrum
68	Stopping distrib. peaked $\frac{1}{3}$ into target	741	560	32
70	B = 1.96 T	952	780	50
71	B = 2.04 T	879	708	45
72	TECs-in, nominal beam	926	771	49
73	Muons stopped far upstream	1113	451	-
74	Nominal A	580	497	32
75	Nominal B	834	759	49
76	Steered beam A	685	610	39
80	Muons stopped far upstream	363	267	-
82	TECs-in, spread beam	861	612	-
83	Downstream beam package in place	943	830	49
84	Nominal C	1029	727	43
86	Steered beam B	1099	1015	58
87	Nominal D	854	776	45
91	Lower momentum I ($p = 28.75\,\mathrm{MeV/c}$)	225	206	11
92	Lower momentum II ($p = 28.85\,\mathrm{MeV/c}$)	322	272	15
93	Lower momentum III ($p = 28.85\,\mathrm{MeV/c}$)	503	463	26
89	Muons stopped far upstream	708	495	-

5.4 Data quality checks

The previous $P_\mu^\pi \xi$ measurement[22] could not rule out muon beam instabilities as the cause of the dominant systematic uncertainty. The abundant statistics for the current round of measurements allowed extremely conservative rejection of suspicious data, which will now be described.

5.4.1 Data acquisition system

The DAQ (data acquisition system) included the TDCs described in Section 2.12, and the computer to which they were interfaced. This computer used the software `MIDAS`. Runs were excluded due to the following DAQ problems:

- A TDC would sometimes receive too much information due to an electronics problem, and become out-of-synchronisation with the others. This was more common while the TECs were in place.
- A TDC would sometimes stop responding.
- A TDC channel could temporarily become corrupted.
- The computer's event accumulating software would sometimes crash.
- The computer could fail to make a database entry for a run, or write a run with zero events.
- While data were accumulating, a periodic pulser signal was added that could be subsequently analysed to make sure all the TDCs were operating correctly. If a run contained more than three pulser problems it was eliminated.

Any run that completed with too few events was conservatively eliminated, since it was indicative of a DAQ problem.

5.4.2 Chamber signals

During maintenance periods when the proton beam was off for several weeks, the detector was "opened up" to carry out maintenance work such as broken wire repairs, or changing the target foil. This took place in a special clean room, but some contamination due to dust was inevitable. In the days following such a maintenance period the current or voltage in a chamber occasionally exceeded a threshold and the run was stopped; the cause was believed to be residual dust, and these runs were eliminated. More seriously, after maintenance the pre-amplifiers sometimes suffered “electrical oscillations”, and produced signals that overloaded the DAQ. This could make the detector unusable for hours while post amplifier thresholds were adjusted. The periods of time during which this occurred were eliminated.

5.4.3 Rates

The DAQ computer recorded several accumulated counts:

- The number of upstream counts from the muon-counter.
- The number of upstream counts in the annular counter that surrounded the muon-counter (see Section 2.9).
- The number of downstream counts from the new downstream scintillator, which included muons, beam positrons and decay positrons.
- The current from an ion chamber close to the production target, which was proportional to the proton current.

The counts were converted into normalised rates using the ion chamber current, and then used to eliminate periods of beam instability.

5.4.4 Beam line stability

The DAQ recorded the currents and voltages across the M13 elements (dipoles, quadrupoles), and the positions of the slits and jaws. NMRs on the dipoles allowed the field to be automatically regulated to $< 0.01\,\mathrm{mT}$. Occasionally a regulator was not set properly by the operator, or the beam line settings were accidentally changed during a set, or the NMR signal became weak. Such occurrences were easily detected, and the runs eliminated.

There were periods where the operator observed instabilities in several of the beam line quadrupoles (Q4, Q6, Q7). Some of these instabilities were read-back errors, and therefore had no effect on the muon beam. The known genuine instabilities were eliminated, but the questionable runs were kept at this stage, and investigated more carefully using the muon beam itself.

5.4.5 Chamber foil bulging

A control system maintained the differential pressure between the drift chamber cathode foils and the He/N_2 mixture that surrounded the chambers (see Section 2.8). The density of the gas in the pressure sensing tubes depended on the temperature of the experimental hall. If the temperature of the chambers and the experimental hall were too far apart, then the measured differential pressure could become significantly different from the true differential pressure at the centre of the chambers. As a result, the control system set the differential pressure incorrectly, and the foils "bulged" away from their nominal positions.

A temperature difference of ± 3°C between the experimental hall and the chambers was tolerated, which corresponded to a bulge of $\pm 35\,\mu$m[90]. The operators maintained this difference by adjusting the air flow to the experimental hall, and were successful since only 10% of a single set had to be excluded.

5.4.6 Muon stopping distribution

Section 2.9 described a chamber in which the CO_2 and He concentration could be adjusted to control the muon range. This adjustment was automated by analysing a sample of the muons while accumulating data, and maintaining the average last-z position at which the muon registered. Runs were rejected if this average was outside of statistical fluctuations.

5.4.7 Muon beam stability

The set with the TECs in place allowed a high precision measurement of the muon beam stability (position, angle, size, correlation between angle and position *etc.*) over the period of a week. The average positions and angles from the TECs are shown in Fig. 5.2. The beam position was stable to < 0.1 cm, and the angles to < 1 mrad, which are negligible variations.

For the sets without the TECs in place, the internal muon beam analysis was available (see Section 3.7). Figure 5.3 shows the average position from this analysis, which is notably stable to < 0.02 cm, but subject to small systematic steps. The TEC characterisations from the start and end of each set allowed an upper limit on the change in polarisation due to these steps. In practice all the internal muon beam parameters were examined for all sets, and any systematic changes in their values were investigated.

The examination of the internal beam led to several important conclusions. First, the set with the most stable internal beam had the largest change in start/end-of-set TEC characterisations. This strongly suggested that TEC non-reproducibility dominates any systematic uncertainty from muon beam instability. Second, the low polarisation sets were found to be stable, which will be important when the fringe field systematic uncertainty is determined in Chapter 6. Lastly, a handful of runs were excluded due to clear changes in the internal beam.

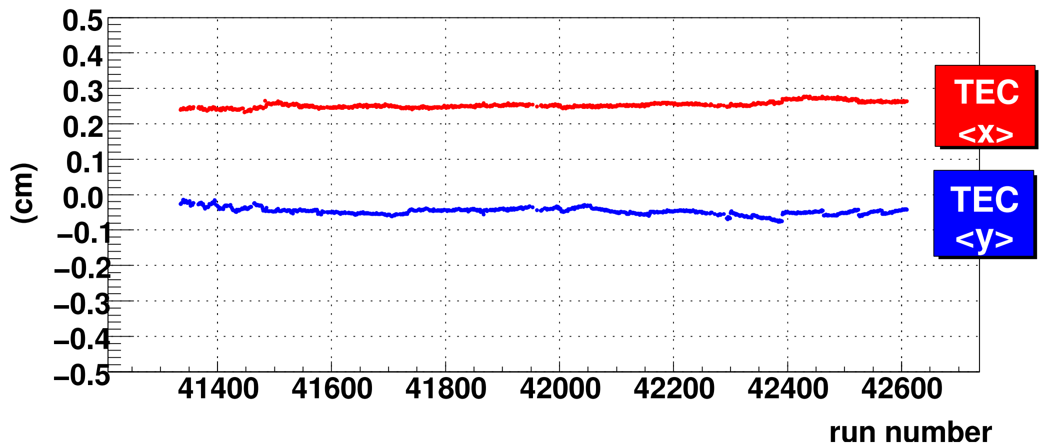

(a) TEC average position.

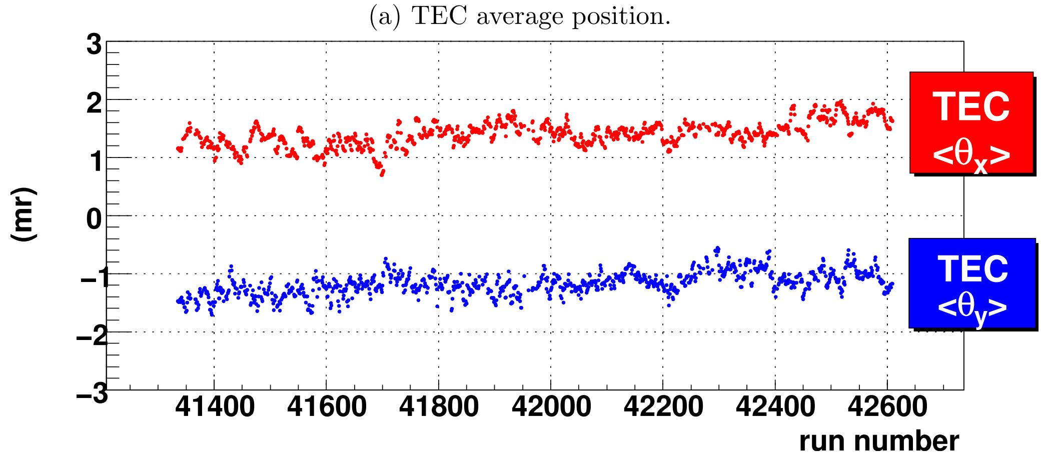

(b) TEC average angle.

Figure 5.2: The mean position and angle of the beam at the TECs were stable over the course of a set ($\approx$ 1 week). The internal beam had no sensitivity to fluctuations within the set that is shown.

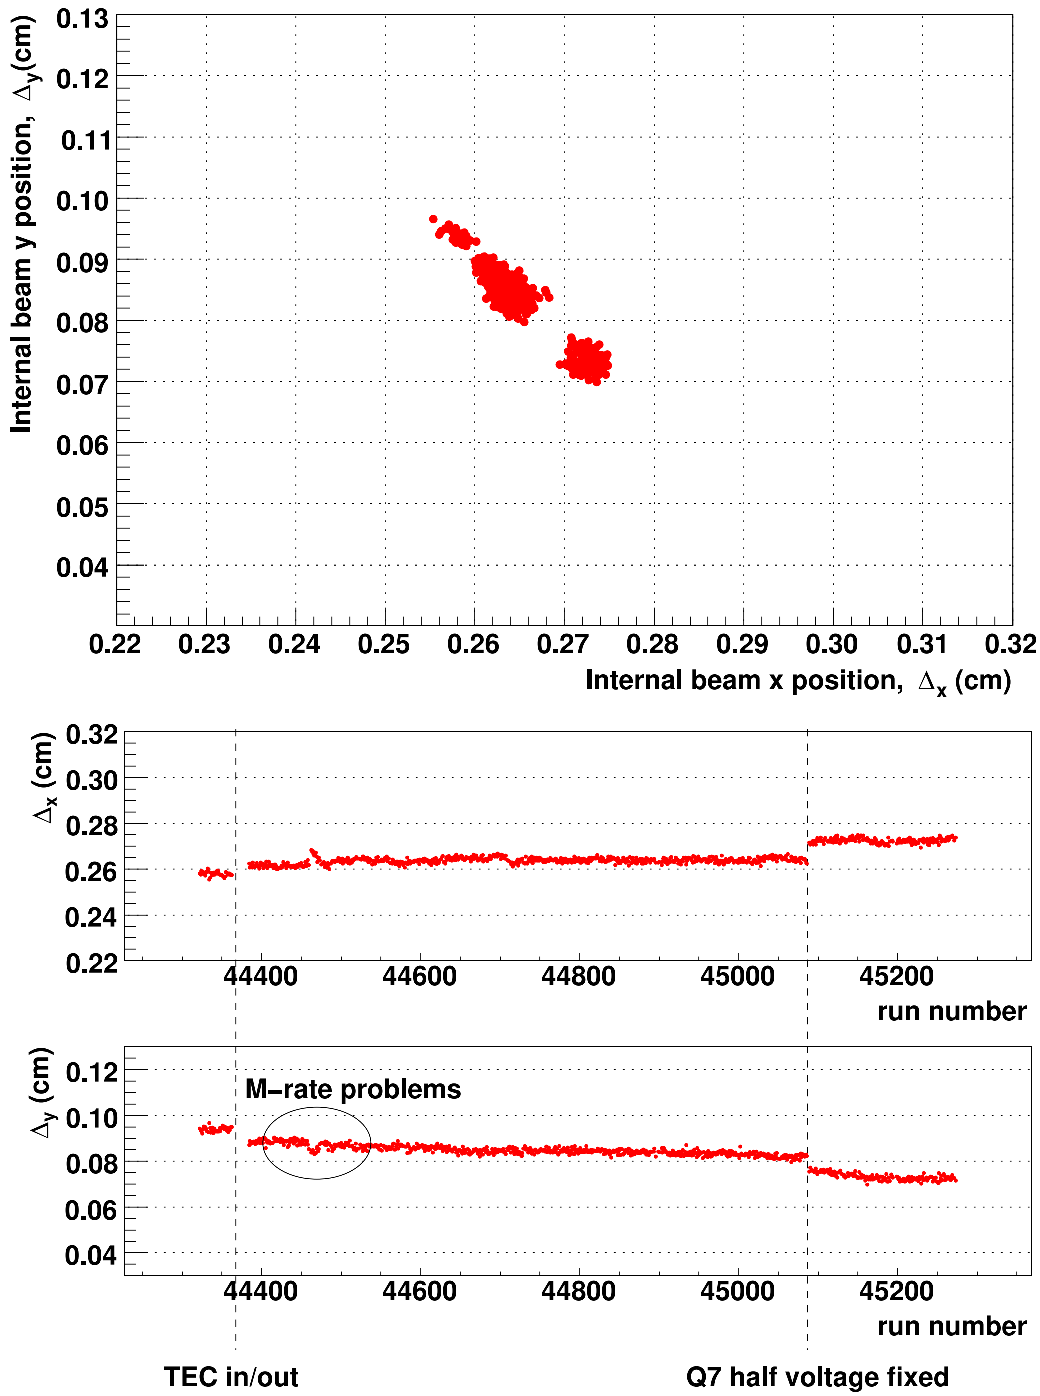

Figure 5.3: Stability of the internal beam measurement for a single nominal data set. The average position, (Δ_x, Δ_y), is stable at the level of $< 0.02\,\mathrm{cm}$. This set is chosen since there are systematic steps observed in the average position.

Chapter 6

Systematic Uncertainties and Corrections

6.1 Introduction

The $P_\mu^\pi \xi$ uncertainties are summarised in Table 6.1. The entries will be described fully in the current chapter. The table lists *systematic* uncertainties that are unique to the polarisation at the time of decay; the other systematic uncertainties are determined simultaneously with the parameters ρ and δ. The largest uncertainty is from the P_μ simulation: the muon beam and magnetic field together determine the muon's polarisation at target entry, and our knowledge of these limited the accuracy of the final result. The table has two *statistical* uncertainties; these could be reduced by accumulating more data and/or simulation under the same running conditions.

There are four corrections to $\Delta P_\mu^\pi \xi$ (the difference in $P_\mu^\pi \xi$ between the data and a hidden simulation value):

1. The simulation does not include depolarisation of the muons while leaving the muon production target, resulting in a correction of $+0.9 \times 10^{-4}$ to the nominal sets, and $+5.9(5.2) \times 10^{-4}$ for the sets at $\langle p \rangle = 28.75\,\mathrm{MeV/c}$ ($28.85\,\mathrm{MeV/c}$). The uncertainty from these corrections is included in Table 6.1.

2. The simulations were generated with an incorrect rate for the time-dependent depolarisation. A correction is applied rather than re-generating the simulation, since this would require several months of computer processing. The correction is $+2.9 \times 10^{-4}$ for silver (2006) and $+2.4 \times 10^{-4}$ for aluminium (2007). The uncertainty of this correction is negligible.

3. The statistics in the simulation spectra exceed the data by a factor of 1.8 to 4.0, depending on the set. This introduces a very small bias in the central value of $\Delta P_\mu^\pi \xi$ during the spectrum fitting (see Section 3.4 for a description of the fitting

technique). As a result, a set-independent correction of -0.5×10^{-4} is applied. The uncertainty of this correction is $\pm 0.5 \times 10^{-4}$, which is too small to enter Table 6.1.

4. The excess statistics in the simulation also causes a bias in the energy calibration procedure. This is applied as a set-dependent correction of between $+1.3 \times 10^{-4}$ and $+2.3 \times 10^{-4}$; the exact numbers will be given in Chapter 7. The uncertainty of this correction is $< 0.5 \times 10^{-4}$, which is again too small to enter Table 6.1.

Corrections 1 and 2 will be described in this chapter. Further detail on corrections 3 and 4 can be found elsewhere[81, 99].

The categories of uncertainty in Table 6.1 are chosen to represent uncorrelated contributions, which allows us to sum them in quadrature. There are two exceptions: the uncertainties due to the drift chamber space-time relationship (Section 6.3.1) and the positron interactions, where a correlation exists with the resolution systematic uncertainty, and an appropriate subtraction is made (see Sections 6.3.1 and 6.5.1 for quantitative results). There is an asymmetric systematic uncertainty for the polarisation: we place our result's central value at the most probable value, and treat each side of the uncertainty as Gaussian distributions with different widths. Note that the total quadrature sums at the bottom of the table are separated into systematic and statistical pieces, and there are two statistical entries in the table, both of which must be included to obtain the correct totals.

Table 6.1: Uncertainties for $\Delta P_\mu^\pi \xi$, the difference in $P_\mu^\pi \xi$ between the data and a hidden simulation value. These are systematic unless marked (stat.).

Category	Section	$\Delta P_\mu^\pi \xi$ uncertainty ($\times 10^{-4}$)		
		This analysis	MacDonald '08 [10, 85]	Jamieson '06 [22, 59]
Polarisation				
Magnetic field map and μ^+ beam	6.2.1 to 6.2.9	**-4.0,+15.8**	Not eval.[a]	34.0
Stopping material				
Systematic effects	6.2.10	**3.2**	Not eval.	12.0
λ (stat.)	6.2.10	2.4	Not eval.	Not eval.
μ^+ production target	6.2.11	0.3	Not eval.	2.1
Background muons	6.2.12	1.0	Not eval.	1.8
Chamber response				
DC space-time-relationship	6.3.1	0.9	6.0	Not eval.
DC geometric effects	6.3.2	1.3	0.7	2.2
US-DS efficiency	6.3.3	1.4	1.1	1.9
Crosstalk	6.3.4	0.5	Not eval.	Not eval.
Wire time offsets	6.3.5	0.8	0.4	8.9
Detector alignment				
z length scale	6.4	0.0	0.7	2.2
u/v width scale	6.4	0.2	0.2	Not eval.
Positron interactions				
δ-electron rate	6.5.1	0.1	1.4	2.9
Bremsstrahlung rate	6.5.1	0.5	0.03	
Outside material	6.5.2	0.4	0.6	0.2
Resolution	6.6	1.5	0.7	Not eval.
Momentum calibration				
Tracking B-field	6.7.1	0.3	1.1	0.9
Kinematic endpoint				
Model uncertainty	6.7.2	1.4	Not eval.	Not eval.
Propagation to bulk	6.7.2	0.5	0.01	1.6
Beam stability				
Muon beam intensity	6.8.1	0.3	0.2	1.8
External				
Radiative corrections	6.9.1	0.5	0.5	1.0
η correlation	6.9.2	1.1	1.1	Not eval.
Extraction of $\Delta P_\mu^\pi \xi$ (stat.)[b]	7.1	2.5	3.7	6
Total systematic		**-6.3,+16.5**	-	38
Total statistical		**3.5**	-	6

[a] MacDonald '08 measured ρ and δ, but not $P_\mu^\pi \xi$.

[b] This includes the statistical uncertainty from the kinematic endpoint parameters (see Section 6.7.2).

6.2 Polarisation

6.2.1 Overview

The simulation transports the muon spin from the position of the TECs ($z \approx -191\,\mathrm{cm}$) to the metal stopping target ($z \approx 0\,\mathrm{cm}$). This relies on the accuracies of the muon beam measurement and the magnetic field map, both of which will contribute to the assessment of the $P_\mu^\pi \xi$ uncertainty. The term "fringe field" will be widely used; this refers to the magnetic field from the end of the M13 beam line up to the first drift chamber.

The systematic uncertainties from the muon beam and fringe field are summarised in Table 6.2, in the order they appear in this chapter. First, the uncertainties in the alignment of the muon beam and the magnetic field map will be described. These will be used simultaneously to establish an uncertainty of $^{+6.4}_{-1.2} \times 10^{-4}$. Second, the uncertainties in the transverse magnetic field components will be evaluated using three consistency tests, and these will set an uncertainty of $^{+13.9}_{-0.0} \times 10^{-4}$. Third, uncertainty in the width of the muon beam angular distributions will lead to three uncertainties that total $\pm 3.8 \times 10^{-4}$.

Table 6.2: Summary of muon beam and fringe field uncertainties, for sets with a nominal beam tune.

Description	$\Delta P_\mu^\pi \xi$ uncertainty ($\times 10^{-4}$)
μ^+ beam and magnetic field alignment	-1.2,+6.4
Magnetic field transverse components	-0.0,+13.9
μ^+ beam angular distributions	
Simulation of multiple scattering	$\pm$ 3.1
Noise from TEC electronics	$\pm$ 1.7
Aging of TEC sense planes	$\pm$ 1.5
Quadratic sum	**-4.0,+15.8**

6.2.2 Measures of polarisation

Throughout this chapter the muon beam's polarisation is always considered to be an average of the z-components of the muons' spins, *not* the average of the spin projected onto the momentum vector. Two measurements of polarisation will now be described[37].

In the simulation, the absolute polarisation is available *before* any time-dependent target depolarisation has taken place. This is the polarisation after passing through the solenoidal magnetic field, and it is calculated by simply averaging the z-components of the muons' spins. From here on the simulation's absolute polarisation is written as $P_\mu(0)$, where the (0) is a reminder that it is effectively determined at $t = 0$. Since the simulation knows the spin of each muon exactly, $P_\mu(0)$ is known with high precision using a relatively small number of muons.

A fit between two decay positron spectra (data or simulation) measures $\Delta P_\mu^D \xi$, where P_μ^D is the decay polarisation of the muon *after* time-dependent depolarisation has taken place. If the spectra have the same ξ and time dependence, then this quantity is directly comparable to changes in $P_\mu(0)$. Spectrum fits between the data with a nominal and steered beam will later be used to evaluate the leading systematic uncertainty for this measurement.

6.2.3 Initial beam position and angle alignment

A muon beam measurement was made with the TECs at the beginning and end of most data sets. These two measurements have small differences in position and angle, resulting in different values of $P_\mu(0)$. Since the `GEANT` simulation can only use one of the beam measurements, the differences must be assessed as a systematic uncertainty.

The changes in average position and angle between the two muon beam measurements are listed in Table 6.3, where differences of up to 0.18 cm in position and 3 mrad in angle are observed. Note that these changes are smaller than the observations from the previous TWIST $P_\mu^\pi \xi$ analysis[22]. The cause of the changes could be muon beam instability, a variation in the TEC drift cell response, or a limitation in the reproducibility of the position and angle of the TECs. Each of these possible causes will now be discussed.

Muon beam instabilities could originate from an instability in the M13 beam line

[37] A third measurement is possible using the time dependence of the forward-backward asymmetry. This was described in Section 3.6: $P_\mu(t) = P_\mu^*(0)\exp(-\lambda t)$ is fit to the asymmetry. The difference in $P_\mu^*(0)$ between two data sets is then a measure of their polarisation difference, as long as λ is the same for each set. Polarisation differences from $P_\mu^*(0)$ are *not* used in this chapter.

Table 6.3: Muon beam differences for the beginning and end of set TEC measurements. The temperature difference between the measurements is denoted by ΔT.

Set	Target	Description	$\Delta\langle x\rangle$ (cm)	$\Delta\langle y\rangle$ (cm)	$\Delta\langle\theta_x\rangle$ (mrad)	$\Delta\langle\theta_y\rangle$ (mrad)	ΔT^{a} (°C)
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	0.11	-0.05	0.2	-3.2	-0.3
70	Ag	B = 1.96 T	0.03	0.00	1.0	-0.4	-1.2
71	Ag	B = 2.04 T	0.09	-0.05	0.0	0.1	2.4
74	Ag	Nominal A[b]	-	-	-	-	-
75	Ag	Nominal B	0.04	-0.10	-0.5	1.5	3.2
76	Ag	Steered beam	-0.04	-0.06	-0.6	1.9	1.3
83	Al	Downstream beam package in place	0.12	-0.09	0.6	0.7	-0.3
84	Al	Nominal C	0.18	-0.15	0.2	1.4	-0.4
86	Al	Steered beam B	0.04	-0.01	1.0	-0.01	-0.4
87	Al	Nominal D	0.13	-0.11	-0.1	0.7	-1.3
91/92/93	Al	Lower momentum[b]	-	-	-	-	-

[a] $\Delta T = T_{\text{end}} - T_{\text{start}}$, so that $\Delta T > 0$ indicates a temperature rise between measurements.
[b] These sets only had one TEC measurement.

elements (*i.e.* quadrupoles, dipoles, slits, jaws, asymmetric currents for quadrupole steering), or a change in the proton beam upstream of the muon production target. We do not believe there were measurable M13 beam line instabilities since all elements were monitored with a slow control system, and runs where an element fluctuated were not analysed (see Section 5.4). The proton beam was ruled out as a significant source of instability. In detail, a special test was carried out with the proton beam displaced at the production target by ± 0.1 cm vertically, which is about five times larger than the beam could have moved during normal operation[38]. For this test, the largest observed TEC changes in the muon beam were $\Delta\langle y\rangle = \pm 0.07$ cm in position and $\Delta\langle\theta_y\rangle = \pm 1.0$ mrad in angle, which are negligible after scaling down by a factor of five.

We have good evidence that the muon beam was sufficiently stable throughout a set, and was therefore not the cause of the variations in Table 6.3. First, we acquired a whole week of data with the TECs in place throughout (set 72), and found that the average muon beam position and angle were stable to < 0.02 cm and < 1 mrad, respectively (see

[38]The proton beam was surrounded by four monitor plates (top, bottom, left, right). In order to steer the beam vertically by 0.1 cm and avoid destroying the top or bottom plate, the proton beam current was reduced from the nominal setting of $\approx 100\,\mu$A to $\approx 20\,\mu$A. This suggests that during normal operation the proton beam could not have moved by more than $\approx 20/100 \times 0.1$ cm $= 0.02$ cm.

Fig. 5.2). Second, when the TECs were not in place the muon beam measurement from the wire chambers was used to monitor the beam's stability; a typical monitoring plot was shown in Fig. 5.3. The sensitivity of this measurement was calibrated by deliberately changing the currents in each quadrupole and dipole by ±5%. This allowed us to conclude that the muon beam instabilities that were observed are all small and uncorrelated with any beginning/end of set changes in Table 6.3.

The space-time-relationship (STR) in the TEC drift cells depended on temperature. This was ruled out as the cause of the differences in Table 6.3. In detail, new STRs were prepared that corresponded to a ±3°C variation, and all the data from the TECs were then re-analysed. The average reconstructed positions changed by between 0.028 cm and 0.050 cm, depending on the proximity of the beam to the sense plane of the TEC[39]. The predicted changes in angle were all < 0.05 mrad, except for set 76, which was still predicted to change by only 0.4 mrad. Clearly these changes in position and angle cannot explain the larger differences observed in Table 6.3.

Lastly, the insertion and removal of the TECs required the beam line elements to be switched off, and a breaking of the vacuum in the beam line, which then had to be pumped down again before data could be taken with the TECs. This process exerted significant forces on the beam line components and the box containing the TECs, and these forces are the prime candidate for the measured variation in initial position and angle. Therefore we conclude that the muon beam itself was stable and no systematic uncertainty is necessary for its fluctuations. Instead an uncertainty from the initial position and angle of the muon beam is needed, since we cannot be sure that the alignment of the TECs was reproducible.

Later we will evaluate the $P_\mu(0)$ sensitivity to position changes of ±0.2 cm and angle changes of ±3 mrad. These are the limits of the observations in Table 6.3, but they are not overly conservative for a number of reasons. First, there are not enough entries in Table 6.3 to establish whether the changes in position and angle follow a predictable distribution. Second, the TECs were aligned to the drift chambers with a systematic uncertainty of about 2 mrad. And third, the long-term stability measurements of the TECs found movements of 0.1 cm, but they only had an accuracy of about 0.1 cm (see Section 2.13 for more information on the alignments and stability measurements).

[39]The x-positions increased with temperature and the y-positions decreased; see Ref. [78] for further detail.

6.2.4 Magnetic field position and angle alignment

The fringe field map for the analysis was generated with the `Opera` software package[87]. The generated map was compared to B_z measurements that were taken with Hall probes (see Appendix D). This comparison could not produce a precise translational alignment in x and y. Instead the position of the muon beam inside the detector was used. On a set-by-set basis, a field translation was determined such that the data and simulation positions matched. On average this required a translation of the entire map by[40] $(\Delta x, \Delta y) = (0.21, 0.15)$ cm. The uncertainties in this translation from the muon beam and an error in the strength of the magnetic field components were evaluated, and found to be ± 0.04 cm. This is smaller than the uncorrelated uncertainty from the Hall probe alignment; these were aligned in the yoke's coordinate system system to < 0.1 cm in x and y and < 1 mrad in θ_x and θ_y[100]. The next section will use the uncorrelated uncertainties from the Hall probe alignment to determine the systematic uncertainty for $P_\mu(0)$.

6.2.5 Alignment uncertainties from beam and magnetic field

Section 6.2.3 established that the alignment of the muon beam is uncertain at the level of ± 0.2 cm in both the x and y position, and ± 3 mrad in both the θ_x and θ_y position. Similarly, Section 6.2.4 established the alignment of the magnetic field map to be uncertain at the level of ± 0.1 cm in x and y, and ± 1 mrad in θ_x and θ_y. All of these uncorrelated uncertainties must be combined in order to determine the sensitivity of $P_\mu(0)$ to possible misalignments.

The dependence of $P_\mu(0)$ on each alignment is non-linear. This is demonstrated for a nominal beam in Fig. 6.1. In addition, since the experimental beam is optimised with respect to the solenoid axis, the quantity $P_\mu(0)$ is also optimised, and introducing a combination of misalignments is therefore more likely to reduce $P_\mu(0)$ than increase it. As a result of these complications, the uncertainties are combined using a Monte Carlo approach. One hundred simulations are run with each misalignment selected from an independent Gaussian distribution with the standard deviation equal to the alignment uncertainty. For example, the initial beam position in x is drawn from a Gaussian with a mean of 0.0 cm, and standard deviation equal to 0.2 cm. The resulting $P_\mu(0)$ distribution

[40] In practice, an early analysis found a preliminary translation of $(\Delta x, \Delta y) = (0.20, 0.20)$ cm, and this was used to generate all the simulations. The correction between the preliminary translation of $(\Delta x, \Delta y) = (0.20, 0.20)$ cm and the final value of $(\Delta x, \Delta y) = (0.21, 0.15)$ cm was evaluated to have a negligible effect on $P_\mu(0)$ ($< 0.6 \times 10^{-4}$ for the nominal sets.)

from these simulations is asymmetric, and the width above and below the central value is used to establish a $P_\mu(0)$ systematic uncertainty of $^{+1.2}_{-6.4} \times 10^{-4}$ for the nominal sets. The uncertainty for $\Delta P_\mu^\pi \xi$, the difference between the data and a hidden simulation value, is then $^{+6.4}_{-1.2} \times 10^{-4}$.

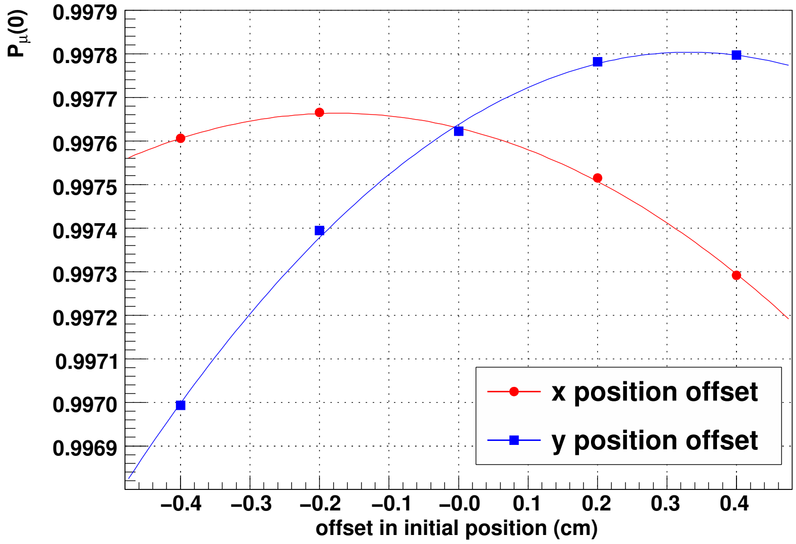

(a) Initial position sensitivity.

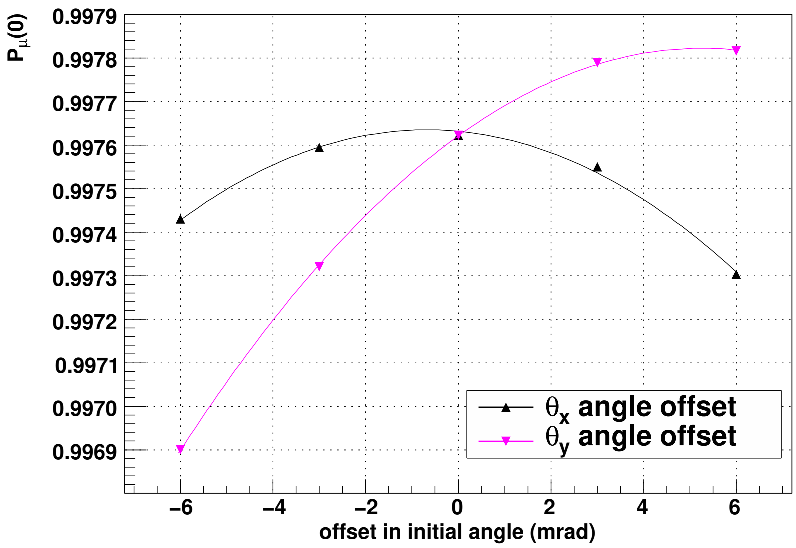

(b) Initial angle sensitivity.

Figure 6.1: The polarisation of the muons after the fringe field, $P_\mu(0)$, has a quadratic dependence on changes in the initial position and angle of the beam. The results of simulating a nominal beam profile (set 74) are shown.

6.2.6 Overview of the magnetic field map

The magnetic field map for the analysis was generated using the `Opera` software package[87]. The resulting components of $\vec{B}$ are shown in Figs. 6.2 and 6.3. The longitudinal components (B_z) increase steadily up to the drift chamber (DC) tracking region. The transverse components (B_x, B_y) are less than 1.5 mT while on-axis ($x = y = 0$), but increase significantly off-axis; Fig. 6.3(a) demonstrates this by including the components for ($x = y = 1$ cm). The transverse components have approximate radial symmetry and are maximised just inside the door of the yoke. They are also strongly linear within ≈ 5 cm of the axis, which fully encloses the volume through which the muons pass (see Fig. 6.3(c)). The transverse components are closely related to the depolarisation, which is shown for the simulation of a nominal profile in Fig. 6.3(b). For example, the onset of rapid depolarisation coincides with the maximisation of the transverse components, and the field's inflection at $z = -100$ cm is accompanied by an inflection in the depolarisation. Clearly the quality of the fringe field downstream of the door is important since it controls the rate of depolarisation. In addition, the field upstream of the door must be known since it affects the part of the fringe field which the beam is transported through.

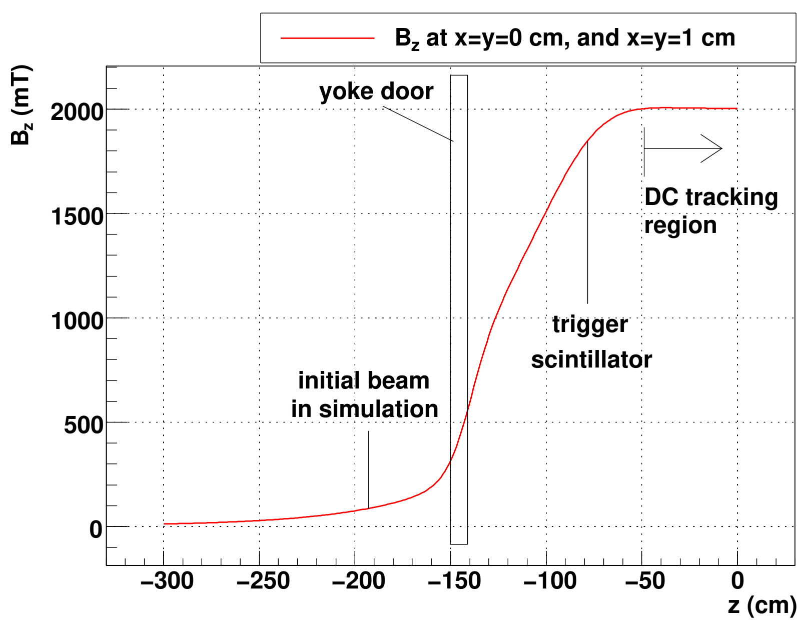

Figure 6.2: Longitudinal (B_z) components from the field map used for the analysis. These are produced from an `Opera` finite element analysis[87].

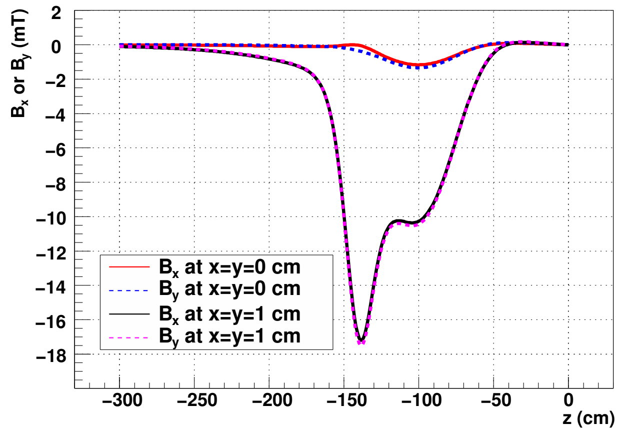

(a) B_x and B_y components.

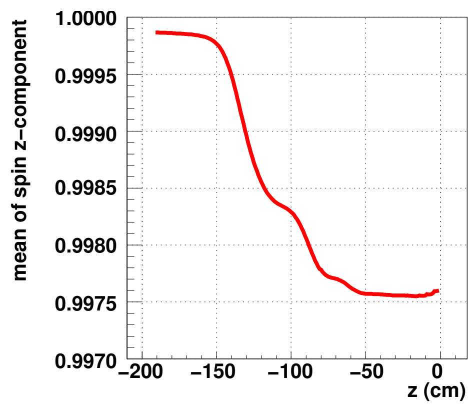

(b) Average of spin z-component, for a nominal beam profile.

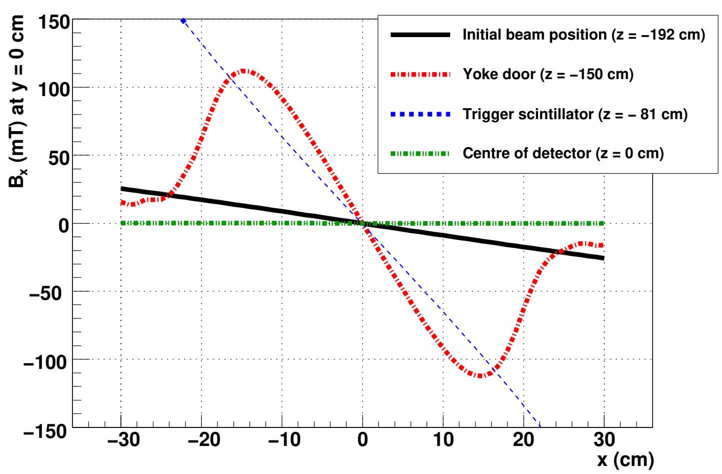

(c) B_x components for $y = 0\,\mathrm{cm}$.

Figure 6.3: Transverse fringe field components, B_x and B_y, from the field map used for the analysis. The average spin is also shown since its behaviour is closely related to the transverse components.

The B_z components from `Opera` are compared to Hall probe measurements in Fig. 6.4, where discrepancies of up to 6 mT are observed. The same comparison is not possible for the smaller B_x and B_y components since the Hall probes were single axis and did not measure them. We have identified three possible causes for the discrepancies. First, the `Opera` software did not include all the boundary conditions, such as the steel in the floor of the M13 area and the final M13 quadrupoles; an attempt was made to include these quadrupoles, but it did not find a significant effect for the B_z components. Second, there were several inputs to `Opera` that had to be tuned within their measured accuracy. For example, adjustments were made to the solenoid's coil positions in $x/y/z$, the radii of these coils and their current densities, the $B - H$ curve for the iron yoke, and the position of the yoke door in z[101, 102, 103, 104]. There was freedom of up to 0.2 cm in the positions of the solenoid coils[101], since their positions were not measured, and were instead initially placed according to a sketch from the magnet manufacturer. The previous $P_\mu^\pi \xi$ measurement found that variations in the `Opera` inputs affected $P_\mu(0)$ by 3×10^{-4} at most[59]. As a result of this low sensitivity the `Opera` inputs were not tuned further for the current measurement. Third, the `Opera` software used a finite element method to solve Maxwell's equations. This is expected to have limitations when modelling the 40 cm diameter circular hole in the yoke door at $z = -150$ cm, which is at a critical region for the *transverse* field components and hence the depolarisation. This is because there are several orders of length scale involved, and it is reasonable to expect that finite element analyses will suffer from accuracy problems under these conditions. In this case, the muons are within $\lesssim 5$ cm of the axis, there is a circular hole in the yoke of diameter 40 cm, and the whole map must be determined over a z-length of about 5 m.

If the discrepancy in the B_z components is eliminated, then the $P_\mu(0)$ estimates are changed by $< 1 \times 10^{-4}$ for all profiles; this test does not obey Maxwell's equations, but it does demonstrate that a discrepancy of a few mT in only the B_z components is unimportant at the 10^{-4} level. In contrast, a change of a few mT in the B_x and B_y components has a significant effect on $P_\mu(0)$. Since there are no Hall probe measurements of the B_x and B_y components, the uncertainty from them must be evaluated indirectly. We carry out this evaluation by comparing how well the simulation reproduces large differences in polarisation between sets with a nominal and a steered muon beam. The next sections will describe this technique.

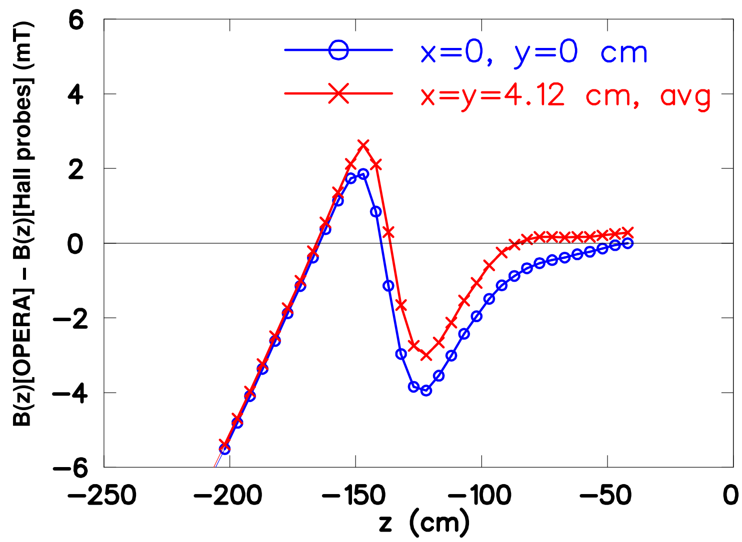

Figure 6.4: Difference in B_z between the magnetic field maps from `Opera` and the Hall probes. Two comparisons are shown: on-axis ($x = y = 0$) and an off-axis average of $x = \pm 4.12\,\mathrm{cm}, y = \pm 4.12\,\mathrm{cm}$. Note that muons start in the simulation at $z = -191.944\,\mathrm{cm}$, and stop in the target at $z = 0\,\mathrm{cm}$.

6.2.7 Muon beam tunes for evaluating the uncertainty

The nominal beam tune was described in Section 5.2; initially the position of the beam at the time expansion chambers (TECs) was steered to be close to $x = y = 0$, with the angles θ_x and θ_y minimised, but ultimately the tune was chosen to place the muon beam spots within the detector along a straight line, corresponding to minimised transverse momentum. The "envelope" of a simulated nominal muon beam, defined as the mean plus or minus one standard deviation, is shown in Figs. 6.5(a) and 6.5(b). Most of the beam remains within 1.0 cm of the solenoid's axis, and is focussed by the field to a few millimetres in extent at $z = -100\,\mathrm{cm}$. After this focus the envelope develops oscillations in its mean position and size.

Three additional beam tunes are shown in Fig. 6.5. These will be used to assess the systematic uncertainties for the fringe field. They are set 76, where the muon beam was steered to have $\langle\theta_y\rangle \approx 28\,\mathrm{mrad}$ at the TECs, set 86, where the beam was positioned off-axis and pointed away from the axis ($\langle x\rangle \approx -1.0\,\mathrm{cm}$, $\langle\theta_x\rangle \approx -10\,\mathrm{mrad}$), and set 72 where the TECs were in place throughout, which increased the muon beam's emittance due to the additional multiple scattering. In the range where rapid depolarisation takes place, $(-150 < z < -100)$ cm, the beams for sets 76 and 86 sample the fringe field at a different location to the nominal beam, with a non-zero average angle. They undergo a considerable focus, corresponding to an intersection of the field lines at a steep angle. As a result, these beams will be seen to undergo significantly more depolarisation, and the ability of the simulation to reproduce the depolarisation observed in the data will provide a stringent test of the fringe field model; such a validation was not available for the previous TWIST $P_\mu^\pi \xi$ analysis[22].

The three pairs of sets are listed in Table 6.4, and from here-on they are labelled as comparisons I, II and III. The first column in Table 6.4 shows the differences in polarisation from data. These are determined by fitting the nominal and steered decay spectra against each other, and the uncertainties are purely statistical. The next column in Table 6.4 shows the simulation's prediction for the polarisation difference using the `Opera` magnetic field map. These are evaluated using a Monte Carlo technique that is similar to Section 6.2.5. For each Monte Carlo trial, the initial position and angle of the beam are selected independently for the nominal and steered set, since we believe that variations originate from the reproducibility of the TECs. In contrast, the magnetic field position and angle must be the same for a nominal and steered set, so the same random values are used for both sets in each Monte Carlo trial.

In terms of standard deviations, the simulation is seen to *underestimate* the depolarisation by $2.0\,\sigma$ and $0.6\,\sigma$ for comparisons I and II, and to *overestimate* the depolarisation by $1.0\,\sigma$ for comparison III. The result for comparison I would not usually be a cause for concern; however, there is good reason to believe that an additional systematic uncertainty from the B_x and B_y components is missing.

Figures 6.3(a) and 6.3(c) show how the transverse field components depend on position. The muons are confined to within < 5 cm of the axis, and Fig. 6.3(c) indicates that the field has a linear dependence on x (and y) over this range. It is reasonable to expect that `Opera` has the correct linear shape, but may suffer from errors in the magnitude of the components. The sensitivity of comparison I to a scaling of B_x and B_y by a common factor, f, is shown in Fig. 6.6. The discrepancy of $2.0\,\sigma$ between data and simulation can be reduced by increasing f, which increases the central value of the simulation, and enlarges the uncertainties in $P_\mu(0)$ from alignments.

The systematic uncertainty is evaluated by considering the effect of the scaling, f, on all three comparisons. The number of standard deviations by which data and simulation differ is shown in Fig. 6.7(a), as f is varied. A decrease in the transverse field components is limited by comparison I, since the simulation rapidly diverges from the data as

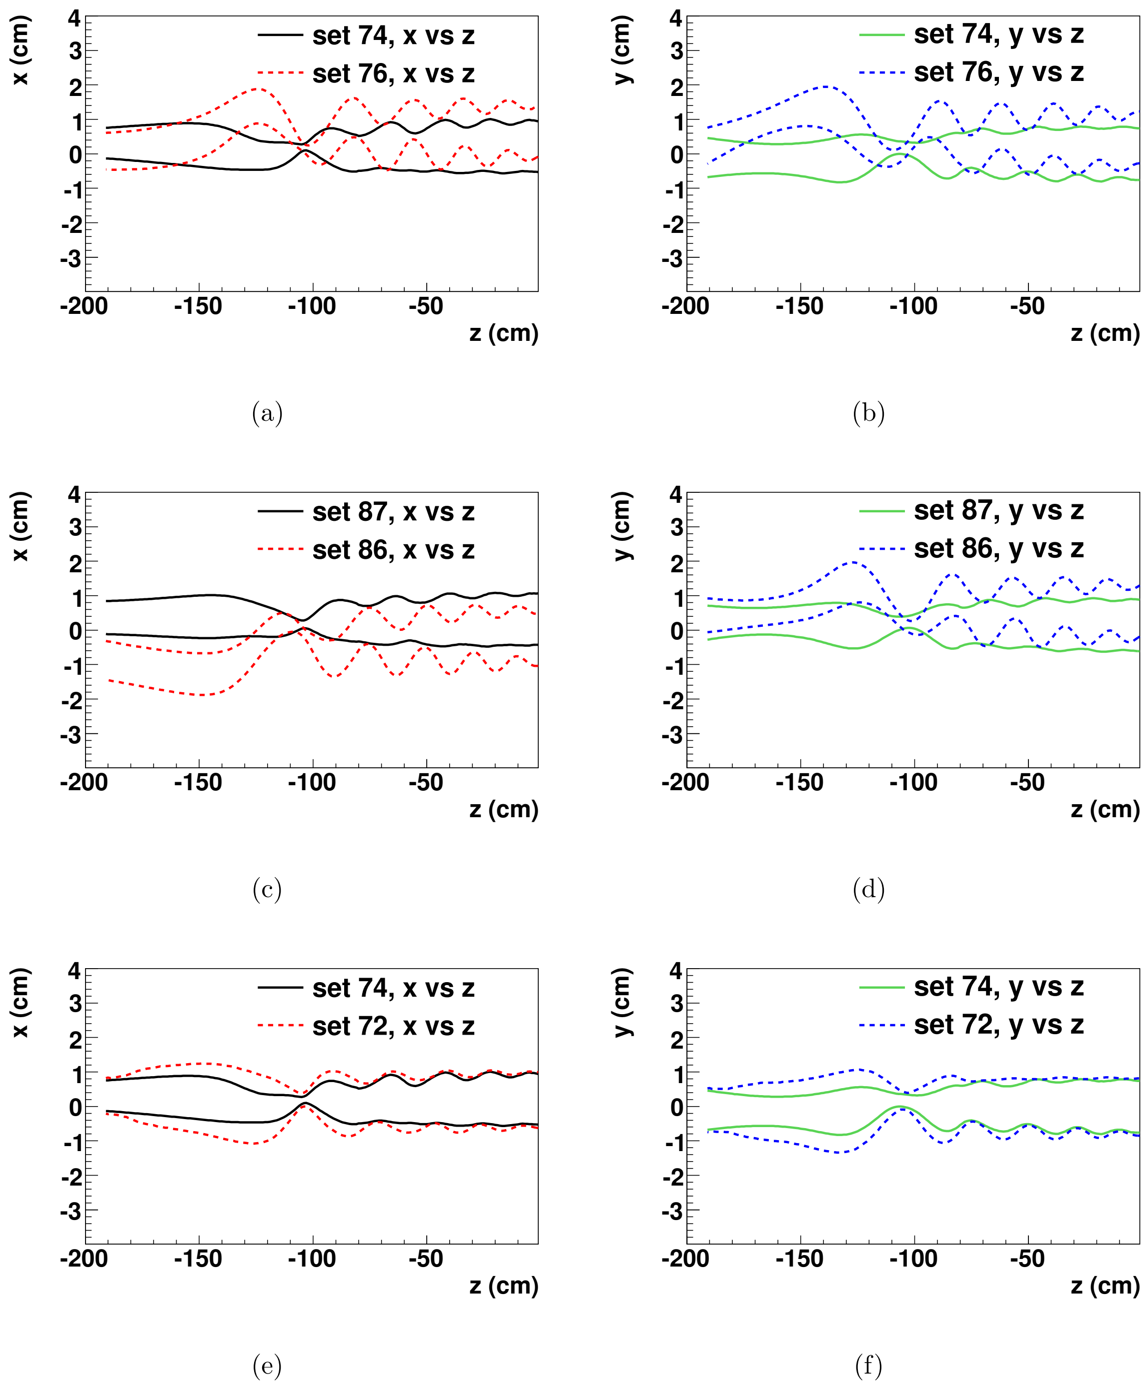

Figure 6.5: Muon beam *envelopes* from the simulation. These are defined as the mean position plus or minus one standard deviation.

Table 6.4: Validation of the fringe field map by establishing how well the simulation reproduces large polarisation differences in the data. Ideally the value of simulation minus data should be consistent with zero.

		Difference in polarisation ($\times 10^{-4}$)		
Comparison	Sets	Data	Simulation	Simulation minus data
I	74,76 (steered in θ_y)	105 ± 9	56^{+23}_{-18}	-48^{+24}_{-20}
II	87,86 (steered in x, θ_x)	62 ± 8	47^{+23}_{-16}	-15^{+24}_{-18}
III	74,72 (TECs-in)	18 ± 9	28^{+12}_{-5}	$+10^{+15}_{-10}$

f is decreased. Similarly, an increase in the transverse field components is limited by comparison III.

We have no reason to favour one of the comparisons over the others. Therefore we combine the differences between data and simulation for each comparison, and seek to minimise this combination by scaling the transverse field components. Specifically the χ^2 statistic is calculated as

$$\chi^2 = \sum_{\text{comparisons}} \frac{(\bar{x}_{\text{data}} - \bar{x}_{\text{simulation}})^2}{\sigma^2_{\text{data}} + \sigma^2_{\text{simulation}}} \tag{6.1}$$

where for convenience the definition $x = \Delta P_\mu(0)$ has been made, and $\Delta P_\mu(0)$ is the difference in polarisation between a nominal and a steered set. The resulting χ^2 statistic is shown in Fig. 6.7(b); a minimisation occurs when the transverse field components are scaled by $f = 1.10$. A correction of the result to this minimum is *not* made, since we cannot prove that the deviations between the `Opera` and true fields are fully covered by a simple scaling of the B_x and B_y components. Instead the f range corresponding to $(\chi^2_{\text{minimum}}+1)$ is taken to present a reasonable maximum of our error due to not measuring the transverse field components. This leads to a range in f of 1.01 to 1.20. All nominal sets are evaluated at these two f-values, and the average difference in $P_\mu(0)$ is 13.9×10^{-4}. This is assigned as a one sided uncertainty for $P_\mu(0)$ of ${}^{+0.0}_{-13.9} \times 10^{-4}$. The uncertainty for $\Delta P^\pi_\mu \xi$, the difference between the data and a hidden simulation value, is then ${}^{+13.9}_{-0.0} \times 10^{-4}$.

We note that early studies used a different field map to investigate the systematic

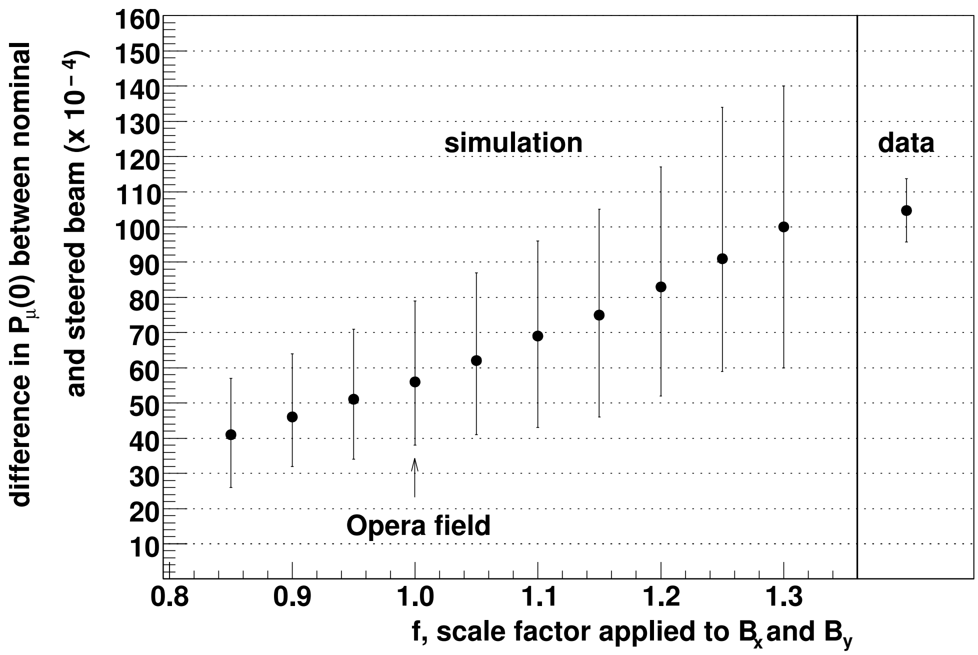

Figure 6.6: Comparison I is shown. With the original `Opera` field ($f = 1.0$) the data and simulation differ by about two standard deviations. The agreement improves as the B_x and B_y components are scaled by a factor, f.

uncertainty from the magnetic field. The fields from three on-axis coils were added to the original `Opera` map, and their currents were tuned to almost completely eliminate the B_z mismatch in Fig. 6.4. The coils also increased the B_x and B_y components, and the effect on the nominal set polarisation was 14.1×10^{-4}, on average. This is very close to the result derived from the simple scaling approach.

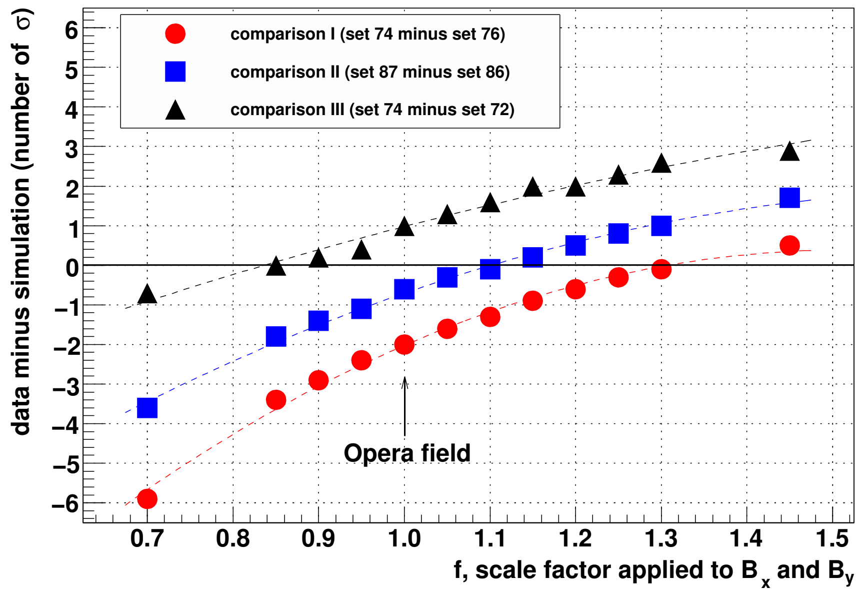

(a) Data-simulation agreement for all three comparisons.

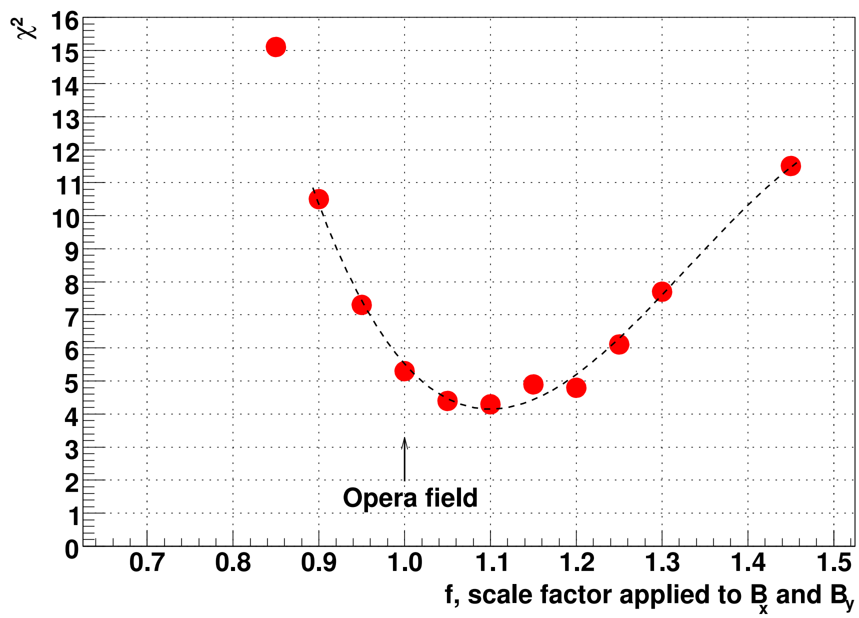

(b) Combining the three comparisons into a χ^2.

Figure 6.7: The transverse magnetic field components (B_x and B_y) are multiplied by a factor, f. The difference in polarisation between a nominal and a steered set is sensitive to f. An overall better agreement is possible with $f = 1.1$, corresponding to a 10% increase in the strength of the transverse components.

6.2.8 Muon beam angular distribution width

The muons were multiple scattered as they passed through the TECs, resulting in a measured angle distribution that is larger than the distribution in the absence of the TECs. A `GEANT3` simulation of the TECs finds that the root mean square of the angles should be reduced by a factor of $c_x = 0.64$ in the $x-$module, and $c_y = 0.48$ in the $y-$module to account for this multiple scattering. The correction in the y-module is larger since it is located downstream of the x-module. The dependence of $P_\mu(0)$ on the choice of c_x is shown in Fig. 6.8 for a nominal and a steered beam. To an acceptable approximation, $dP_\mu(0)/dc_x$ and $d^2P_\mu(0)/dc_x^2$ are independent of the beam steering. For example, if $c_x = 0.64$ then a variation of $\pm 10\%$ in c_x changes $P_\mu(0)$ by ${}^{-1.8}_{+2.1} \times 10^{-4}$ for the nominal case, and ${}^{-1.9}_{+2.0} \times 10^{-4}$ for the steered beam. As a result, the choice of c_x and c_y has no bearing on the comparison of polarisation differences between data and simulation, and the systematic uncertainties from these factors can be treated as orthogonal to those already evaluated.

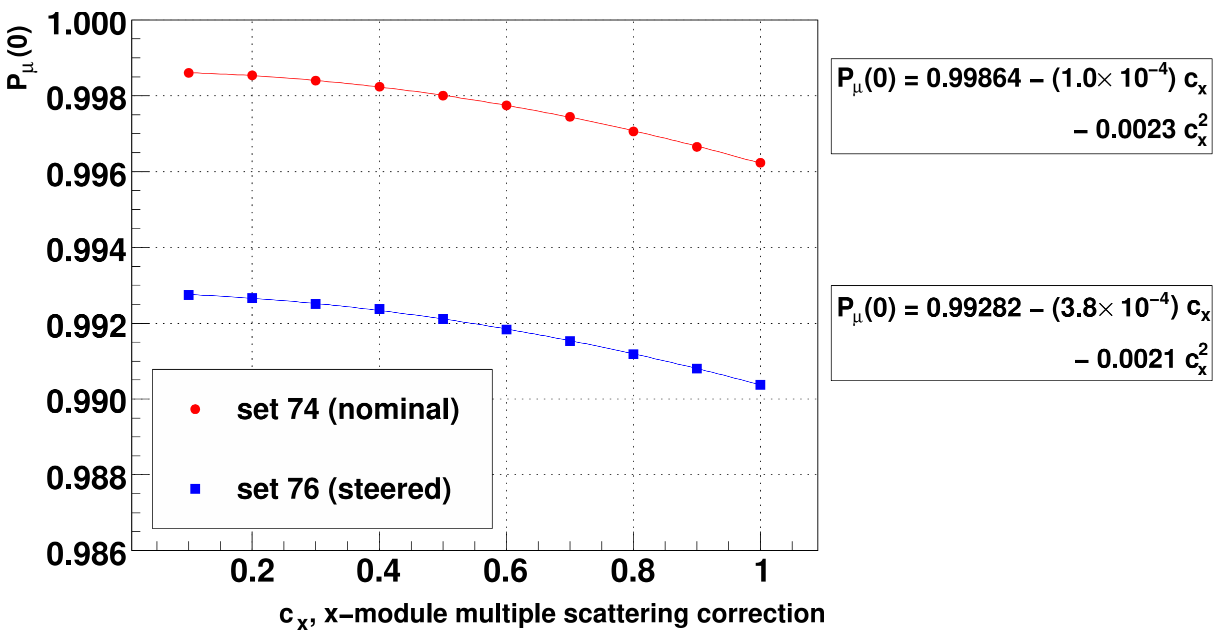

Figure 6.8: Sensitivity of the simulation's $P_\mu(0)$ to c_x, the multiple scattering correction factor in the x-module. The ratio $c_x/c_y = 0.64/0.48$ was maintained.

Note that although the simulation predicts the systematic uncertainties in this section are asymmetric, this asymmetry is at the 10^{-5} level, which is too small to be significant; the systematic uncertainties are therefore quoted as *symmetric*, with the magnitude set to the average of the upper and lower bounds.

The parameters c_x and c_y relied on the accuracy of multiple scattering in `GEANT3`. The author is unaware of any validation studies for the multiple scattering of muons with $p \approx 30\,\mathrm{MeV/c}$ in thin materials (the entire TEC apparatus was equivalent to just $\approx 7\,\mathrm{mg/cm^2}$ of material). Our most direct test of the `GEANT3` accuracy used five runs with the upstream window on the TECs changed from the nominal $6\,\mu$m of Mylar to a thicker $25\,\mu$m ($3.2\,\mathrm{mg/cm^2}$) window. The additional $19\,\mu$m of material increased the scattering distribution so that[41]

$$\theta_{\mathrm{TECs}+25\,\mu\mathrm{m}} \approx \sqrt{\theta^2_{\mathrm{TECs}+6\,\mu\mathrm{m}} + \theta^2_{19\,\mu\mathrm{m}}}. \tag{6.2}$$

The results for $\theta_{19\,\mu\mathrm{m}}$ from data and simulation are shown in Table 6.5, where the simulation overestimates the root mean square width of the scattering distributions by 18.3% in the x-module and 15.6% in the y-module. This implies that the c_x and c_y factors were reliable to 17.0% (the average of the overestimate in each module), which results in a systematic uncertainty of $\pm 3.1 \times 10^{-4}$. The observed discrepancy of 17.0% must not be taken as a formal validation of multiple scattering in `GEANT3`, since there are systematic uncertainties associated with the values in Table 6.5 that have not been evaluated.

Table 6.5: Width of reconstructed angle distributions for Mylar windows of thickness $6\,\mu$m and $25\,\mu$m. These are placed on the upstream window of the TEC gas box. The bracketed number indicates the statistical uncertainty in the final digit.

Mylar window	RMS of θ_x (mrad)		RMS of θ_y (mrad)	
thickness (μm)	data	simulation	data	simulation
6	14.50 (5)	14.10 (7)	19.79 (7)	19.89 (6)
25	17.25 (8)	17.90 (9)	22.01 (6)	22.8 (1)
$\Rightarrow$ 19	**9.3 (2)**	**11.0 (2)**	**9.6 (2)**	**11.1 (2)**

The TEC analysis code was reviewed for this measurement, and the accuracy of the reconstruction algorithm was found to be limited by noise from the electronics. This did not affect the mean position/angle, only the width of the angular distributions. An attempt to overcome this limitation resulted in two variants of the algorithm that are systematically different (see Appendix G.3.4). Since an event-by-event investigation could not distinguish which variant was the most accurate, the systematic difference between

[41] The multiple scattering distribution is non-Gaussian, and adding the widths of layers in quadrature will systematically underestimate the total width[3]. However, in this case we are comparing data and simulation where the same error is made in both.

the two is taken as a systematic uncertainty. For all sets this difference in $P_\mu(0)$ was less than 1.7×10^{-4}, except for set 76 (steered) where $P_\mu(0)$ changed by 6.3×10^{-4} between the two variants of the algorithm. A conservative systematic uncertainty of $\pm 1.7 \times 10^{-4}$ is assigned for the nominal sets.

The width of the angular distributions depended on the mean number of hits in the final track ($\langle n_x \rangle$ in the x-module, $\langle n_y \rangle$ in the y-module), which decreased depending on the length of time that the sense planes were exposed to the beam. The same c_x and c_y correction factors were used for all muon beam measurements, despite differences in the age of the planes, and this resulted in a systematic uncertainty. The values of $\langle n_x \rangle$ and $\langle n_y \rangle$ for each set are shown in Fig. 6.9. The c_x and c_y factors were tuned using set 75, which had $\langle n_x \rangle = 15.0$ and $\langle n_y \rangle = 15.9$. For all the sets, the ranges of $\langle n_x \rangle$ and $\langle n_y \rangle$ were 13.0 to 16.7, and 14.4 to 18.5 respectively, which is almost symmetric about the values used for tuning. The set 75 data were reanalysed, with hits removed at random to reduce $\langle n_x \rangle$ to 13.0 and $\langle n_y \rangle$ to 16.7. The root-mean-square decreased by 7.9% in the x-module, and 3.7% in the y-module. If the larger of these is used, then a systematic uncertainty for the nominal sets due to sense plane aging is $\pm 1.5 \times 10^{-4}$.

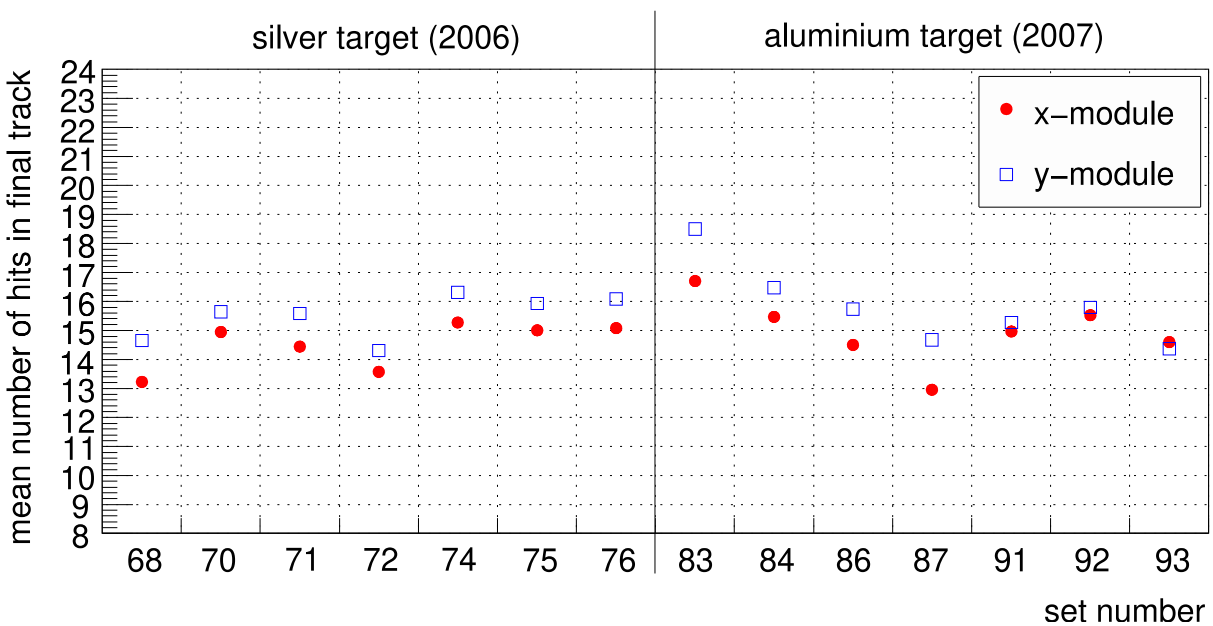

Figure 6.9: Mean number of hits in final TEC track, for each module. The c_x and c_y correction factors were tuned using set 75, which had $\langle n_x \rangle = 15.0$ and $\langle n_y \rangle = 15.9$.

6.2.9 Summary of muon beam and fringe field uncertainties

First the discrepancies between beginning and end of set muon beam measurements were considered (Section 6.2.3). These are blamed on a problem with the reproducibility of the TECs' alignment, rather than a muon beam instability or a variation in the TEC drift cell response. The resulting alignment uncertainties are $\pm 0.2\,\mathrm{cm}$ for position and $\pm 3\,\mathrm{mrad}$ for angle. Second, the alignment of the magnetic field in position and angle was considered (Section 6.2.4), leading to uncertainties of $\pm 0.1\,\mathrm{cm}$ for position and $\pm 1\,\mathrm{mrad}$ for angle. These are combined in Section 6.2.5 using a Monte Carlo approach, and produce an uncertainty of $^{+6.4}_{-1.2} \times 10^{-4}$ for $\Delta P_\mu^\pi \xi$.

The accuracy of the depolarisation through the magnetic field relies on the simulation's ability to reproduce large polarisation changes in three tests. The polarisation was lowered by either steering the muon beam away from the solenoid's symmetry axis (comparison I and II), or by inserting the TECs so that the muon beam's emittance was increased (comparison III); see Section 6.2.7 for a full description. A Monte Carlo approach was used to see how well the simulation reproduced the polarisation differences in the data. Using the original `Opera` field map, the largest discrepancy between data and simulation was $2.0\,\sigma$, which would not usually be a cause for concern. However, the `Opera` map is expected to differ from reality due to the quality of the initial conditions given, and the accuracy of the finite element method that is used to solve Maxwell's equations. Also the transverse field components (B_x and B_y) from `Opera` could not be validated directly since they were not measured. Therefore although we believe that the original `Opera` map is close to the truth, we must assign an uncertainty that covers a potential error in the field map. This was evaluated by scaling the transverse magnetic field components, and observing where the agreement in the comparisons is minimised. The resulting uncertainty for $\Delta P_\mu^\pi \xi$ is $^{+14.1}_{-0.0} \times 10^{-4}$.

Finally there are orthogonal systematic uncertainties from the degree to which the angle distributions of the muons could be reconstructed. There are contributions from the accuracy of multiple scattering in `GEANT3` ($\pm 3.1 \times 10^{-4}$), from the algorithm used to reconstruct the muon trajectories ($\pm 1.7 \times 10^{-4}$), and from the aging of the sense planes in the TECs ($\pm 1.5 \times 10^{-4}$).

6.2.10 Stopping material

Overview

About 80% of the muons stopped in a metal target, which also served as the shared cathode foil for the proportional chambers PC6 and PC7 (see Fig. 2.13). Events were accepted only if the muon produced a signal in PC6, but not in PC7. Muons that stopped in the PC6 gas (a mixture of CF_4 and isobutane) or the wires were removed by cutting on the muon pulse width in PC5 and PC6 (see Section 3.3.3). This selected a sample of muons that stopped in the metal foil. The relaxation was determined using a weighted asymmetry measurement (see Section 3.6). Each data set was fit with

$$P_\mu(t) = P_\mu^*(0)\exp\left(-\lambda t\right). \tag{6.3}$$

The relaxation rate from the data was used in the simulation.

There is one correction, one statistical uncertainty and four systematic uncertainties related to the stopping material. A correction to $\Delta P_\mu^\pi \xi$ is necessary since the simulation used an incorrect λ value that was derived from an early analysis; re-generating the simulation would have required several months, so a correction must be made instead. A statistical uncertainty results from the precision with which λ can be determined from the data ($\pm 2.4 \times 10^{-4}$). There is a systematic uncertainty from the degree to which stops in the PC6 chamber gas are eliminated ($\pm 0.3 \times 10^{-4}$), and another uncertainty from muons that pass through the target and stop in PC7, but do not have enough energy to produce a signal ($\pm 0.9 \times 10^{-4}$). A small fraction of muons are scattered from the target back into PC6, resulting in a systematic uncertainty of $\pm 0.2 \times 10^{-4}$. Lastly, there is an indication of a small bias in the asymmetry analysis, and this introduces a systematic uncertainty of ($\pm 3.0 \times 10^{-4}$). After adding in quadrature the total systematic uncertainty is $\pm 3.2 \times 10^{-4}$. Note that for this analysis there is no uncertainty due to the choice of model for $P_\mu(t)$ in the metals; see Section 1.6.3 for the theoretical arguments that demonstrate a strong preference for a single exponential form.

Correction and statistical uncertainty

The results of the asymmetry analysis applied to the data are shown in Table 6.6, where the fit is made over the nominal time range of $(1.05 < t < 9.00)\,\mu$s. A weighted average of these relaxation rates finds

$$\lambda_{\mathrm{Ag}}^{\mathrm{data}} = (0.813 \pm 0.079)\,\mathrm{ms}^{-1}, \tag{6.4}$$
$$\lambda_{\mathrm{Al}}^{\mathrm{data}} = (1.279 \pm 0.086)\,\mathrm{ms}^{-1}. \tag{6.5}$$

Note that these are consistent with the μ^+SR results from Appendix H.8:

$$\lambda_{\mathrm{Ag}}^{\mu\mathrm{SR}} = (0.9 \pm 0.2\,(\mathrm{stat.}) \pm 0.2\,(\mathrm{syst.}))\,\mathrm{ms}^{-1}, \tag{6.6}$$
$$\lambda_{\mathrm{Al}}^{\mu\mathrm{SR}} = (1.3 \pm 0.2\,(\mathrm{stat.}) \pm 0.3\,(\mathrm{syst.}))\,\mathrm{ms}^{-1}. \tag{6.7}$$

The simulation used incorrect values of $\lambda_{\mathrm{Ag}} = 0.725\,\mathrm{ms}^{-1}$ and $\lambda_{\mathrm{Al}} = 1.169\,\mathrm{ms}^{-1}$, and as a result the quantity $\Delta P_\mu^\pi \xi$ must be corrected. The effect on the spectrum of a change in λ is calculated using

$$\frac{\int_{t_1}^{t_2} N(t) \cdot P_\mu^*(0) \exp(-\lambda_2 t) dt}{\int_{t_1}^{t_2} N(t) dt} - \frac{\int_{t_1}^{t_2} N(t) \cdot P_\mu^*(0) \exp(-\lambda_1 t) dt}{\int_{t_1}^{t_2} N(t) dt}, \tag{6.8}$$

where $N(t) = N(0)\exp(-t/\tau_\mu)$, τ_μ is the muon lifetime, and λ_1 and λ_2 are the relaxation rates between which the correction is being made. The time range is the same as the normal analysis ($1.05\,\mu\mathrm{s} < t < 9.00\,\mu\mathrm{s}$). The common $P_\mu^*(0)$ factor is close to 1.0, and its choice has a negligible impact on the correction. The quantity $\Delta P_\mu^\pi \xi$ (the difference between the data and a hidden simulation value) is then corrected by $+2.7 \times 10^{-4}$ for Ag ($\lambda_1 = 0.725\,\mathrm{ms}^{-1}$, $\lambda_2 = 0.813\,\mathrm{ms}^{-1}$) and $+3.3 \times 10^{-4}$ for Al ($\lambda_1 = 1.169\,\mathrm{ms}^{-1}$, $\lambda_2 = 1.279\,\mathrm{ms}^{-1}$). The statistical uncertainty in determining λ from the data causes a P_μ uncertainty of $\pm 2.4 \times 10^{-4}$ for both targets, again using Eq. (6.8).

Table 6.6: Relaxation rate λ for each data set. The model $P_\mu(t) = P_\mu^*(0)\exp(-\lambda t)$ has been fit over the nominal time range of $(1.05 < t < 9.00)\,\mu$s.

Set num.	Target	Description	λ (ms^{-1})	Fit quality χ^2/ndof	Fit quality confidence
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	0.71 ± 0.24	1.03	0.43
70	Ag	B = 1.96 T	0.89 ± 0.19	0.62	0.90
71	Ag	B = 2.04 T	0.93 ± 0.20	1.58	0.05
72	Ag	TECs-in, nominal beam	0.85 ± 0.19	1.03	0.42
74	Ag	Nominal A	1.02 ± 0.24	0.63	0.90
75	Ag	Nominal B	0.85 ± 0.19	0.72	0.81
76	Ag	Steered beam A	0.38 ± 0.22	0.52	0.96
	Ag	**Weighted average:**	$\mathbf{0.813 \pm 0.079}$		
83	Al	Downstream beam package in place	1.46 ± 0.19	1.85	0.01
84	Al	Nominal C	1.10 ± 0.21	1.08	0.36
86	Al	Steered beam B	1.21 ± 0.18	1.08	0.36
87	Al	Nominal D	1.22 ± 0.20	0.85	0.65
91	Al	Lower momentum I	1.63 ± 0.40	1.06	0.39
92	Al	Lower momentum II	1.35 ± 0.35	0.98	0.49
93	Al	Lower momentum III	1.32 ± 0.27	0.75	0.78
	Al	**Weighted average:**	$\mathbf{1.279 \pm 0.086}$		

Muon stops in gas

The distribution of the pulse widths in the PCs immediately before the target was described in Section 3.3.3. It is shown here in Fig. 6.10. The majority of the distribution is made up of stops in the metal target, which primarily contribute to zones 1 and 2, and there is a "gas band" that runs approximately parallel to the cut-B line and contributes mostly to zones 2 and 3. The analysis selects only zone 1; the cut-B line removes almost all of the gas contamination, but a small fraction still makes it to zone 1.

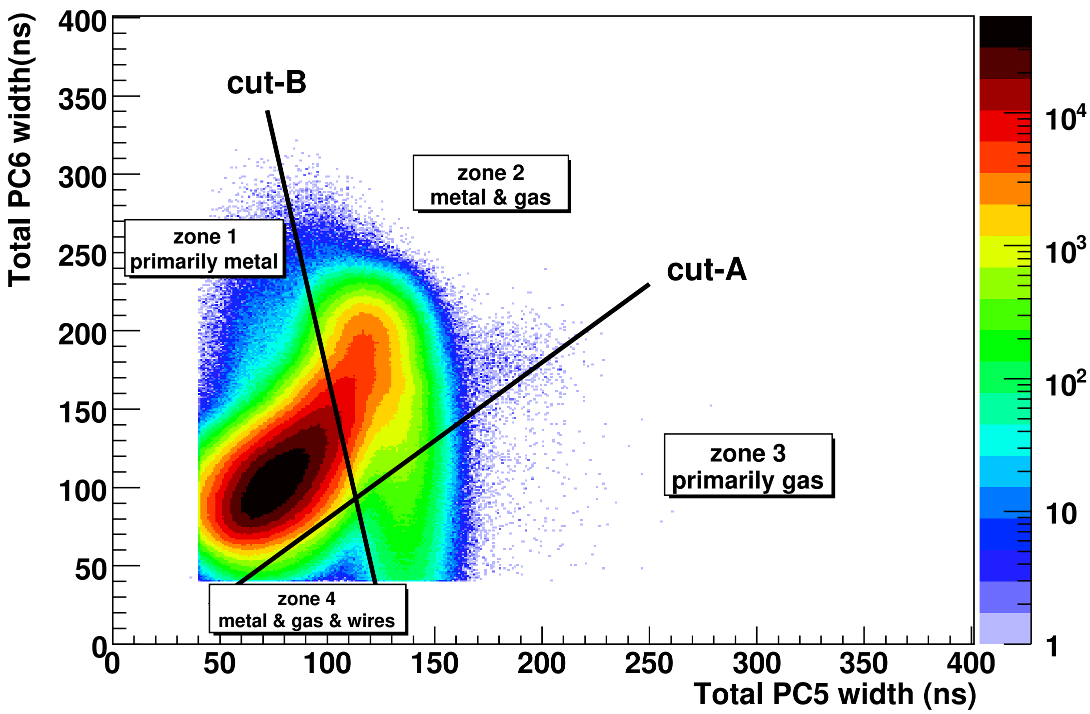

Figure 6.10: Muon pulse widths in PC5 and PC6 (the proportional chambers immediately before the metal stopping stopping target). For data, the stops in gas appear as a band that is approximately parallel to cut-B, primarily contributing to zones 2 and 3.

A systematic uncertainty was estimated using only the data. This required three numbers that are described in Appendix J: the depolarisation in gas (8.0%), the fraction of the total PC6+target stops in gas (f_g), which is between 4% and 7% depending on the set, and the fraction of the gas distribution in zone 2 that leaks into zone 1 (tuned to be $< 0.5\%$ by adjusting cut-B). If the number of muons stopping in the gas in zone i is denoted N_i^g, then

$$\begin{aligned} f_g N_{\text{total}} &= N_1^g + N_2^g + N_3^g + N_4^g, && (6.9)\\ N_1^g &< 0.5\% \times (N_1^g + N_2^g), && (6.10)\end{aligned}$$

which gives

$$\frac{N_1^g}{N_{\text{total}}} < 0.5\% \times \left[f_g - \frac{N_3^g + N_4^g}{N_{\text{total}}} \right]. \qquad (6.11)$$

If all the gas stops are conservatively assumed to be in zones 1 and 2 so that $N_3^g = N_4^g = 0$, and f_g is set to its maximum value of 7%, then an upper limit on the systematic uncertainty from muons stopping in PC6 gas is $8.0\% \times 0.5\% \times 7\% = 0.3 \times 10^{-4}$.

A special simulation found that no more than 0.11% of muons entered PC7 (the wire chamber *after* the metal stopping target), but did not have enough energy to produce a signal. These stop in the PC gas, and experience a depolarisation of 8.0% (see above). An upper limit for the sytematic uncertainty due to these stops is therefore $0.11\% \times 8.0\% = 0.9 \times 10^{-4}$. The same simulation determined that 0.02% of muons enter the target volume, but are scattered back into PC6, and pass all the analysis cuts. These contribute a negligible systematic uncertainty of $0.02\% \times 8.0\% = 0.2 \times 10^{-4}$.

Analysis bias

The asymmetry analysis is applied to the simulation to confirm that the input values of λ can be recovered ($\lambda_{\text{Ag}}^{\text{input}} = 0.725\,\text{ms}^{-1}$ and $\lambda_{\text{Al}}^{\text{input}} = 1.169\,\text{ms}^{-1}$). The results of this study are shown in Fig. 6.11, where the simulation that accompanies set 74 (nominal) is seen to have a λ value that is 3.4 standard deviations below the input value. As a result, the *weighted average* of λ for the silver simulations is 2.4 standard deviations from the input; the weighted average for the aluminium simulations is consistent with the input λ.

An exhaustive investigation attempted to find the cause of the discrepancies for the silver simulations. The results were found to be robust to the binning, weighting strategy and time range used in the asymmetry analysis. The confidence levels from the asymmetry analysis were spread evenly between 0 and 1, indicating that the λ uncertainties are properly determined. A separate analysis that made fits to the upstream and downstream time distributions (*i.e.* an asymmetry was not constructed) confirmed the λ results, but with inferior statistical precision. The simulation for set 74 only differed from the other nominal sets in its input beam profile and muon/positron rates; the beam profile produced a reasonable value for $P_\mu(0)$, and the rate differences from the other sets were insignificant.

We note that an additional statistically independent simulation of set 74 (not used here in the weighted averages) produced $\lambda = (0.59 \pm 0.17)\,\mathrm{ms}^{-1}$, which is consistent with both the input value ($\lambda = 0.725\,\mathrm{ms}^{-1}$) and the anomalous result from the earlier set 74 simulation ($\lambda = (0.32 \pm 0.12)\,\mathrm{ms}^{-1}$); clearly significant additional statistics are needed to make a firm conclusion. Since the investigation could not find a reason to reject the anomalous λ result, we are forced to admit the possibility that a systematic λ uncertainty exists due to an analysis bias, or the λ statistical uncertainties from the asymmetry analysis are underestimates.

We are not aware of a credible mechanism that would create a different bias for the silver and aluminium targets. Therefore the average discrepancy between the input and measured λ from the simulation ($0.5 \times (2.4 + 0.0)\sigma = 1.2\sigma$) is taken as an additional uncertainty that must be applied to the *data*. Using Eqs. (6.4), (6.5) and (6.8), this results in a systematic uncertainty of $\pm 2.9 \times 10^{-4}$ for silver and $\pm 3.1 \times 10^{-4}$ for aluminium. The average of these ($\pm 3.0 \times 10^{-4}$) is applied to all sets as the systematic uncertainty due to a potential analysis bias.

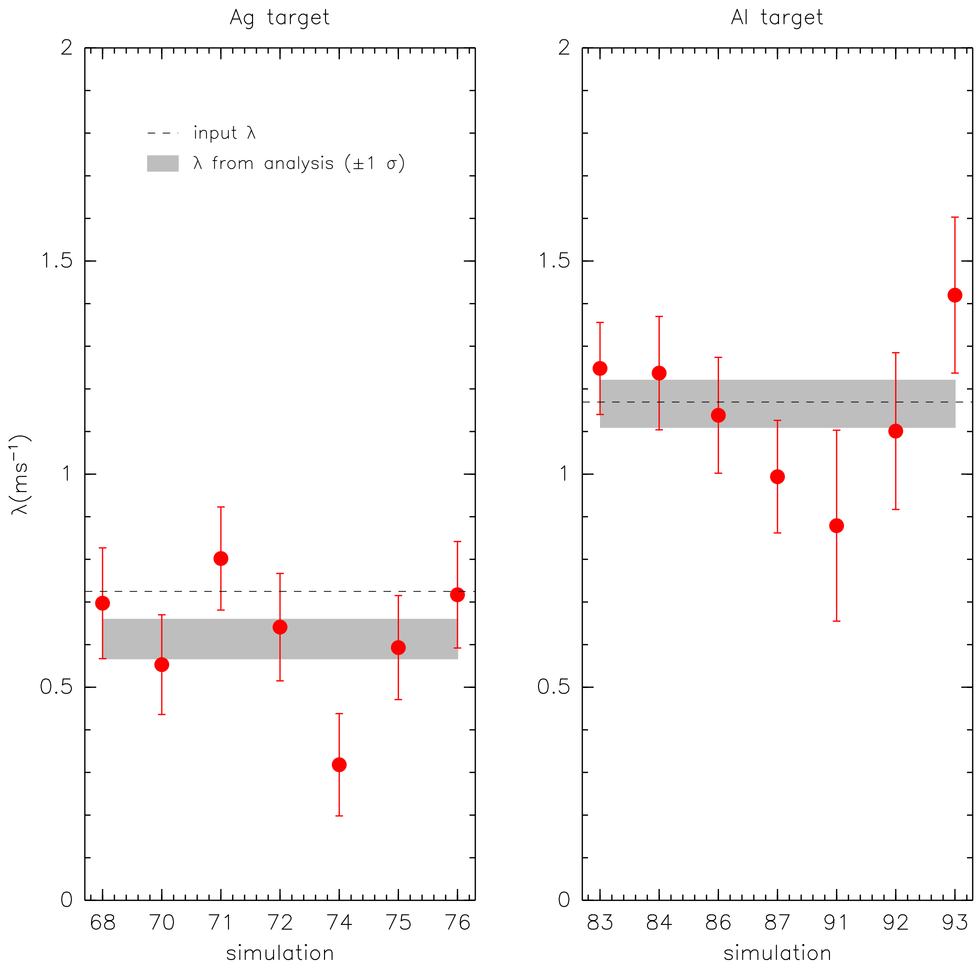

Figure 6.11: The result of the asymmetry analysis applied to the simulations[105]. The relaxation rate (λ) for the set 74 simulation is $3.4\,\sigma$ away from the input value.

6.2.11 Muon production target

The simulation generates muons with anti-parallel spin and momentum vectors, starting from the position of the TECs. This neglects multiple scattering in the muon production target and the beam line vacuum window, which changes the momentum vector but not the spin. This is treated here as a systematic correction with an associated uncertainty. (The difference in precession frequencies of the momentum and spin through the M13 beam line is neglected since it introduces an error of $< 10^{-8}$; see Appendix I).

Surface muons are produced with $p \approx 29.79\,\mathrm{MeV/c}$, but the beam line was nominally tuned to accept muons with an average momentum of $\langle p \rangle = 29.6\,\mathrm{MeV/c}$. Therefore the muons lost 0.19 MeV/c of momentum on average, which is equivalent to $\approx 3.8\ \mathrm{mg/cm^2}$ in graphite. (The $3\,\mu$m beam line vacuum window can be safely neglected since it corresponds to just 0.3 $\mathrm{mg/cm^2}$ of material.) The width of the resulting multiple scattering distribution[42], θ_0, is found to be 9.3 mrad using a `GEANT4` simulation[105]. The uncertainty is conservatively estimated as $\pm 17\%$ based on the observed discrepancy for `GEANT3` from Section 6.2.8. The degree to which the spin is depolarised with respect to the momentum is then estimated by $\cos(\theta^{\mathrm{rms}}_{\mathrm{space}})$, where $\theta^{\mathrm{rms}}_{\mathrm{space}} = \sqrt{2}\,\theta_0$. This results in a correction to the simulation's P_μ of $(-0.9 \pm 0.3) \times 10^{-4}$, which is a correction to $\Delta P^\pi_\mu \xi$ of $(+0.9 \pm 0.3) \times 10^{-4}$.

The `GEANT4` simulations are repeated for the lower momentum sets. The widths are $\theta_0 = 24.2\,\mathrm{mrad}$ for $\langle p \rangle = 28.75\,\mathrm{MeV/c}$, and $\theta_0 = 22.9\,\mathrm{mrad}$ for $\langle p \rangle = 28.85\,\mathrm{MeV/c}$. These correspond to $\Delta P^\pi_\mu \xi$ corrections of $(5.9 \pm 1.9) \times 10^{-4}$ and $(5.2 \pm 1.6) \times 10^{-4}$, respectively. The uncertainty for this correction will be applied set-dependently in Chapter 7, although in practice this detail makes no difference to the final answer.

[42] θ_0 is the standard deviation of a Gaussian fit to the central 98% of the plane-projected multiple scattering distribution.

6.2.12 Background muon contamination

The true muon stopping distribution is not available for the data, and we can only measure the last plane that registered a muon hit. In the previous TWIST $P_\mu^\pi \xi$ analysis[59] this distribution did not agree well in data and simulation unless an additional source of pion decays was included in the upstream beam package area (see Fig. 6.12(a)). An upper limit on the effect of these "background muons" was calculated as 1.8×10^{-4}, assuming they have opposite polarisation to the surface muons. For the current analysis the agreement is better *without* the additional pion source; this is demonstrated in Fig. 6.12(b), where improvements in the classification have increased the number of downstream events in the simulation relative to the data. If anything, the simulation now has *more* downstream muons for planes 35 onwards[43], and even this excess is at the level of $< 10^{-5}$. Therefore the background muon contamination is now reduced to a negligible level, and this part of the systematic uncertainty is zero.

However, the data-simulation discrepancies in planes 35 to 53 (and in planes 10 to 20) do introduce a systematic uncertainty; the simulation's stopping distribution must match the data, since high angle muons that undergo more depolarisation are preferentially stopped further upstream. Figure 6.12(b) indicates that the mean stopping position and/or its shape are not reproduced by the simulation. The most accurate measurement of the mean muon stopping distribution comes from the energy calibration procedure (see Section 3.5). The reconstructed spectrum endpoint depends on the thickness of target (and detector) material traversed, which results in a dependence on $1/\cos\theta$. The gradient of this relationship, α_{up} or α_{down} for upstream and downstream respectively, is therefore a measure of the material traversed by the upstream and downstream decay positrons. The stopping distribution can then be measured using

$$\alpha_{\text{diff.}} = \alpha_{\text{up}} - \alpha_{\text{down}}. \tag{6.12}$$

The energy calibration procedure finds the difference in reconstructed spectrum endpoint between data and simulation, which has a gradient equal to

$$\alpha_{\text{up}}^{\text{data}} - \alpha_{\text{up}}^{\text{sim.}}, \tag{6.13}$$

[43] Planes 53 to 56 in Fig. 6.12(b) suggest there are more muons in the data than the simulation. However, these are beam positrons that are mistakenly identified as muons in the data due to an imperfect separation of pulse width in the proportional chambers. This does does not occur in the simulation, where the separation between muons and beam positrons is perfect.

in the upstream half of the detector, and

$$\alpha^{\text{data}}_{\text{down}} - \alpha^{\text{sim.}}_{\text{down}} \tag{6.14}$$

in the downstream half. The difference between the gradients in the upstream and downstream half (*i.e.* Eq. (6.13) minus Eq. (6.14)) is then

$$\left(\alpha^{\text{data}}_{\text{up}} - \alpha^{\text{data}}_{\text{down}}\right) - \left(\alpha^{\text{sim.}}_{\text{up}} - \alpha^{\text{sim.}}_{\text{down}}\right) = \alpha^{\text{data}}_{\text{diff.}} - \alpha^{\text{sim.}}_{\text{diff.}}, \tag{6.15}$$

which measures how well the muon stopping distributions agree in data and simulation.

The stopping distribution measure from Eq. (6.15) is shown in Fig. 6.13. For five of the simulations, the mean of the muon last plane distribution was matched by adding an extra 1.9 mg/cm^2 of material to the simulation (see Section 2.11), and the resulting disagreement is (5.5 ± 1.4) keV/c on average. For the other simulations no additional material was added, and this improved the data-simulation disagreement to (0.2 ± 1.1) keV/c (We note that although $\alpha^{\text{data}}_{\text{diff.}} - \alpha^{\text{sim.}}_{\text{diff.}}$ is an accurate measurement of the stopping distribution, it is clearly not precise. For this reason it could not be used to tune the simulation so that it matched the data.)

The simulation can determine the relationship between $P_\mu(0)$ and $(\alpha^{\text{data}}_{\text{diff.}} - \alpha^{\text{sim.}}_{\text{diff.}})$. This depends on the beam profile, and is within the range $(0.04-0.25)\times10^{-4}/(\text{keV/c})$, with an average of $0.13 \times 10^{-4}/(\text{keV/c})$. The systematic uncertainty from matching the stopping distribution is therefore $\pm5.5 \times 0.13 \times 10^{-4} = \pm0.7 \times 10^{-4}$ for the sets with the extra 1.9 mg/cm^2 of material. This result is an approximation since it averages over the silver and aluminium targets, where differences are expected in the data-simulation agreement. Since this is a small number and only applies to 5 of the 14 simulations, it is treated as a conservative systematic uncertainty rather than a correction. An additional piece of the same magnitude is added in quadrature to account for a small difference in stopping distribution shape, resulting in a total systematic uncertainty of $\pm1.0 \times 10^{-4}$.

The remaining systematic uncertainties are related to the decay positron reconstruction. They are mostly evaluated by exaggerating a setting in the simulation or analysis. The exaggerated spectrum is then fit to the original spectrum in order to determine the change in the muon decay parameters (MPs). The exaggeration factors are made as large as possible to obtain statistically meaningful MP changes, while maintaining a linear relationship with the MPs. The changes in MPs are then scaled down according to how large the effect could actually be, resulting in the systematic uncertainty. When the original *and exaggerated spectra are highly correlated (i.e. they contain a large number of events with identical energy and angle), the spectrum fit will have a reduced* χ^2 *that is much smaller than one, and the uncertainty in the* $P_\mu^\pi \xi$ *change will therefore be too large. In this case the uncertainties in the MPs are scaled down so that the reduced* χ^2 *is equal to one, corresponding to multiplication by* $\sqrt{\chi^2/ndof}$.

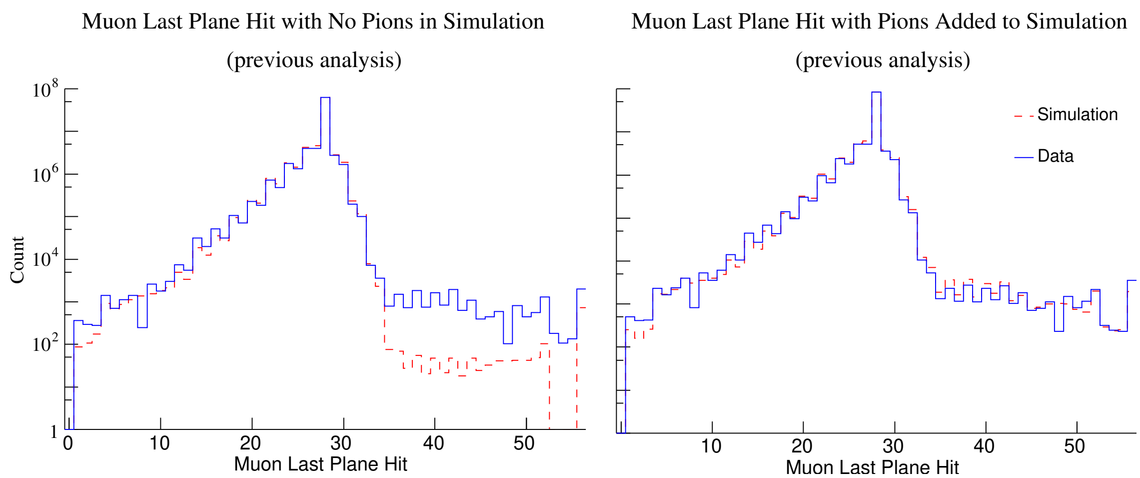

(a) Distribution of last plane that measures the muon, for data and simulation, from the previous analysis (originally Fig. 6.9 from Ref. [59]). The mismatch is resolved by including an additional source of pion decays at the location of the upstream beam package.

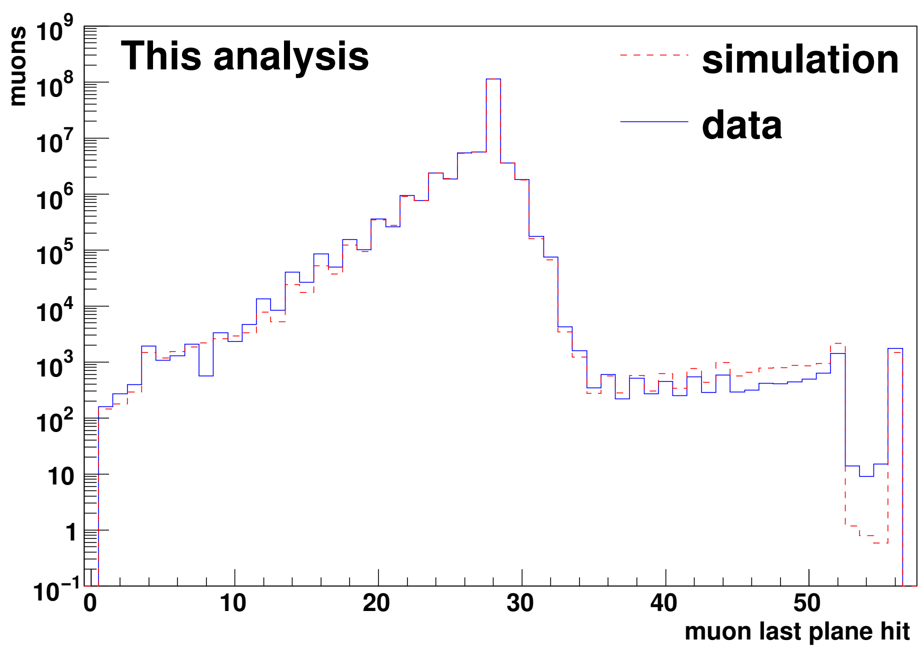

(b) The same figure re-made for the current analysis, where *no* additional pion decays are added.

Figure 6.12: Background muon contamination in the two $P_\mu^\pi \xi$ analyses.

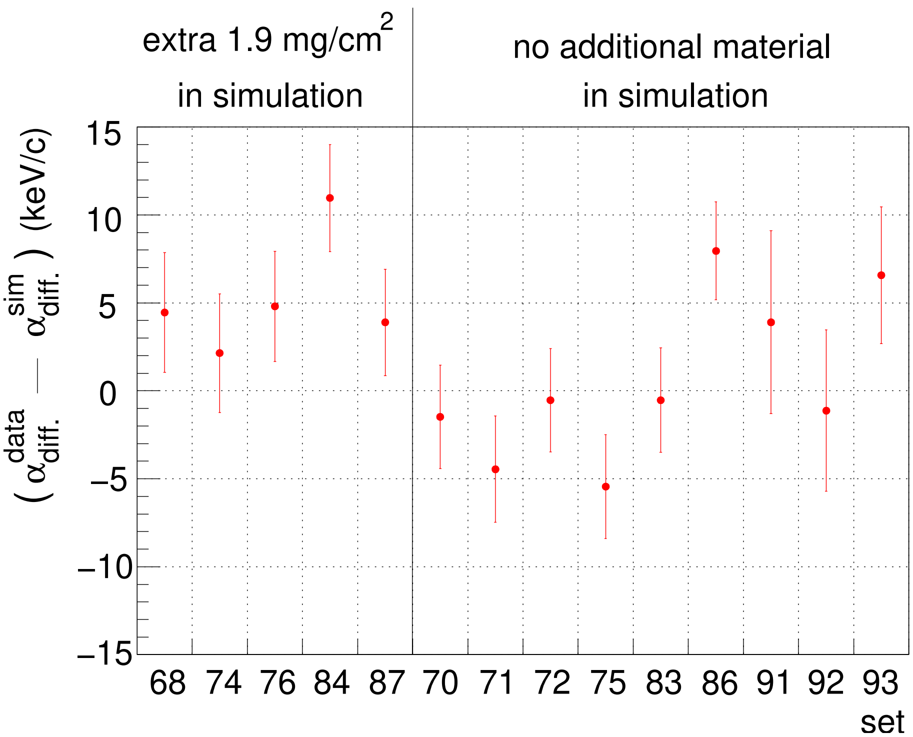

Figure 6.13: The muon stopping distribution is most accurately measured in data and simulation using the endpoint calibration; see the main text for a description of the ordinate. The data and simulation distributions are in agreement when $\alpha_{\text{diff.}}^{\text{data}} - \alpha_{\text{diff.}}^{\text{sim.}} = 0$.

6.3 Chamber response

6.3.1 Drift chamber space-time relationship

The space-time-relationships (STRs) in the drift cells were optimised by minimising the residual between the hit times from the drift cell, and the times that best fit the helix trajectories (see Section 3.2.7). In the simulation, where it was sufficient to use a single STR cell for all wires and planes, this procedure effectively absorbs a small bias from the helix fitting into the STRs. In data, where a separate STR cell was obtained for each plane[44], the procedure corrects for plane-to-plane construction and response differences, in addition to any small bias from the helix fitting algorithm.

The refinement procedure was carried out iteratively, with the STR forced to remain smooth at each step. After convergence, there were residuals in the drift cell corresponding to regions where manipulating the STRs could not bring the drift time closer to the fitted trajectory. The amount by which these residuals differ in data and simulation is the basis of the chamber response systematic uncertainty.

The difference between the data and simulation residuals at the final iteration is shown in Fig. 6.14. The data results are averaged over all planes. Only half a cell is shown, and in practice this is reflected about the line $uv = 0$. The largest discrepancies between data and simulation are at the edge of the cell (u or $v = 0.2\,\text{cm}$), where there were low statistics and the single hit resolution was degraded. In the rest of the cell the variations are at the impressive level of $< 4\,\text{ns}$.

The following approach exaggerated the differences between data and simulation without breaking the smoothness of the STRs:

1. Generate a separate refined STR cell for each plane in the simulation.

2. For each plane, find the difference in residuals at the final iteration, just like Fig. 6.14. Fit this distribution with a fifth order polynomial function.

3. Exaggerate the polynomial function until the χ^2/ndof of positron helix fits becomes a factor of two worse. This corresponds to the single hit resolution being degraded by $\approx (\sqrt{2} - 1) \approx 40\%$, and requires the polynomial function to be exaggerated by a factor of ten.

[44]One STR cell was used for the whole plane. The cell is therefore an average over all the cells within the plane.

4. Re-analyse the simulation with the scaled STRs, and compare to the nominal simulation.

5. Reduce the change in $P_\mu^\pi \xi$ by the scale factor of ten.

The change in $P_\mu^\pi \xi$ from this procedure is 2.0×10^{-4}. However, the exaggeration of errors in the STRs also worsens the reconstruction resolution, for which there is already an orthogonal systematic uncertainty. In order to avoid double-counting with the resolution systematic uncertainty, a subtraction of 1.1×10^{-4} is necessary, leading to a final systematic uncertainty from the STRs of $\pm 0.9 \times 10^{-4}$.

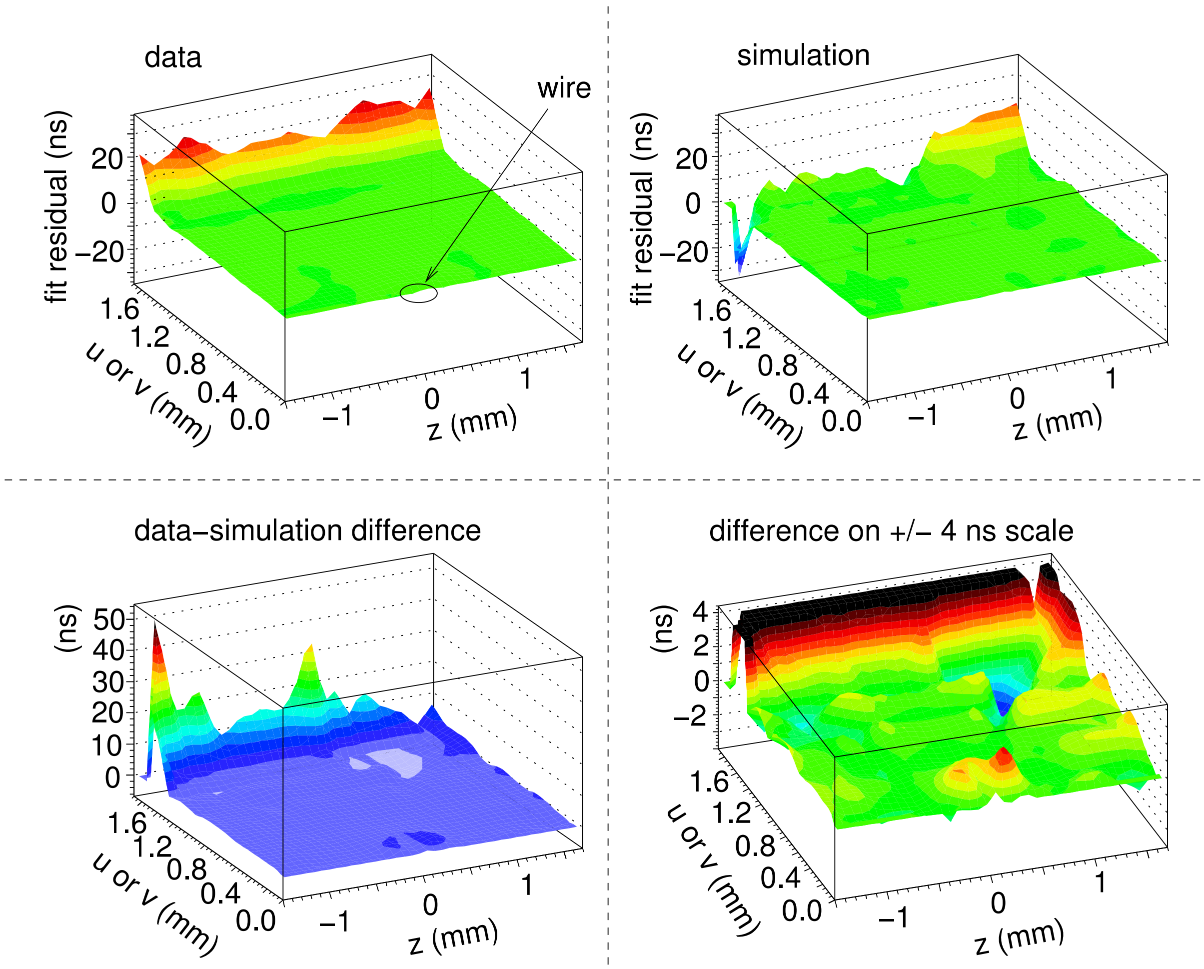

Figure 6.14: An example of the time residuals from the helix fit, after refining the space-time-relationships. Only half a cell is shown, which in practice is reflected about $uv = 0$.

6.3.2 Drift chamber geometric effects

In the simulation, the drift chamber foils are always at the same, well-known positions. In the real chambers there are three uncertainties in the shape and position of the foils: i) the absolute foil positions are only known to better than 100 μm, ii) the foils have an absolute bulge (inwards or outwards) of up to 60 μm, and iii) the bulge in the foils varies by up to 35 μm during data acquisition. A change in the foil position can reduce or increase the number of hits that are available for the helix fitting algorithm, and this introduces a systematic uncertainty.

An upper limit on the systematic uncertainty uses a simulation with the chamber foils moved inwards[45] by 500 μm. This changes $\Delta P_\mu^\pi \xi$ by $(6.7 \pm 5.4) \times 10^{-4}$. Using the largest foil position uncertainty of 100 μm, this corresponds to a systematic uncertainty of $(100/500) \times 6.7 \times 10^{-4} = \pm 1.3 \times 10^{-4}$.

6.3.3 Upstream-downstream efficiency

The simulation must reproduce the difference in track reconstruction efficiency (TRE) between the upstream and downstream halves of the detector. This is measured in data and simulation using a special "upstream stops" analysis, where muons are stopped close to the trigger scintillator, and the decay positrons are reconstructed independently in each half of the detector. The TRE is then calculated based on how often a positron is reconstructed in one half of the detector, but not in the other.

The difference in TRE between data and simulation is shown in Fig. 6.15, where a cut has been placed at $(23 < p < 29)$ MeV/c to remove the area of phase space that is contaminated with a beam positron background. There is evidence that the TRE depends on $|\cos\theta|$, but not on p. The systematic uncertainty is evaluated by multiplying the number of upstream counts by i) a constant to measure the p-contribution, and ii) a linear function to measure the $|\cos\theta|$-contribution. After averaging over the aluminium and silver targets, the p-contribution is 1.3×10^{-4}, and the $|\cos\theta|$-contribution is 0.2×10^{-4}. These are added in quadrature to give a total efficiency systematic uncertainty of $\pm 1.4 \times 10^{-4}$.

[45] The systematic uncertainty is symmetric since the real cells can be thicker or thinner than their nominal positions. The systematic uncertainty is evaluated using simulation, where it is only straightforward to move the foils inward.

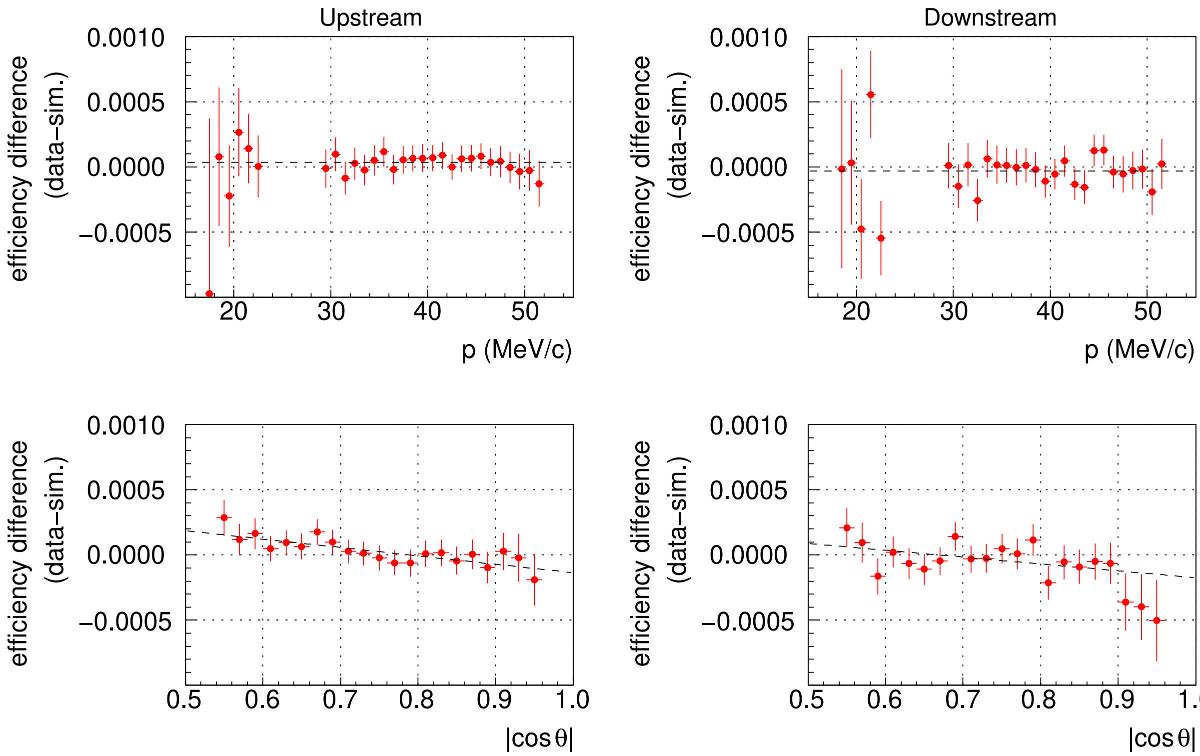

Figure 6.15: Track reconstruction efficiency for the kinematic fiducial. The results for the aluminium target are shown. The upstream efficiency is determined by how often a positron is reconstructed in the downstream half of the detector, but not in the upstream half. The definition is reversed for the downstream efficiency. A cut has been placed at $(23 < p < 29)$ MeV/c in order to remove beam positrons. The dashed lines are simple linear fits that are used to characterise the dependence on angle.

6.3.4 Crosstalk

The origin and removal of electronic crosstalk is described in Section 3.2.2. A limit was set on a systematic uncertainty from crosstalk by re-analysing a data set with the crosstalk removal routine disabled. The change in $\Delta P_\mu^\pi \xi$ was 0.5×10^{-4}, and this is used here as the systematic uncertainty (the scale factor is effectively one).

6.3.5 Wire time offsets

In the previous TWIST $P_\mu^\pi \xi$ analysis[22] the wire time offsets were only calibrated at the beginning and end of the run period. The calibration data were acquired with the magnetic field off, using 120 MeV/c pions and a special downstream trigger. There were significant differences between the two calibrations, and the $P_\mu^\pi \xi$ result changed by 9×10^{-4}

depending on which calibration was used[59].

For this measurement a reliable downstream trigger was in place throughout the acquisition of data. The wire time offsets in each half of the detector were calibrated on a set-by-set basis using the decay positrons (see Section 3.2.1). Beam positrons that passed through the entire detector were then used to determine the relative timing of the upstream and downstream halves of the detector.

There are two systematic uncertainties from this procedure: a bias from the calibration technique, and the degree to which the relative upstream/downstream timing could be established. The systematic uncertainty due to a bias was evaluated as $\pm 0.04 \times 10^{-4}$, and will therefore not be discussed here (see Ref. [81] for further details). The relative timing of the upstream and downstream halves was determined to better than 0.10 ns. An exaggerated upstream-downstream shift of 10 ns changes $\Delta P_\mu^\pi \xi$ by $(82 \pm 7) \times 10^{-4}$. After scaling down by $10\,\text{ns}/0.10\,\text{ns} = 100$, this results in a systematic uncertainty of $\pm 0.8 \times 10^{-4}$.

6.3.6 Negligible chamber response uncertainties

Three additional chamber response systematic uncertainties were considered, but they were not large enough to require recognition in Table 6.1, although they were included in the previous TWIST $P_\mu^\pi \xi$ analysis[22].

When a decay positron crossed the same drift cell as the muon, there was reduced gain for the positron due to lingering muon ionisation (the chamber had a "dead zone"). The simulation included this dead zone by removing hits within 0.06 cm of the muon hit along the wire, during a recovery time of 3.0 μs. These parameters were tuned according to the fraction of hits removed by the dead zone in data. In the simulation this corresponded to 0.03% of hits, tuned to an accuracy of about 20% (*i.e.* $\pm 0.006\%$ of the total number of hits) A special simulation was run with an enhanced dead zone: instead of only removing positron hits within 0.06 cm of the muon hit, the entire wire was made insensitive to decay positrons for the 3.0 μs recovery time. Inclusion of the enhanced dead zone changes $\Delta P_\mu^\pi \xi$ by $(25 \pm 2) \times 10^{-4}$, and removes 0.99% of hits. Compared to the nominal case, this special simulation exaggerates the removed hits by $0.99\% - 0.03\% = 0.96\%$, resulting in a scale factor of $0.96\%/0.006\% = 160$. The change in $\Delta P_\mu^\pi \xi$ is then $(1/160) \times 25 \times 10^{-4}$, which is a systematic uncertainty of just $\pm 0.2 \times 10^{-4}$.

For the previous TWIST $P_\mu^\pi \xi$ analysis[22], the wires in the simulation were placed at exactly the centre of the cathode foils. In reality, they were determined to be offset from the centre by 150 μm, and this introduced a systematic uncertainty of $\pm 2.2 \times 10^{-4}$. The current analysis includes the offset in the simulation, eliminating this uncertainty altogether.

In the simulation the gas density inside the drift chambers is constant, but for the data it tracks the atmospheric pressure and the exterior temperature, which results in drift cell space-time-relationships (STRs) that vary over a data set. The previous analysis used the same STRs for all the data, and a systematic uncertainty of $(0.2 \pm 0.2) \times 10^{-4}$ was established based on the RMS variation in density[85], and $(1.7 \pm 1.0) \times 10^{-4}$ based on the largest variation in density[59]. For the current analysis, each run was analysed with STRs that were corrected for the atmospheric pressure and temperature, meaning the systematic uncertainty is much smaller than $\pm 0.2 \times 10^{-4}$. Therefore the change in STRs due to gas density is a negligible effect, and is no longer evaluated.

6.4 Detector alignment

Two systematic misalignments of the drift chambers have been previously considered[85]: a "shear", where each detector plane is offset in u or v from the previous one by a constant amount, and a "corkscrew", where each detector plane is rotated about the z-axis from the previous one by a constant angle. Stringent limitations from the detector design meant these systematic uncertainties had negligible values of $\pm 0.009 \times 10^{-4}$ and $\pm 0.020 \times 10^{-4}$ for the shear and corkscrew, respectively. These were not re-evaluated for the current measurement, and are simply discarded. Random misalignments are not considered, since these smear the tracking residuals and degrade the resolution, which is handled as an orthogonal systematic uncertainty (see Section 6.6).

The length (z) and width (u, v) scales are used to determine the momentum components of the reconstructed helices. The systematic uncertainties from these scales are re-evaluated for the current measurement. The z length scale is known to 50 μm out of 100 cm, which is a fractional uncertainty of 5.0×10^{-5}. A special analysis made a fractional change of 1×10^{-3} to the z-component of the momentum, which is a factor of 20 larger than the true uncertainty. The change in $P_\mu^\pi \xi$ is $(0.5 \pm 5.2) \times 10^{-4}$, which leads to a systematic uncertainty of $(1/20) \times 0.5 \times 10^{-4} = \pm 0.03 \times 10^{-4}$.

For the width scale, the wires were positioned in a plane to better than 5 μm; a wire plane is 32 cm in width, corresponding to a fractional uncertainty of 2×10^{-5}. The sensitivity of $P_\mu^\pi \xi$ was determined in a similar way to the length scale: a special analysis made a fractional change of 1×10^{-3} in both the $u-$ and $v-$components of the momentum, finding that $P_\mu^\pi \xi$ changes by $(7.8 \pm 5.2) \times 10^{-4}$. This time the systematic uncertainty is scaled down to $(2 \times 10^{-5})/(1 \times 10^{-3}) \times 7.8 \times 10^{-4} = \pm 0.2 \times 10^{-4}$.

6.5 Positron interactions

6.5.1 Bremsstrahlung and δ-electron rates

Uncertainties from the simulation's continuous energy loss model are part of the energy calibration uncertainties in Section 6.7.2. In this section the uncertainties from discrete processes are described; the most important of these are δ-electron production (where an electron is knocked out of an atomic orbital) and Bremsstrahlung ("braking radiation", where one or more photons are radiated during deceleration, significantly changing the energy of the positron). The simulation must accurately reproduce these processes since the extra tracks interfere with the reconstruction. Specifically there are three effects: first, the number of chamber hits will differ in data and simulation; second, the δ-electrons and Bremsstrahlung represent invisible energy loss contributions that cause the positron to be reconstructed with a different energy and angle than the original decay kinematics; third, the processes can "break" the track by introducing a large angle change, and this prevents reconstruction. Approximations in the `GEANT3` physics and uncertainties in material thicknesses can also cause the simulation's rates to differ from data.

The δ-electron rate is compared in data and simulation by selecting events where the decay positron trajectory is broken in two, with an additional electron track originating from the point where the track is broken. The momentum of the electron track is well correlated with the momentum difference between the two broken track halves, indicating that genuine δ-electrons are being measured (see Fig. 6.16). A special simulation with the production of δ-electrons disabled confirmed that the background of the measurement is small.

The momentum distributions of the reconstructed electrons are compared for nominal data and simulation on the left of Fig. 6.17, where the reconstruction efficiency is seen to decrease below 6 MeV/c. Using the range $(6 < p_\delta < 16)$ MeV/c and all the available data sets, the ratio of δ-electrons in data and simulation is 1.007 ± 0.009 (*i.e.* the central value indicates a deficit of δ-electrons in the simulation, but this is not statistically supported).

A special simulation is used with the δ-electron probability increased by a factor of three, changing $P_\mu^\pi \xi$ by $(25 \pm 7) \times 10^{-4}$. For this simulation, the reconstructed δ-electrons

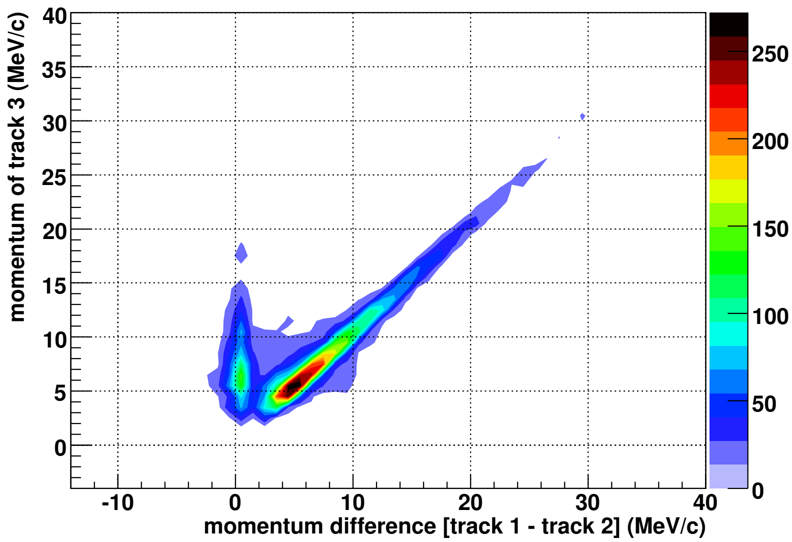

Figure 6.16: Validation of the δ-electron measurement method. Events are selected with a decay positron track that is broken in half, becoming tracks 1 and 2, with an additional negative particle originating from the break point (track 3). The observed correlation indicates that δ-electrons are being measured.

are compared for the nominal and special simulation on the right of Fig. 6.17, where the ratio of counts is 2.80 ± 0.04; the ratio is not exactly 3.0 due to second order effects. The systematic uncertainty is then

$$\frac{1.007 - 1.0}{3.0 - 1.0} \times 25 \times 10^{-4} = \pm 0.1 \times 10^{-4}. \tag{6.16}$$

Note that using 2.80 instead of 3.0 in the denominator does not change the systematic uncertainty by a significant amount.

The Bremsstrahlung rate is compared in data and simulation using a similar approach to the δ-electrons: events are selected with a "broken" decay positron trajectory, and the momentum difference between the two halves of the track is shown on the left side of Fig. 6.18. Another special simulation is used with the Bremsstrahlung probability increased by a factor of three, and this changed $P_\mu^\pi \xi$ by $(55 \pm 7) \times 10^{-4}$; the effect on the broken track momentum difference is shown on the right side of Fig. 6.18. Using the momentum range of $(15 < p < 35)$ MeV/c, the ratio of Bremsstrahlung events in data and simulation averaged over all sets is 1.024 ± 0.004. The ratio between the nominal and increased

Bremsstrahlung rate simulations is 2.82 ± 0.02. The systematic uncertainty is then

$$\frac{1.024 - 1.0}{3.0 - 1.0} \times 55 \times 10^{-4} = \pm 0.7 \times 10^{-4}. \tag{6.17}$$

However, exaggerating the Bremsstrahlung also degrades the reconstruction resolution. To avoid double counting the resolution uncertainty, the contribution from the resolution must be subtracted, which results in a final Bremsstrahlung systematic uncertainty of $\pm 0.5 \times 10^{-4}$.

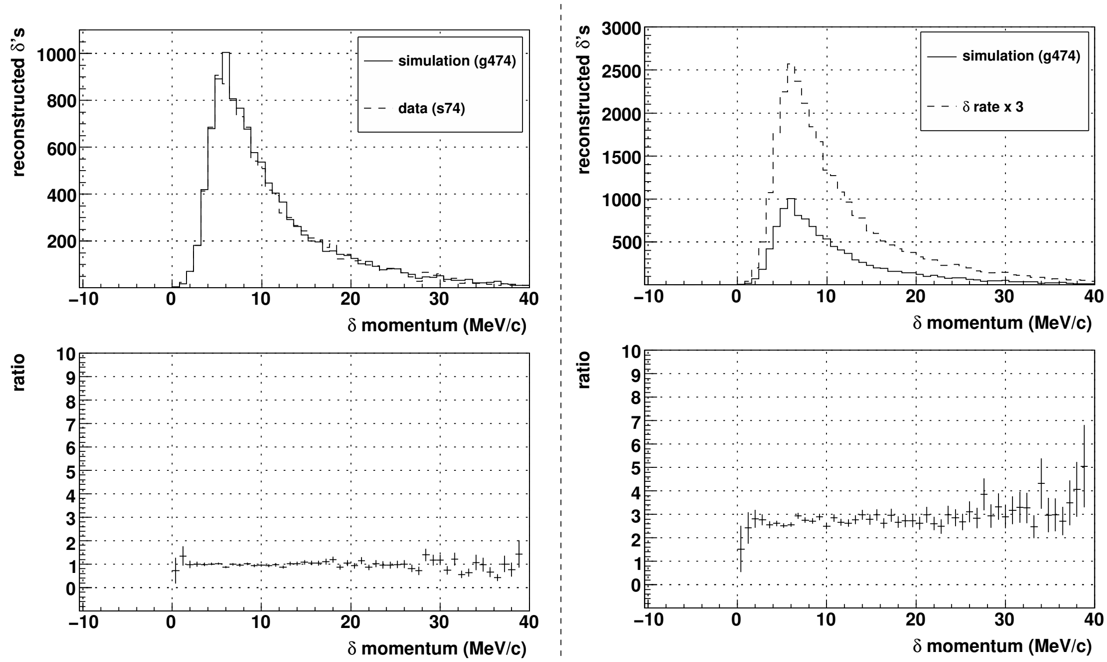

Figure 6.17: The left hand side shows the distribution of reconstructed δ-electrons for data and simulation. The right hand side shows the simulation where the δ production rate was increased by a factor of three. The range $(6 < p_\delta < 16)$ MeV/c is used to evaluate the systematic uncertainty.

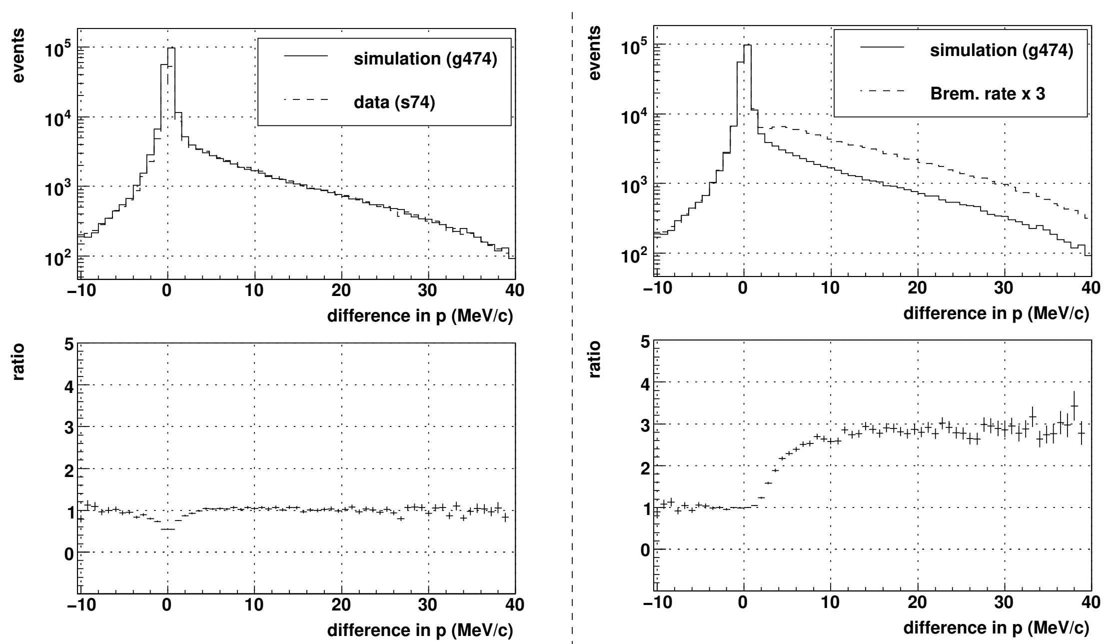

Figure 6.18: The left hand side shows the energy difference for the two halves of a broken track, for data and simulation. The right hand side shows the same distribution for a simulation where the Bremsstrahlung rate was increased by a factor of three. The range $(15 < p < 35)$ MeV/c is used to evaluate the systematic uncertainty.

6.5.2 Outside material

In the *upstream* half of the detector the positrons could be backscattered from the beam pipe and the upstream beam package (see Section 2.9). The door of the steel yoke was also a source of backscatters, but most of the yoke was shielded by the upstream beam package. In the *downstream* half of the detector, during nominal operation there was no corresponding downstream beam package, so positrons could only be backscattered from the steel yoke.

These backscatters cause extra hits that interfere with the decay positron reconstruction. They are well reproduced upstream in the simulation since the beam pipe and most of the upstream beam package (excluding, for example, the light guides) are included. However, the steel of the yoke is not included since the additional showering was expected to significantly increase the computation time. This lack of steel yoke and imperfections in placement and/or thickness of materials are expected to introduce a systematic uncertainty.

The degree to which backscatters match in data and simulation can be compared by selecting the time window containing the decay positron, and then finding the difference in average times between the PCs at the far upstream and downstream ends of the detector. A backscatter will cause extra hits in either the upstream or downstream PCs, resulting in an additional peak in this time distribution. The upper and middle plots in Fig. 6.19 show the time distribution for the windows where the decay positron is upstream and downstream respectively. The simulation shows evidence of a surplus in backscattered upstream decay positrons, and a deficit in backscattered downstream decay positrons that is consistent with the steel yoke being disabled.

The effect on $P_\mu^\pi \xi$ can be estimated using two simulations, with and without the *downstream* beam package in place. This exaggerates the number of downstream decay positrons that are backscattered. The difference in $P_\mu^\pi \xi$ between these simulations is $(2.4 \pm 4.0) \times 10^{-4}$, and the two time distributions are shown in the bottom panel of Fig. 6.19. The change in $P_\mu^\pi \xi$ is then scaled down according to the ratio of differences in counts; specifically, the difference between data and simulation is divided by the difference between the simulations with and without the downstream beam package. After averaging over all sets, this results in scale factors of 7 and 14 for the upstream and downstream backscatters respectively. When added in quadrature the systematic uncertainty is then $\sqrt{(2.4/7)^2 + (2.4/14)^2}$, which is $\pm 0.4 \times 10^{-4}$.

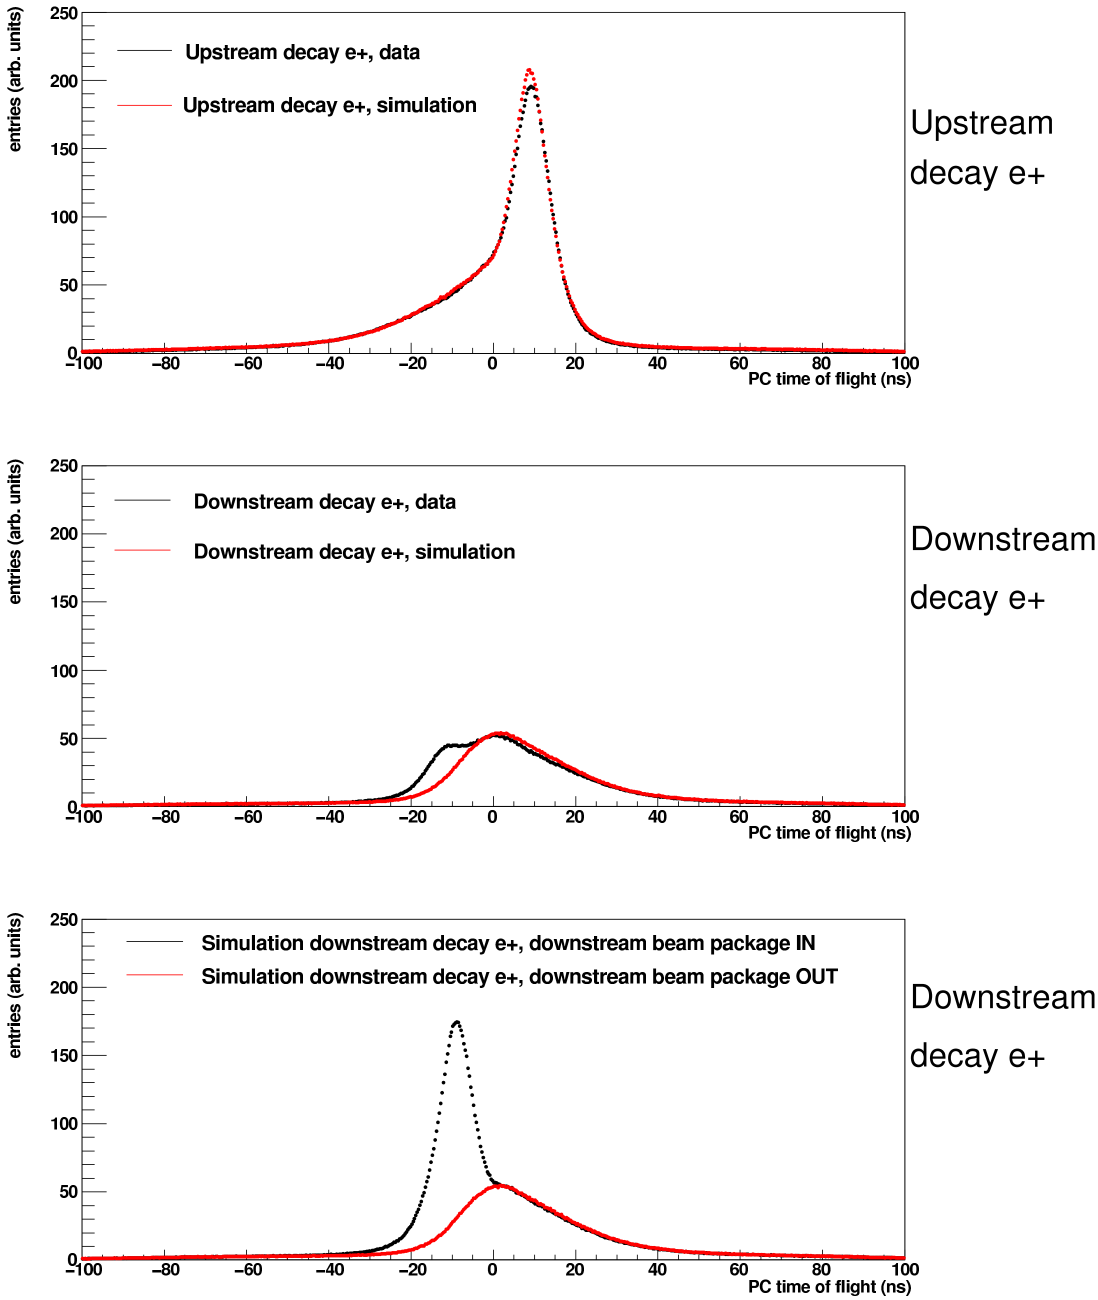

Figure 6.19: The distributions are the difference between the average upstream and downstream PC times. The backscatters from upstream and downstream decay positrons are compared for data and simulation in the upper and middle distributions. The effect of including a downstream beam package in the simulation is shown at the bottom.

6.6 Resolution

The momentum (p) and angle ($\cos\theta$) reconstruction resolutions might differ in data and simulation. This discrepancy is relatively unimportant over the kinematic fiducial since the decay spectrum is smooth. However, the analysis uses the sharp kinematic endpoint to energy calibrate the spectrum, and a significant difference in resolution at this endpoint will introduce a systematic uncertainty.

The resolution can be compared in data and simulation indirectly using special "upstream stops" data, where muons are stopped at the entrance of the detector, and the decay positron is reconstructed independently in each half of the spectrometer. The reconstructed p and θ will differ in each half due to energy loss and multiple scattering through the target module. The distribution of the momentum/angle differences is dominated by a Gaussian resolution function, so that the difference in σ widths between data and simulation is a measure of how well the simulation reproduces the detector's resolution. Even though this difference is also sensitive to discrepancies in the simulation's positron interactions physics, we conservatively assign the full difference as a resolution problem.

The Gaussian widths are shown for a limited p range in Fig. 6.20. The absolute width is seen to depend on $1/\sin\theta$, but the *difference* between data and simulation is well approximated by a constant for both the momentum and angle dependence. A previous analysis[85] found this difference in width had a non-trivial dependence on p and $\cos\theta$; the current analysis is improved due to the use of superior drift cell space-time-relationships (see Section 3.2.7).

For each point in $(p, \cos\theta)$, the difference in p and θ width between data and simulation is constructed according to

$$\Delta\sigma = \begin{cases} \sqrt{\sigma_{\text{data}}^2 - \sigma_{\text{sim}}^2} & , \sigma_{\text{data}} > \sigma_{\text{sim}} \\ -\sqrt{\sigma_{\text{sim}}^2 - \sigma_{\text{data}}^2} & , \sigma_{\text{sim}} > \sigma_{\text{data}} \end{cases} \tag{6.18}$$

The weighted average of this quantity over all $(p, \cos\theta)$ is shown in Table 6.7. The systematic uncertainty is measured by exaggerating the largest differences from the table by a factor of five; specifically an additional smearing of 58 keV/c in momentum and 6 mrad in angle is added to the data. The momentum and angle smearing are carried out separately. The momentum smearing changes $P_\mu^\pi \xi$ by $(7.6 \pm 3.4) \times 10^{-4}$, resulting in a systematic uncertainty of $(7.6/5) \times 10^{-4} = \pm 1.5 \times 10^{-4}$. The angle smearing changes $P_\mu^\pi \xi$ by 0.2×10^{-4}, resulting in a negligible systematic uncertainty after scaling down by five.

The width of the decay spectrum endpoint provides an additional independent measurement of the resolution at $\sin\theta = 1$ (extrapolated). The data-simulation difference in this width is $(0.0 \pm 0.3)\,\mathrm{keV/c}$ for silver and $(0.1 \pm 0.3)\,\mathrm{keV/c}$ for aluminium. Since these are smaller than the differences in Table 6.7, the systematic uncertainty of $\pm 1.5 \times 10^{-4}$ is already sufficient.

Table 6.7: Difference in resolution between data and simulation, as defined in the text.

Target	Difference in σ, defined by Eq. (6.18)	
	Momentum (keV/c)	Angle (mrad)
Aluminium	-6.3	-0.07
Silver	-11.5	1.1

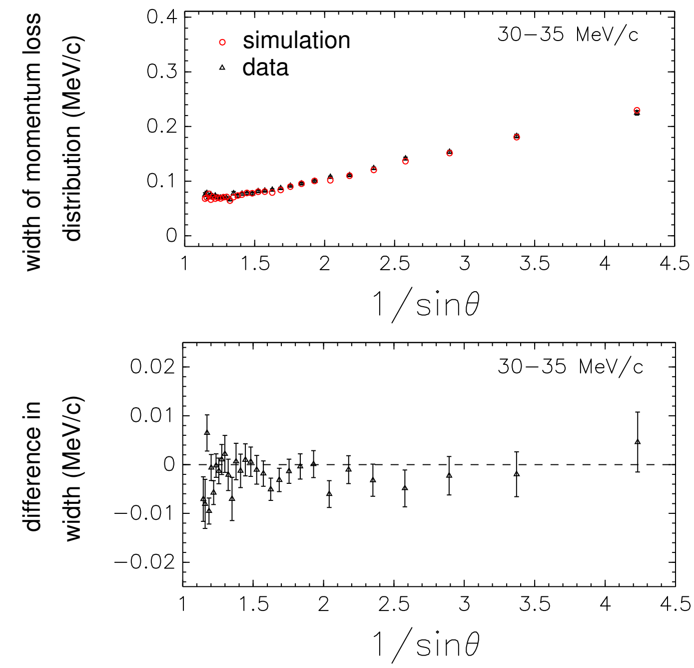

(a) Momentum resolution.

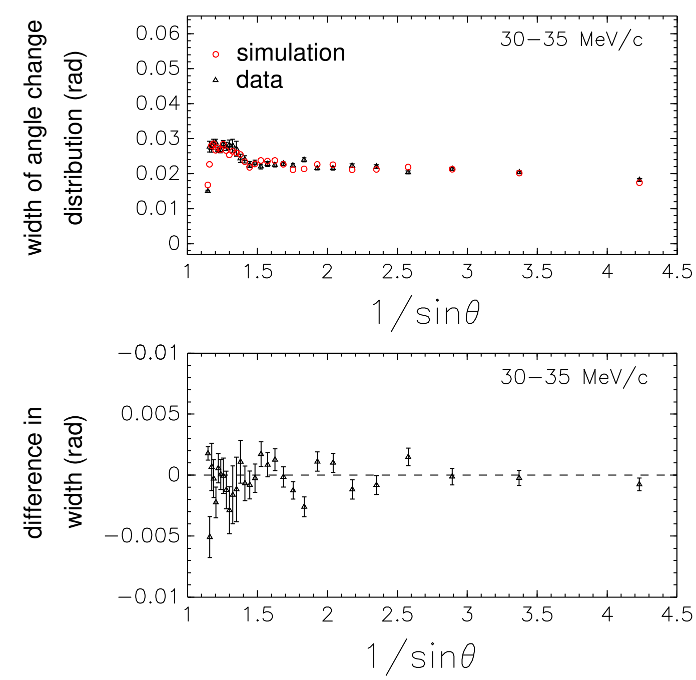

(b) Angle resolution.

Figure 6.20: Comparison of the reconstruction resolution in data and simulation, derived from a special "upstream stops" analysis: muons stop at the entrance of the detector, and are reconstructed independently in each half of the spectrometer; the width of the energy loss and angle change distributions are sensitive to the resolution.

6.7 Momentum calibration

6.7.1 Magnetic field shape

The simulation was self-consistent since it used the same `Opera` field map for generating and reconstructing the positrons. In contrast, the data were analysed with the `Opera` field map, but this had known discrepancies with the mapper measurements within the tracking region, which can introduce a systematic uncertainty.

The difference in field shape between `Opera` and the mapper measurements is well approximated by

$$\delta B_z = C_2 z^2 + C_3 z^3 + C_r r, \tag{6.19}$$

where the optimum C_2, C_3 and C_r values are recorded in Table 6.8. The three previous TWIST analyses have used the same coefficients for the nominal 2.0 T field. The coefficients were re-evaluated for the current measurement by weighting the mapper measurements more carefully[106].

Table 6.8: Coefficients of Eq. (6.19). These relate the `Opera` magnetic field to the results from the mapper measurements.

Parameter	Nominal 2.0 T field		1.96 T	2.04 T
	Previously[a]	**This analysis**		
C_2 ($\times 10^{-8}$ T/cm^2)	-6	-1.7 ± 0.4	$+11.4 \pm 0.3$	-19.7 ± 0.5
C_3 ($\times 10^{-10}$ T/cm^3)	-4	-7.8 ± 0.9	$+2.0 \pm 0.5$	-2.2 ± 0.7
C_r ($\times 10^{-6}$ T/cm)	-12.5	-8.3 ± 0.3	-1.1 ± 0.6	-2.8 ± 0.9

[a] The same parameters were used in Refs. [59, 85, 86, 88].

A new 2.0 T field was produced with the coefficients in Eq. (6.19) exaggerated by a factor of 20. Maxwell's $\nabla \cdot \delta\vec{B} = 0$ equation was satisfied by modifying the radial field components according to

$$\delta B_r = -\left(C_2 r z + \frac{3}{2} C_3 z^2 r \right). \tag{6.20}$$

A data set was then re-analysed with the exaggerated field, and the change in $P_\mu^\pi \xi$ is $(5.1 \pm 7.8) \times 10^{-4}$. After scaling down by a factor of 20, the change in $P_\mu^\pi \xi$ is 0.3×10^{-4}.

This is too small to justify re-analysing all the data with a corrected map, and is instead taken as a systematic uncertainty of $\pm 0.3 \times 10^{-4}$. Previous evaluations of this uncertainty used a smaller scale factor of ten, and did not apply Eq. (6.20), but still found an effect below 1×10^{-4}[59, 85].

(The data sets with a central magnetic field strength of 1.96 T and 2.04 T are not included for $P_\mu^\pi \xi$. Therefore the systematic due to magnetic field shape for these sets is not discussed here. See Refs. [81, 99] for further information.)

6.7.2 Use of the kinematic endpoint

The motivation for an energy calibration and its implementation were described in Section 3.5. In summary, the reconstructed momenta of the data and simulation disagree at the kinematic endpoints by about 10 keV/c. This difference is shown in Fig. 6.21 for the kinematic fiducial. The upstream ($1/\cos\theta < 1$) and downstream ($1/\cos\theta > 1$) points are fit separately with a straight line, yielding two slopes ($a_{\mathrm{up}}, a_{\mathrm{down}}$) and two intercepts ($b_{\mathrm{up}}, b_{\mathrm{down}}$), for a total of four "endpoint parameters". The calibration must be applied to the entire data spectrum by shifting or scaling each reconstructed momentum. Note that the energy calibration procedure is applied to every systematic uncertainty test, which improves the robustness of the $P_\mu^\pi \xi$ measurement.

There are three uncertainties from the energy calibration: i) a statistical part since only a limited region of the spectrum is used to establish the required correction, ii) a systematic part since either a shift or scale must be used to propagate the correction to the rest of the spectrum, and iii) a systematic part due to the model imposed on the endpoint behaviour with angle. The statistical part is propagated separately for each set, and is applied later in Section 7.1. For the systematic part due to a shift or scale, the two extremes considered were a *shift* according to

$$p_{\mathrm{corrected}} = p_{\mathrm{reconstructed}} - \left(b - \frac{a}{|\cos\theta|} \right), \tag{6.21}$$

and a *scale* according to

$$p_{\mathrm{corrected}} = \frac{p_{\mathrm{reconstructed}}}{1 + \frac{1}{W_{e\mu}} \left(b - \frac{a}{|\cos\theta|} \right)}, \tag{6.22}$$

where $W_{e\mu}$ is the maximum kinematic positron momentum (52.83 MeV/c). The central value of $P_\mu^\pi \xi$ is placed half way between the shift and scale extremes, and a systematic uncertainty of half the difference is assigned to cover both possibilities ($\pm 0.5 \times 10^{-4}$).

An additional systematic uncertainty is assigned due to the linear model imposed on the endpoint behaviour with $\cos\theta$ (see Fig. 6.21). The reduced χ^2 for the linear fits average to 1.27. A decrease in the χ^2 to 1.00 corresponds to larger uncertainties for the endpoint parameters, and this ultimately results in a systematic uncertainty of $\pm 1.4 \times 10^{-4}$ for $P_\mu^\pi \xi$.

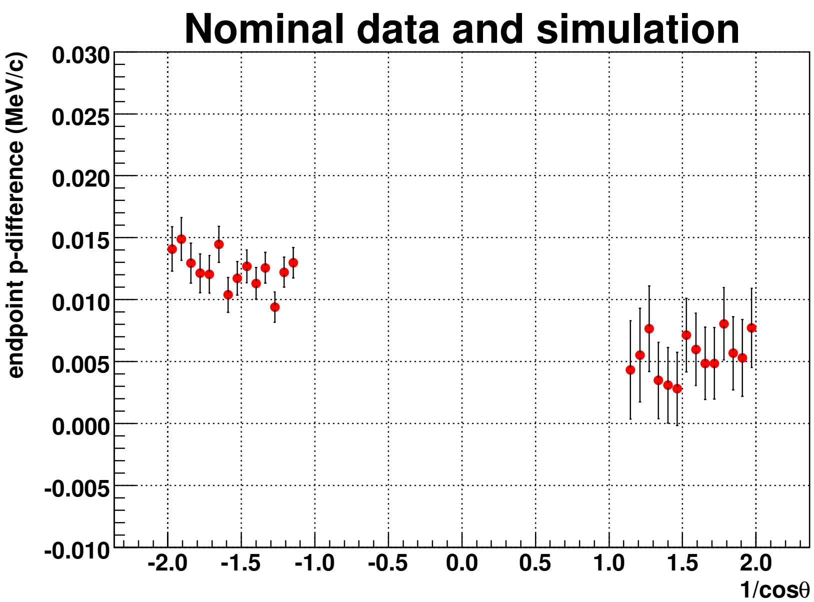

Figure 6.21: Difference between reconstructed momentum in data and simulation at the endpoint of the muon decay spectrum.

6.8 Beam stability

6.8.1 Muon beam intensity

The simulation's muon and beam positron rates were tuned to match the data; a disagreement can change the number of hits in the wire chambers, which affects the performance of the classification. The simulation's muon rate was tuned to match approximately the trigger rate in data. Similarly, the simulation's beam positron rate was tuned to match the probability of a beam positron entering an event.

Increasing the muon rate in the simulation by a factor of 10 changes $\Delta P_\mu^\pi \xi$ by $(10 \pm 7) \times 10^{-4}$. This must be scaled down by repeating the procedure from the previous analysis[85]:

1. For each data set and its accompanying simulation, calculate the following ratio of event types:

$$R_\mu = \frac{(\text{more than one}\, \mu^+)}{(\text{more than one}\, \mu^+) + (\text{one}\, \mu^+, \text{one decay}\, e^+)}. \tag{6.23}$$

   This is a measure of the probability of an event having more than one muon.

2. Calculate the relative ratio of R_μ for data and simulation,

$$\frac{R_\mu\,(\text{sim.}) - R_\mu\,(\text{data})}{R_\mu\,(\text{data})}. \tag{6.24}$$

3. Multiply the relative ratio by the average trigger rate from the data, to estimate the error (in s^{-1}) made in the simulation.

4. Divide the simulation's error by the exaggeration ($24579\,\text{s}^{-1}$).

The calculations from each of these steps appear in Table 6.9. An upper limit on the effect of data-simulation rate discrepancies uses the smallest scale factor (28.7, from set 92), resulting in a systematic uncertainty of $(1/28.7) \times 10 \times 10^{-4} = 0.3 \times 10^{-4}$.

The systematic uncertainty due to a beam positron rate mismatch in data and simulation is just $\pm 0.2 \times 10^{-4}$, and is therefore not discussed further.

Table 6.9: Scale factors for the systematic uncertainty due to beam intensity. R_μ is a measure of the probability of an event with more than one muon. Sets 68-76 were accumulated in 2006 using a silver target. Sets 83-93 were accumulated in 2007 using an aluminium target, with a muon rate that was intentionally higher. Simulation is abbreviated as Sim.

Set	R_μ Data	R_μ Sim.	$\frac{R_\mu\,(\text{sim.}) - R_\mu\,(\text{data})}{R_\mu\,(\text{data})}$	Avg. data trigger (s^{-1})	Sim. error (s^{-1})	Scale factor
68	0.00553	0.00495	-0.10615	2066.0	-219.3	112.1
70	0.00630	0.00541	-0.14129	2324.8	-328.5	74.8
71	0.00660	0.00619	-0.06305	2582.9	-162.8	150.9
72	0.00718	0.00552	-0.23151	2674.2	-619.1	39.7
74	0.00672	0.00617	-0.08089	2592.9	-209.7	117.2
75	0.00739	0.00641	-0.13216	2686.6	-355.1	69.2
76	0.00903	0.00640	-0.29114	2740.9	-798.0	30.8
83	0.01220	0.01029	-0.15728	4221.8	-664.0	37.0
84	0.01291	0.01103	-0.14616	4452.3	-650.7	37.8
86	0.01353	0.01205	-0.10914	4972.3	-542.7	45.3
87	0.01188	0.00977	-0.17768	4024.0	-715.0	34.4
91	0.01218	0.01004	-0.17531	4202.5	-736.7	33.4
92	0.01251	0.01002	-0.19897	4305.3	-856.6	28.7
93	0.01077	0.01003	-0.06851	3809.5	-261.0	94.2

6.9 External

6.9.1 Radiative corrections

The simulation uses the following radiative corrections: full first order, $O(\alpha^2 L^2)$ and $O(\alpha^2 L^1)$ from the second order, and $O(\alpha^3 L^3)$ from the third order. The term $O(\alpha^2 L^0)$ was not used, despite becoming available in 2007 [16]; its effect on the decay parameters will now be shown as negligible.

Over the TWIST kinematic fiducial, the $O(\alpha^2 L^0)$ term has a similar shape to the $O(\alpha^2 L^1)$ term, and the ratio between the terms never exceeds 0.2; this is demonstrated in the paper where the $O(\alpha^2 L^0)$ term is calculated[16]. The $P_\mu^\pi \xi$ sensitivity to excluding the $O(\alpha^2 L^0)$ term was estimated by adding a pure $O(\alpha^2 L^1)$ spectrum to the nominal spectrum. This changes $P_\mu^\pi \xi$ by $(2.9 \pm 0.1) \times 10^{-4}$. The ratio of counts in the nominal and combined spectrum is 1.11. Therefore the effect on $P_\mu^\pi \xi$ of only adding 0.2 of the pure $O(\alpha^2 L^1)$ spectrum is

$$(0.2/1.11) \times 2.9 \times 10^{-4} = 0.5 \times 10^{-4}, \tag{6.25}$$

and $\pm 0.5 \times 10^{-4}$ is used as the systematic uncertainty.

6.9.2 Correlation with η

The muon decay spectrum does not allow a precise measurement of the parameter η. Therefore η was fixed to its world average value[46] of $(-36 \pm 69) \times 10^{-4}$[7], and the correlation between η and $P_\mu^\pi \xi$ is assessed here as a systematic uncertainty. The correlation was found to be $d\xi/d\eta = 0.01528$, so that $\Delta\eta = \pm 69 \times 10^{-4}$ corresponds to a $P_\mu^\pi \xi$ systematic uncertainty of $\pm 1.1 \times 10^{-4}$.

[46]The evaluation of η in Ref. [7] uses a global analysis that includes the TWIST experiment's result for ρ and δ from the 2002 datasets. The evaluation does *not* use the TWIST experiment's more recent ρ, δ and $P_\mu^\pi \xi$ measurements from the 2004 datasets.

Chapter 7

Results

7.1 Blind results

The blind results after all corrections, including set-dependent uncertainties, are shown in Table 7.1. The uncorrected results, the corrections and the set-dependent uncertainties are listed separately in Tables 7.2 and 7.3. The results are determined independently for each of the energy calibration approaches (shift and scale) that were described in Section 6.7.2. The results for the energy calibration applied as a shift are shown in Fig. 7.1, where there is good consistency over the chosen sets. Sets 72 (TECs-in), 76 (muon beam steered in θ_y) and 86 (muon beam steered in x and θ_x) were used to evaluate the systematic uncertainty for the polarisation; they have significantly larger polarisation uncertainties that have not been evaluated, and are therefore not included in the final result. Sets 70 and 71 were taken at different central magnetic field strengths; they are excluded from the final $P_\mu^\pi \xi$ result since the fringe field validation was only carried out at a central field strength of 2.0 T. The weighted average of $\Delta P_\mu^\pi \xi$ is 79.8×10^{-4} when the energy calibration is applied as a shift, and 80.7×10^{-4} for a scale. Averaging over these results, and including an additional statistical uncertainty from determining the relaxation rate (see Section 6.2.10), the final blind result is

$$\Delta P_\mu^\pi \xi = \left[80.3 \pm 3.5\,(\text{stat.})^{+16.5}_{-6.3}\,(\text{syst.})\right] \times 10^{-4}. \tag{7.1}$$

The fit quality for each set is shown in the table of uncorrected results (Table 7.2). For a nominal set, the normalised residuals are shown explicitly in Fig. 7.2. The fit quality is excellent for all sets, and there is no evidence that the normalised residuals depend on momentum or angle.

Table 7.1: Blind results, *after* all corrections, including set-dependent uncertainties.

Set	Target	Description	Shift energy calib. $\Delta P_\mu \xi$ ($\times 10^{-4}$)	Scale energy calib. $\Delta P_\mu \xi$ ($\times 10^{-4}$)
68	Ag	Stopping distrib. peaked $\frac{1}{3}$ into target	90.2 ± 7.6	90.3 ± 7.6
74	Ag	Nominal A	83.7 ± 7.5	83.9 ± 7.5
75	Ag	Nominal B	85.5 ± 6.4	85.9 ± 6.4
83	Al	Downstream beam package in place	80.7 ± 6.6	82.2 ± 6.6
84	Al	Nominal C	69.1 ± 6.9	70.8 ± 6.9
87	Al	Nominal D	82.5 ± 6.7	83.6 ± 6.7
91	Al	Lower momentum I	82.2 ± 13.0	83.0 ± 13.0
92	Al	Lower momentum II	73.5 ± 11.2	75.0 ± 11.2
93	Al	Lower momentum III	62.1 ± 9.2	63.8 ± 9.2
70	Ag	B = 1.96 T	79.0 ± 6.3	80.0 ± 6.3
71	Ag	B = 2.04 T	93.2 ± 6.6	93.3 ± 6.6
72	Ag	TECs-in, nominal beam	90.7 ± 6.4	91.1 ± 6.4
76	Ag	Steered beam A	33.2 ± 7.0	33.6 ± 7.0
86	Al	Steered beam B	52.7 ± 6.2	54.6 ± 6.2

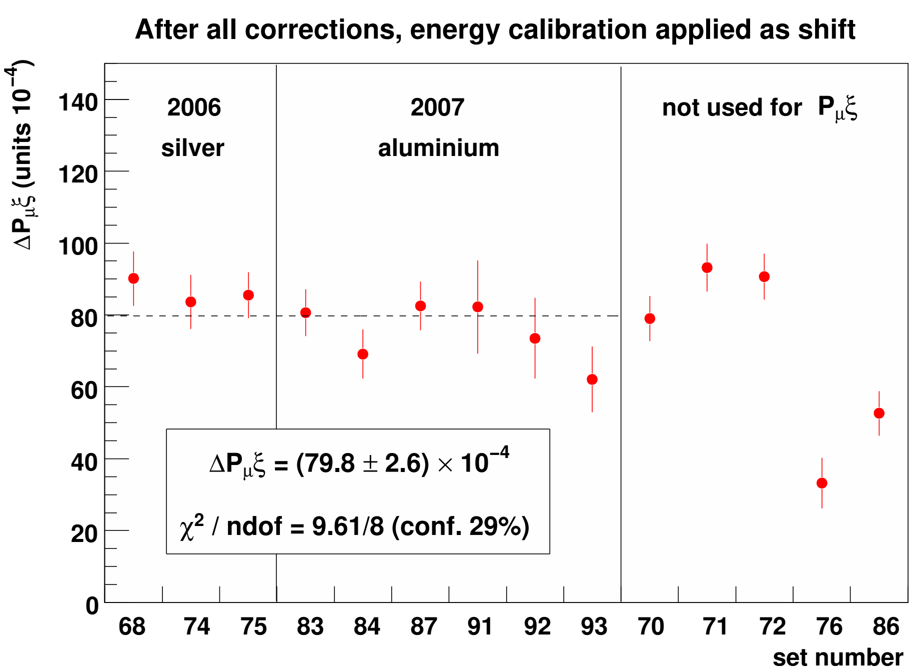

Figure 7.1: Consistency of $\Delta P_\mu^\pi \xi$, the difference in $P_\mu^\pi \xi$ between the data and a hidden value in the simulation.

Table 7.2: Blind results, *before* any corrections.

Set	Energy calibration applied as a shift			Energy calibration applied as a scale		
	χ^2/ndf	Confidence	$\Delta P_\mu\xi$ ($\times 10^{-4}$)	χ^2/ndf	Confidence	$\Delta P_\mu\xi$ ($\times 10^{-4}$)
68	2384/2439	78.3	84.9 ± 7.6	2373/2439	82.7	84.9 ± 7.6
74	2440/2439	49.3	78.3 ± 7.5	2450/2439	43.7	78.5 ± 7.5
75	2453/2439	41.8	80.6 ± 6.4	2452/2439	42.2	80.8 ± 6.4
83	2405/2439	68.5	75.3 ± 6.5	2416/2439	62.5	76.8 ± 6.5
84	2515/2439	13.8	63.6 ± 6.8	2510/2439	15.4	65.2 ± 6.8
87	2411/2439	65.5	77.1 ± 6.7	2407/2439	67.6	78.1 ± 6.7
91	2564/2439	3.9	71.4 ± 12.9	2578/2439	2.5	72.0 ± 12.9
92	2474/2439	30.7	63.3 ± 11.1	2479/2439	28.2	64.7 ± 11.1
93	2504/2439	17.4	52.5 ± 9.0	2518/2439	13.1	54.1 ± 9.0
70	2370/2439	84.0	74.0 ± 6.3	2380/2439	80.0	74.9 ± 6.3
71	2425/2439	57.6	88.1 ± 6.6	2430/2439	54.9	88.1 ± 6.6
72	2513/2439	14.4	85.8 ± 6.4	2508/2439	16.3	86.2 ± 6.4
76	2430/2439	55.2	27.5 ± 7.0	2423/2439	58.9	27.8 ± 7.0
86	2425/2439	57.7	47.7 ± 6.2	2424/2439	58.0	49.6 ± 6.2

Table 7.3: Set-dependent corrections and uncertainties for $\Delta P_\mu^\pi\xi$. These are described in Section 6.1.

Set	Corrections (units 10^{-4})					Set-dependent uncertainties (units 10^{-4})	
	Production target	Relaxation rate	Spectrum fitter	Energy calib. Scale	Energy calib. Shift	Energy calib. statistical	Production target
68	+0.9	+2.7	-0.5	+2.2	+2.3	±0.5	±0.3
74	+0.9	+2.7	-0.5	+2.3	+2.4	±0.5	±0.3
75	+0.9	+2.7	-0.5	+1.9	+2.0	±0.4	±0.3
83	+0.9	+3.3	-0.5	+1.7	+1.8	±0.4	±0.3
84	+0.9	+3.3	-0.5	+1.8	+1.9	±0.4	±0.3
87	+0.9	+3.3	-0.5	+1.8	+1.9	±0.4	±0.3
91	+5.9	+3.3	-0.5	+2.1	+2.2	±0.9	±1.9
92	+5.2	+3.3	-0.5	+2.2	+2.4	±0.7	±1.6
93	+5.2	+3.3	-0.5	+1.7	+1.7	±0.6	±1.6
70	+0.9	+2.7	-0.5	+1.9	+2.0	±0.4	±0.3
71	+0.9	+2.7	-0.5	+2.0	+2.1	±0.4	±0.3
72	+0.9	+2.7	-0.5	+1.8	+1.9	±0.4	±0.3
76	+0.9	+3.3	-0.5	+2.1	+2.2	±0.4	±0.3
86	+0.9	+3.3	-0.5	+1.3	+1.4	±0.4	±0.3

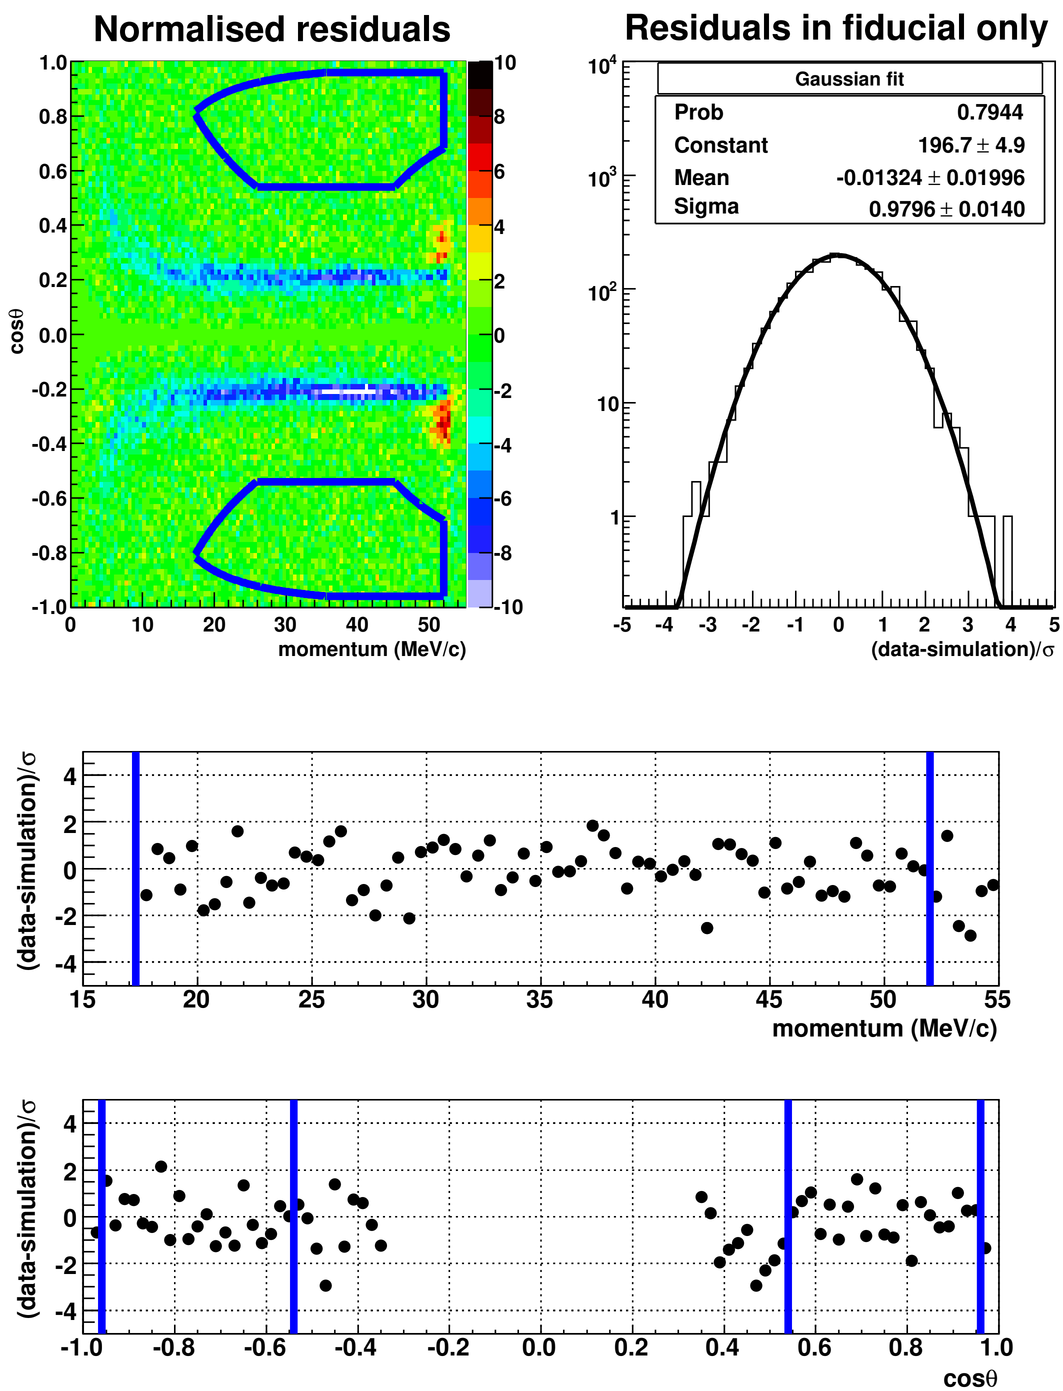

Figure 7.2: Normalised residuals from a spectrum fit of nominal data (set 87) to its accompanying simulation. The thick blue lines indicate the kinematic fiducial boundaries (the region inside the lines is accepted). Within the fiducial region there is no evidence of structure.

7.2 Result revealed

On 29 January 2010, the hidden value of ξ was revealed as

$$\xi_{\text{hidden}} = 0.99281. \tag{7.2}$$

This is added to the blind result in Eq. (7.1) to give

$$P_\mu^\pi\, \xi = 1.00084 \pm 0.00035\,(\text{stat.})^{+0.00165}_{-0.00063}\,(\text{syst.}), \tag{7.3}$$

which agrees with all previous results for $P_\mu^\pi\, \xi$ and $P_\mu^K \xi$, and is consistent with the standard model values of $P_\mu^\pi = \xi = 1$. The new result is a factor of 3.2 more precise than the previous TWIST measurement, $P_\mu^\pi\, \xi = 1.0003 \pm 0.0006\,(\text{stat.}) \pm 0.0038\,(\text{syst.})$[22], and a factor of 7.0 more precise[47] than the pre-TWIST direct measurement, $P_\mu^\pi\, \xi = 1.0027 \pm 0.0079\,(\text{stat.}) \pm 0.0030\,(\text{syst.})$[69]. The new result is also compatible with a recent *indirect* measurement that used TWIST ρ and δ results, $0.99524 < P_\mu \xi \leq \xi < 1.00091$ (90% C.L.)[10]. Note that Eq. (7.3) is the final result of the TWIST $P_\mu^\pi\, \xi$ blind analysis, and this number will be published.

7.3 "White box" consistency test

A simulation was generated using the final result in Eq. (7.3) for ξ. The resulting spectrum was fit against the data. After applying all corrections, the result must come out consistent with zero for the consistency test to be passed. This has been carried out twice (once for the silver target, and once for aluminium), and the results are

$$\Delta P_\mu^\pi \xi\,(\text{Ag}) = (-7.9 \pm 7.5) \times 10^{-4}, \tag{7.4}$$

$$\Delta P_\mu^\pi \xi\,(\text{Al}) = (+7.2 \pm 6.9) \times 10^{-4}. \tag{7.5}$$

These are consistent with zero, which confirms that all corrections have been applied properly.

[47] In units of 10^{-4}, the uncertainty corresponding to $\pm$ one standard deviation for the pre-TWIST direct measurement is $2 \times \sqrt{79^2 + 30^2} = 169$. For the current measurement, the positive error bar is $\sqrt{3.5^2 + 16.5^2} = 16.9$, and the negative error bar is $\sqrt{3.5^2 + 6.3^2} = 7.2$, giving the size of $\pm$ one standard deviation as $16.9 + 7.2 = 24.1$. The improvement factor is then $169/24.1 = 7.0$.

7.4 Physics implications

The physics implications for $P_\mu^\pi \xi$ will now be considered. The simultaneous results for ρ, δ and $P_\mu \xi \delta/\rho$ are beyond the scope of this thesis.

7.4.1 Global analysis of muon decay data

The global analysis of muon decay data was described in Section 1.4.2. At the time of writing this analysis is not yet available in a final form, but a *preliminary* limit is available: the probability of a right-handed muon decaying into a positron of any handedness is reduced from the pre-TWIST value of 0.51% to 0.058% (90% C.L.).

7.4.2 Left-right symmetric models

Several left-right symmetric (LRS) models were described in Section 1.5.2. The most restrictive is the manifest LRS model, where the weak coupling constants and CKM matrices are the same for left- and right-handed particles, and there is no CP violating phase. The generalised LRS model relaxes all of these requirements. An exclusion region for the manifest and generalised LRS models can be constructed at the 90% confidence level. After enforcing the requirement in the LRS models that $P_\mu^\pi \xi \leq 1$, the lower limit at 90% confidence is $\left[P_\mu^\pi \xi\right]_{90} = 0.99922$. This value is used in Eqs. (1.28)-(1.33) to produce Fig. 7.3, which compares the pre-TWIST direct $P_\mu^\pi \xi$ result with the value from the current analysis; there is an additional contour for a measurement of $P_\mu \xi \delta/\rho$. Note that Fig. 7.3 should not be used as the final TWIST exclusion region, since a future publication will include the effect of the new TWIST ρ measurement, and will likely result in a more restrictive plot.

In the generalised LRS model, the current measurement increases the lower limit for the mass of an additional right-handed W-boson from $(g_L/g_R)m_2 \geq 287\,\mathrm{GeV/c^2}$ to $\geq 573\,\mathrm{GeV/c^2}$ (90% C.L.); the only assumption is that right-handed neutrinos are light enough to be produced in muon decay. The mixing angle that relates the weak interaction eigenstates and mass eigenstates (see Eq. (1.26)) is reduced from $|(g_R/g_L)\zeta| \leq 0.065$ to $|(g_R/g_L)\zeta| \leq 0.016$. Note that the Particle Data Group do not list the mass and mixing angle limits under the generalised left-right symmetric model.

In the manifest LRS model, the current measurement increases the lower mass limit in $P_\mu^\pi \xi$ measurements from $m_2 \geq 318\,\mathrm{GeV/c^2}$ to $\geq 634\,\mathrm{GeV/c^2}$ (90% C.L.). As long as $m_{\nu_R} << m_{W_R}$, this can be compared to the limit of $m_{W'} > 1000\,\mathrm{GeV/c^2}$ (95% C.L.)[21]

that originates from high energy W' searches, where the decay $p\bar{p} \rightarrow W'X \rightarrow e\nu X$ is sought. The new result can also be compared to the most restrictive limit on a right-handed W used by the Particle Data Group for averages/fits, $m_{W_R} \geq 715\,\mathrm{GeV/c^2}$ (90% C.L.)[107], which comes from electroweak considerations. In both these cases, the new mass limit is not competitive. The new measurement reduces the manifest mixing angle limit in $P_\mu^\pi \xi$ measurements from $-0.064 < \zeta < +0.053$ to $-0.016 < \zeta < +0.011$. The Particle Data Group do not currently select a value of the mixing angle for averages/fits, which makes a comparison to existing results challenging; we note that the new measurement provides the most restrictive limits from muon decay.

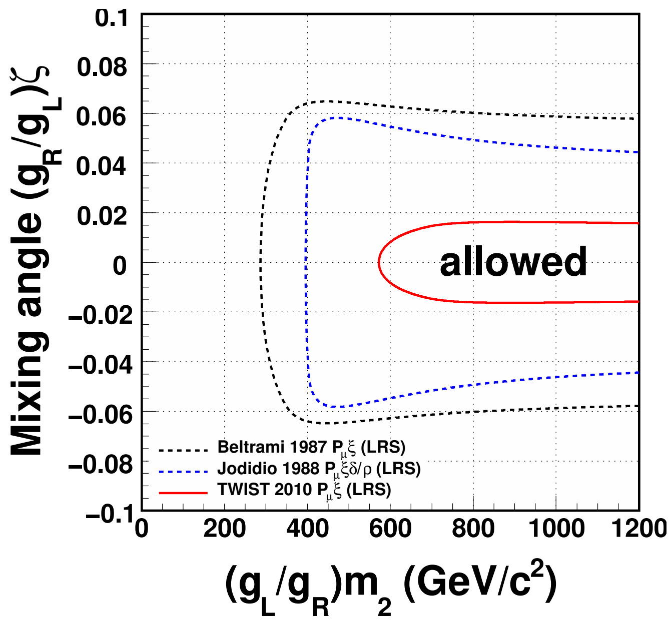

(a) Generalised LRS model.

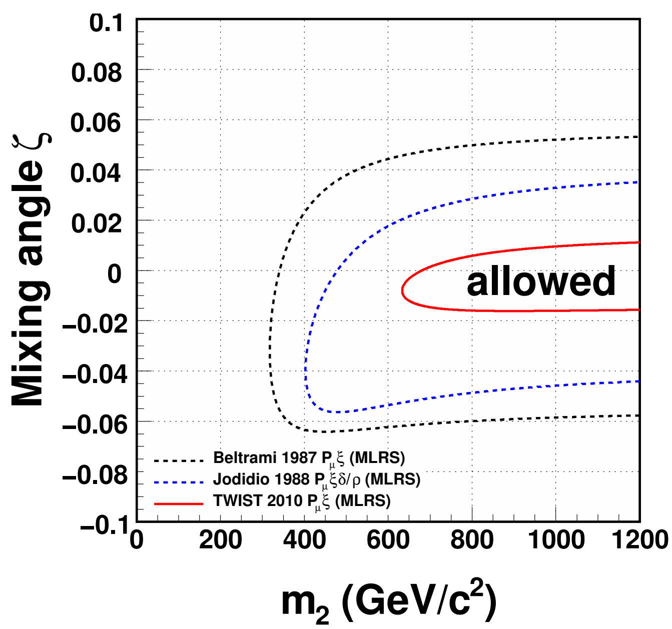

(b) Manifest LRS model.

Figure 7.3: Exclusion region (90% confidence limit) for ζ, the mixing angle between the left- and right-handed W-boson eigenstates, and $(g_L/g_R)m_2$, where g_L and g_R are the weak coupling constants for the predominantly left and right-handed W-bosons, and m_2 is the mass of the predominantly right-handed W-boson. The "Beltrami" entry refers to the pre-TWIST direct $P_\mu^\pi\xi$ measurement[69], and "Jodidio" refers to a measurement of $P_\mu\xi\delta/\rho$[31, 32].

7.5 Future experiments

A $P_\mu^\pi \xi$ measurement with precision greater than 1×10^{-4} is possible with a TWIST-style experiment (*i.e.* longitudinally polarised muons delivered into the centre of uniform magnetic field, with high precision positron tracking). The uncertainties that limited the current measurement will be briefly discussed, with a view to how they could be reduced.

7.5.1 Statistical uncertainty

The statistical uncertainty for this measurement is already at the level of 3.5×10^{-4}, or just 2.5×10^{-4} for the part that originates from the decay spectrum fit. This was achieved with about four months of continuous data acquisition at a surface muon rate of $2000 - 5000\,\mathrm{s}^{-1}$ (a significant amount of tuning over the years 2000 to 2006 was necessary to achieve the required beam quality). The statistical uncertainty of a future measurement could be reduced to commensurate levels by using a channel with an order of magnitude higher flux. For example, a new μE4 channel at PSI already achieves this by placing radiation-hard solenoids close to the muon production target, allowing an acceptance of $\Delta\Omega \sim 135\,\mathrm{msr}$[83] compared to 29 msr from the M13 beam line. A higher rate would require an improvement in the data acquisition electronics, in order to prevent significant pileup. If the TWIST analysis approach were adopted, using an accompanying simulation to include inefficiencies and biases, then the simulation statistics could be significantly increased by taking advantage of faster CPUs.

7.5.2 Magnetic field map uncertainty

The dominant systematic uncertainty from the fringe field could be reduced. We used an MRI magnet surrounded by a custom steel yoke, but a specially constructed magnet could provide a more gradual fringe field by increasing the z-distance over which the field reaches its full strength. Alternatively, or additionally, a higher-rate muon channel would allow the possibility of beam collimation; by selecting low angle muons that undergo very little depolarisation, the uncertainty on that depolarisation would be decreased. However, such collimators could introduce an additional uncertainty from muons scattering off them, and this would have to be carefully assessed; active collimation may help to reduce problems.

The measurement of our magnetic field could have been done better. A future experiment would need alignments of the measuring apparatus under control at the $< 0.5\,\mathrm{mm}$ and $< 1\,\mathrm{mrad}$ level, and should measure all three components of the magnetic field. If the three components are measured with more than one probe, then the relative orientation of the probes must be known with high precision, and a correction may be necessary for the probes not being at exactly the same point in space; a smaller field gradient would also help here. Also, the current experiment would have benefited from field measurements at finer spaced intervals (in all coordinates) over the region that the muons actually traversed.

We used the `Opera` magnetic field simulation to produce the B_x and B_y components of our field map. With all three components measured, it may not be necessary to have a magnetic field simulation at all, although an alternative method of smoothing the field measurements would be necessary. If a simulation is required, it is recommended that more than one piece of software be used; for example, the latest version of `Opera`[87], or the COMSOL Multiphysics (formerly FEMLAB) software[108].

The TWIST approach was to measure the muons before the fringe field, and rely on a `GEANT3` simulation to predict the final polarisation. There are at least two ways to improve the confidence in the final polarisation: first the spin could be transported by one or more independent simulations; second the beam could be steered off-axis in order to lower the polarisation, and a simulation's ability to reproduce the polarisation change from the data would allow confidence to be gained. As seen in this thesis, the alignment of the beam and the field must be under strict control in order for the second approach to work.

The time expansion chambers (TECs) that measured our muon beam had adequate precision, but suffered from alignment uncertainties and aging problems that would be more significant for a future measurement. An improved measurement using a similar device would have to address these issues. A significant uncertainty from the TECs originated from the simulation's ability to correct for the multiple scattering that takes place while the muons pass through the active volume; a subsidiary experiment may be needed to validate the simulation's accuracy in making this correction.

An alternative proposal put forward by a TWIST collaborator is to measure the muon beam inside the strong magnetic field[109, 110]. This would present a greater engineering and analysis challenge, since the device would have to work in a strong magnetic field and the reconstructed trajectories would be helices. If carried out accurately, this approach has the potential to eliminate many of the problems associated with simulating the spin.

7.5.3 Stopping material depolarisation uncertainty

For the current measurement the polarisation's relaxation rate was measured using the normal data. A subsidiary μ^+SR experiment provided a consistent but uncompetitive result. A future experiment should consider an integrated "μ^+SR mode", with a higher beam intensity and a simple analysis that only identifies particles and their times. The goal should be to unequivocally determine the form of $P_\mu(t)$ and its parameters. Since this experimental mode would not measure the absolute polarisation, a Wien filter should be considered to significantly reduce the beam positrons, which would allow a much higher muon rate. (A μ^+SR analysis was considered using the existing TWIST detector. The proportional chambers (PCs) had a timing resolution of $\sim 20\,\mathrm{ns}$, and could identify particles based on their pulse width. This would have allowed us to use decay data below $1\,\mu\mathrm{s}$ to better determine the relaxation rate. However, the suggestion came at a late stage in the analysis and would have required significant software changes to implement.)

If a "μ^+SR mode" is not possible, then a subsidiary μ^+SR experiment should be considered from the outset. Suggestions are made in Section H.9 that would allow a better time differential μ^+SR measurement. Another useful measurement could be provided by a pulsed muon setup such as that of the Rutherford Appleton Laboratory (UK).

For the TWIST polarisation measurements, only aluminium and silver targets were used. Additional targets that produced consistent $P_\mu^\pi \xi$ measurements would strengthen a future result.

We were able to successfully eliminate muons that stopped in the gas before our stopping target; a stricter cut could have further reduced the contamination, with a loss of statistical precision. However, one surprise was our simulation's prediction that 0.11% of muons passed through the metal stopping target and entered PC7, and did not have enough energy to produce a signal. A more careful simulation of the PC response would have allowed us to determine this fraction better, and correct $P_\mu^\pi \xi$ accordingly.

7.5.4 Other uncertainties

The uncertainty from production target scattering can be reduced in three ways: i) by selecting a smaller momentum resolution, which would be feasible with a higher intensity beam line, ii) by a more accurate validation of the multiple scattering within the simulation, iii) if one could select a wide range of sub-surface muon momenta, then muons from much deeper within the production target could be selected; the difference in polarisation between the lower momentum muons and the surface muons would then help to validate the simulation of multiple scattering.

There are theoretical considerations at the $< 1\times10^{-4}$ level that would be important for future measurements. The next level of radiative corrections (full $O(\alpha^3)$) would ideally be evaluated. A calculation of radiative corrections that does not assume an underlying $(V-A)$ interaction would be very welcome, although this would need a suitable renormalisable theory. The pion radiative decay mode should also be considered more carefully; such calculations have been carried out for the purposes of TWIST[111].

The track reconstruction efficiency and resolution were both measured here using a special analysis with the muons stopped at the entrance of the spectrometer, and the decay positron reconstructed separately in each half of the detector. A future experiment should consider designing the beam line to allow a "spread muon tune", where the muons stop close to the detector entrance but are spread out over a much wider area than usual. In addition, the stopping target should be as large as possible to allow a wide range of decay positron phase space to be reconstructed in each detector half. Also the ability to rotate the entire detector (*i.e.* swap the upstream and downstream ends) would provide a more stringent test of measurements that compare the upstream and downstream response of the detector.

The uncertainties from positron interactions (mostly δ-electrons and Bremsstrahlung) will need careful consideration. This may require work by theorists, or a comparison of several simulations that claim to accurately reproduce these processes in the relevant energy range. A future experimenter should consider a subsidiary experiment to help understand these processes better in the low energy range.

Another area requiring thought is the energy calibration. Inevitably a correction or calibration will be needed since the decay positron reconstruction will have subtle biases and systematic errors. The method of measuring and then propagating such a correction will likely be dominant in a future $P_\mu^\pi \xi$ measurement. For the TWIST experiment this correction was due to a complex combination of errors in the magnetic field map, imperfect drift cell space-time-relationships, bias from the helix fitting, the energy-loss model in the simulation, multiple scattering of the decay positron and uncertainties in the stopping distribution; these pieces could not be disentangled, and, as a result, a conservative approach was taken in the propagation of the energy calibration to the bulk of the decay spectrum. A future experiment must consider ways of eliminating these errors, or breaking them into orthogonal pieces; see Ref. [99] for more information.

The remaining uncertainties from Table 6.1 could have easily been reduced. The beam intensity uncertainty could be eliminated by tuning the simulation's muon rate to properly match the data, using the R_μ criteria described in Section 6.8.1. The uncertainty from background muons could be reduced by tuning the stopping distribution based on the $\alpha_{\text{diff.}}$ criteria in Section 6.2.12, and/or adding to the simulation a source of pions at the end of the M13 beam line. The refined space-time-relationships in the DCs and the wire time offsets were already adequate for a measurement at the $< 1 \times 10^{-4}$ level. The strict engineering requirements of the TWIST detector meant that alignment uncertainties were already at a negligible level. The outside material systematic could be eliminated by adding more detail to the geometry of the simulation outside of the active detector region. The η correlation will be reduced for future measurements after a global analysis using this $P_\mu^\pi \xi$ measurement and the simultaneous ρ and δ measurements.

In addition to the goal of extracting $P_\mu^\pi \xi$ (and ρ, δ), a future experiment should consider subsidiary measurements that may even benefit the main experiment. An η measurement from the decay spectrum would provide a validation of the results that use the transverse polarisation of the decay positron, although positron interactions would have to be thoroughly understood since η affects the low momentum end of the spectrum. Some extensions to the standard model postulate additional parameters to describe the decay spectrum; see Ref. [81] for a more detailed discussion. The *negative* muon decay spectrum for each stopping target could be produced using the same analysis software; see Ref. [112] for a measurement using the TWIST apparatus. Lastly, if there was a possibility to switch between muons sourced from pions and kaons, then the resulting $P_\mu^\pi \xi$ and $P_\mu^K \xi$ measurements would provide a more complete test of the standard model.

7.6 Conclusions

The quantity $P_\mu^\pi\,\xi$ has been measured as

$$P_\mu^\pi\,\xi = 1.00084 \pm 0.00035\,(\text{stat.})^{+0.00165}_{-0.00063}\,(\text{syst.}). \tag{7.6}$$

This is the final direct $P_\mu^\pi\,\xi$ measurement from the TWIST collaboration, and is a factor of 7.0 more precise that the pre-TWIST result[69]. As discussed in Section 7.4, this result improves the limits on the mass of an additional right-handed W-boson in generalised left-right symmetric models, in addition to the mixing angle between weak interaction and mass eigenstates. New limits for the mass and mixing angle are also possible under the manifest left-right symmetric model, although only the mixing angle is improved in this more restrictive model. A preliminary global analysis of muon decay data, including the new $P_\mu^\pi\,\xi$ measurement, indicates that the probability of a right-handed muon decaying into a positron of any-handedness is reduced from a pre-TWIST value of 0.51% to 0.058% (90% C.L.).

Bibliography

[1] S. Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19(21):1264–1266, Nov 1967.

[2] Donald Hill Perkins. *Introduction to High-Energy Physics; 4th ed.* Cambridge Univ. Press, Cambridge, 2000.

[3] C. Amsler *et al.* Review of Particle Physics. *Phys. Lett. B*, 667(1-5):1–6, 2008. Review of Particle Physics.

[4] D. Griffiths. *Introduction to Elementary Particles.* Wiley, 1987.

[5] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of Neutrinos. *Phys. Rev.*, 109:1015 – 1017, 1958.

[6] D. Binosi and L. Theuβl (sic.). JaxoDraw: A graphical user interface for drawing Feynman diagrams. *Comput. Phys. Commun.*, 161:76–86, August 2004.

[7] C.A. Gagliardi, R.E. Tribble and N.J. Williams. Global analysis of muon decay measurements. *Phys. Rev. D*, 72(7):073002, Oct 2005.

[8] W. Fetscher, H.-J. Gerber and K.F. Johnson. Muon decay: complete determination of the interaction and comparison with the standard model. *Phys. Lett. B*, 173(1):102–106, 1986.

[9] S. Eidelman *et al.* Review of Particle Physics. *Phys. Lett. B*, 592(1-4):1–5, 2004. Review of Particle Physics.

[10] R.P. MacDonald *et al.* Precision measurement of the muon decay parameters ρ and δ. *Phys. Rev. D*, 78(3):032010, 2008.

[11] A. Sirlin, New York University, USA. Private communication, 2007.

[12] A. B. Arbuzov. First-order radiative corrections to polarized muon decay spectrum. *Phys. Lett. B*, 524(1-2):99 – 106, 2002.

[13] A. Arbuzov, A. Czarnecki and A. Gaponenko. Muon decay spectrum: Leading logarithmic approximation. *Phys. Rev. D*, 65:113006, 2002.

[14] A. Arbuzov and K. Melnikov. $\mathcal{O}(\alpha^2 \ln(m_\mu/m_e))$ corrections to electron energy spectrum in muon decay. *Phys. Rev. D*, 66:093003, 2002.

[15] A. Arbuzov. Higher order QED corrections to muon decay spectrum. *J. High Energy Phys.*, 2003(03):063–063, 2003.

[16] K. Melnikov C. Anastasiou and F. Petriello. The electron energy spectrum in muon decay through $\mathcal{O}(\alpha^2)$. *J. High Energy Phys.*, (09):014, 2007.

[17] J. Ng, TRIUMF, Vancouver, Canada. Private communication, 2010.

[18] N. Danneberg *et al.* Muon Decay: Measurement of the Transverse Polarization of the Decay Positrons and its Implications for the Fermi Coupling Constant and Time Reversal Invariance. *Phys. Rev. Lett.*, 94:021802, 2005.

[19] C. Gagliardi, Texas A&M University, Texas, USA. Private communication, 2008.

[20] P. Herczeg. On muon decay in left-right-symmetric electroweak models. *Phys. Rev. D*, 34(11):3449–3456, Dec 1986.

[21] V.M. Abazon *et al.* Search for W' Bosons Decaying to an Electron and a Neutrino with the D0 Detector. *Phys. Rev. Lett.*, 100:031804, 2008.

[22] B. Jamieson *et al.* Measurement of $P_\mu \xi$ in polarized muon decay. *Phys. Rev. D*, 74(7):072007, 2006.

[23] J.D. Jackson. *Classical Electrodynamics.* John Wiley & Sons, 1999.

[24] L.H. Thomas. The kinematics of an electron with an axis. *Phil. Mag. Series 7*, 3(13):1–21, 1927.

[25] V. Bargmann, L. Michel and V.L. Telegdi. Precession of the Polarization of Particles Moving in a Homogeneous Electromagnetic Field. *Phys. Rev. Lett.*, 2:435–436, 1959.

[26] A. Balakin, V. Kurbanova and W. Zimdahl. Precession of a particle with anomalous magnetic moment in electromagnetic and gravitational pp-wave fields. *Gravity Cosmology Supplement*, 82:6–9, 2002.

[27] P. Depommier. The BMT equation. Presentation to TWIST collaboration, May 2006.

[28] P. Depommier. Muon depolarization in multiple scattering (TN100). Technical report, TWIST collaboration, TRIUMF, 2005.

[29] W.H. Koppenol. Names for muonium and hydrogen atoms and their ions(IUPAC Recommendations 2001). *Pure Appl. Chem.*, 73:377–379, 2001.

[30] M. Senba. Muon spin depolarization in noble gases during slowing down in a longitudinal magnetic field. *J. Phys. B: At. Mol. Opt. Phys.*, 31:5233–5260, 1998.

[31] A. Jodidio *et al.* Search for right-handed currents in muon decay. *Phys. Rev. D*, 34(7), 1986.

[32] A. Jodidio *et al.* Erratum: Search for right-handed currents in muon decay. *Phys. Rev. D*, 37(1), 1988.

[33] J. H. Brewer. Muon spin rotation/relaxation/resonance. In *Encyclopedia of Applied Physics 11*, pages 23–53. 1994.

[34] S.F.J. Cox. Implanted muon studies in condensed matter science. *J. Phys. C: Solid State Phys*, 20:3187–3319, 1987.

[35] O. Hartmann *et al.* Diffusion of positive muons in some cubic metals. *Phys. Lett. A*, 61(2):141 – 142, 1977.

[36] J. Brewer, University of British Columbia, Canada. Private communication, 2009.

[37] W.B. Gauster *et al.* Measurement of the depolarization rate of positive muons in copper and aluminum. *Solid State Commun.*, 24(9):619–622, 1977.

[38] W. Schilling. The physics of radiation damage in metals. *Hyperfine Interact.*, 4:636–644, 1978.

[39] D.K. Brice. Lattice atom displacements produced near the end of implanted μ^+ tracks. *Phys. Lett. A*, 66:53–56, 1978.

[40] P. Dalmas de Réotier and A. Yaouanc. Muon spin rotation and relaxation in magnetic materials. *J.Phys: Condens. Matter*, 9:9113–9166, 1997.

[41] P. Dalmas de Réotier and A. Yaouanc. Quantum calculation of the muon depolarization function: effect of spin dynamics in nuclear dipole systems. *J.Phys: Condens. Matter*, 4:4533–4556, 1992.

[42] A. Abragam. *Principles of Nuclear Magnetism.* International series of monographs on physics. Oxford University Press, 1986.

[43] R.S. Hayano *et al.* Zero- and low-field spin relaxation studied by positive muons. *Phys. Rev. B*, 20:850, 1979.

[44] W.J. Kossler *et al.* Diffusion and Trapping of Positive Muons in Al:Cu Alloys and in Deformed Al. *Phys. Rev. Lett.*, 41:1558–1561, 1978.

[45] O. Hartmann *et al.* Coherent Propagation and Strain-Induced Localization of Muons in Al. *Phys. Rev. Lett.*, 41:1055–1058, 1978.

[46] K.W. Kehr *et al.* Muon diffusion and trapping in aluminum and dilute aluminum alloys: Experiments and comparison with small-polaron theory. *Phys. Rev. B*, 26:567–589, 1982.

[47] K.W. Kehr. Empirical information on quantum diffusion. *Hyperfine Interact.*, 17-19:63–74, 1984.

[48] O. Hartmann. Diffusion and trapping of muons in aluminum: New experiments and comparison with Kondo theory. *Phys. Rev. B*, 37:4425–4440, 1988.

[49] O. Hartmann. New results on diffusion in fcc metals. *Hyperfine Interact.*, 64:641–648, 1990.

[50] O. Hartmann *et al.* Studies of μ^+ Localization in Cu, Al, and Al alloys in the Temperature Interval $0.03 - 100\,\mathrm{K}$. *Phys. Rev. Lett.*, 44:337–340, 1980.

[51] D.P. Stoker *et al.* Search for Right-Handed Currents by Means of Muon Spin Rotation. *Phys. Rev. Lett.*, 54:1887–1890, 1985.

[52] D.P. Stoker, University of California, Irvine, USA. Private communication, 2009.

[53] J. Korringa. Nuclear magnetic relaxation and resonnance (sic.) line shift in metals. *Physica*, 7-8:601–610, 1950.

[54] S.J. Blundell and S.F.J. Cox. Longitudinal muon spin relaxation in metals and semimetals and the Korringa law. *J. Phys.: Condens. Matter*, 13:2163–2168, 2001.

[55] J.H. Brewer *et al. Positive muons and muonium in matter*. Muon Physics. Academic Press Inc (London) Limited, 1975.

[56] S.F.J. Cox *et al.* Muon Korringa relaxation. *Physica B*, 289-290:594–597, 2000.

[57] R.H. Heffner. Muon spin depolarization in nonmagnetic metals doped with paramagnetic impurities. *Hyperfine Interact.*, 8:655–662, 1981.

[58] J.A. Brown *et al.* Muon Depolarization by Paramagnetic Impurities in Nonmagnetic Metals. *Phys. Rev. Lett.*, 47:261–264, 1981.

[59] B. Jamieson. *Measurement of the muon decay asymmetry parameter with the TWIST spectrometer*. PhD thesis, University of British Columbia, 2006.

[60] T.D. Lee and C.N. Yang. Question of Parity Conservation in Weak Interactions. *Phys. Rev.*, 104(1):254–258, Oct 1956.

[61] ARGUS collaboration. Determination of the Michel Parameters ρ, ξ, and δ in τ-Lepton Decays with $\tau \rightarrow \rho\nu$ Tags. *Phys. Lett. B*, 431:179–187, 1998.

[62] L3 Collaboration. Measurement of the Michel parameters and the average tau-neutrino helicity from tau decays at LEP. *Phys. Lett. B*, 438:405–416, 1998.

[63] R. Bartoldus. Measurements of the Michel Parameters in Leptonic Tau Decays using the OPAL Detector at LEP. *Nucl. Phys. B (Proc. Suppl.)*, 76:147–157, 1999.

[64] R.L. Garwin, L.M. Lederman and M. Weinrich. Observations of the Failure of Conservation of Parity and Charge Conjugation in Meson Decays: the Magnetic Moment of the Free Muon. *Phys. Rev.*, 105:1415, January 1957.

[65] G. R. Lynch, J. Orear, and S. Rosendorff. Muon decay in nuclear emulsion at 25 000 gauss. *Phys. Rev.*, 120(6):2277, Dec 1960.

[66] M. Bardon, D. Berley, and L. M. Lederman. Asymmetry parameter in muon decay. *Phys. Rev. Lett.*, 2(2):56–57, Jan 1959.

[67] Richard J. Plano. Momentum and asymmetry spectrum of μ-meson decay. *Phys. Rev.*, 119(4):1400–1408, Aug 1960.

[68] Akhmanov *et al.* Asymmetry of Positrons from the Decay $\pi^+ \to \mu^+ \to e^+$ in the Longitudinal 140 000 Oe Magnetic Field. *Sov. J. Nucl. Phys*, 6:230, 1968.

[69] I. Beltrami *et al.* Muon decay: Measurement of the integral asymmetry parameter. *Phys. Lett. B*, 194(2):326–330, 1987.

[70] Ali-Zade *et al.* Asymmetry of angular distribution of $\mu^+ \to e^+$ decay electrons in a 27,000 gauss magnetic field. *Soviet Phys.-JETP*, 36 (9):940, 1959.

[71] R.S. Hayano *et al.* Search for Right-Handed Currents in the Decay $K^+ \to \mu^+\nu$. *Phys. Rev. Lett.*, 52(5):329–332, Jan 1984.

[72] J. Imazato *et al.* Search for right-handed currents in the decay chain $K^+ \to \mu^+\nu$, $\mu^+ \to e + \nu\bar{\nu}$. *Phys. Rev. Lett.*, 69(6):877–880, Aug 1992.

[73] T. Yamanaka *et al.* Search for right-handed currents in the decay $K^+ \to \mu^+\nu$. *Phys. Rev. D*, 34(1):85–96, Jul 1986.

[74] C.A. Coombes *et al.* Polarization of μ^+ Mesons from the Decay of K^+ Mesons. *Phys. Rev.*, 108(5):1348–1351, Dec 1957.

[75] D. Cutts, T. Elioff, and R. Stiening. Muon Polarization and Energy Spectrum in $K^+ \to \pi^0 + \mu^+ + \nu$. *Phys. Rev.*, 138(4B):B969–B979, May 1965.

[76] D. Cutts, R. Stiening, C. Wiegand, and M. Deutsch. Measurement of the Muon Polarization Vector in $K^+ \to \pi^0 + \mu^+ + \nu$. *Phys. Rev.*, 184(5):1380–1392, Aug 1969.

[77] R.S. Henderson *et al.* Precision planar drift chambers and cradle for the TWIST muon decay spectrometer. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 548:306–335, August 2005.

[78] J. Hu *et al.* Time expansion chamber system for characterization of TWIST low-energy muon beams. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 566:563–574, October 2006.

[79] C.J. Oram, J.B. Warren and G.M. Marshall. Commissioning of a new low energy $\pi - \mu$ channel at TRIUMF. *Nucl. Instrum. Methods*, 179:95–103, January 1981.

[80] J. Doornbos. The tuning of M13 for TWIST, theory and some results. Internal TWIST bulletin boards, December 2000.

[81] A. Hillairet. *Measurement of δ (in preparation) (TWIST).* PhD thesis, University of Victoria, 2010.

[82] R.E. Mischke, TRIUMF, Vancouver, Canada. Private communication, 2009.

[83] T. Prokscha et al. The new μE4 beam at PSI: A hybrid-type large acceptance channel for the generation of a high intensity surface-muon beam. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 595.

[84] J.E. Draper. Beam Steering with Quadrupole and with Rectangular Box Magnets. *Rev. Sci. Instrum.*, 37:1390, October 1966.

[85] R.P. MacDonald. *A Precision Measurement of the Muon Decay Parameters ρ and δ (TWIST).* PhD thesis, University of Alberta, 2008.

[86] J. Musser. *Measurement of the Michel parameter ρ in muon decay (TWIST).* PhD thesis, Texas A&M University, 2005.

[87] Vector Fields. OPERA 3D simulation software. http://vectorfields.com.

[88] A. Gaponenko. *A Precision Measurement of the Muon Decay Parameter δ (TWIST).* PhD thesis, University of Alberta, 2005.

[89] D. Wright. TN1: DME / He WC Gas Comparison. Technical report, TWIST collaboration, TRIUMF, October 1996.

[90] R. Openshaw, TRIUMF, Vancouver, Canada. Private communication, 2009.

[91] LeCroy Corporation. 1877 Multihit Time-to-Digital Converter: Specification. `http://www.lecroy.com/lrs/dsheets/1877.htm`.

[92] LeCroy Corporation. LeCroy 1877 Fastbus TDC Manual.

[93] R. Brun and F. Rademakers. ROOT - An Object Oriented Data Analysis Framework. *Nucl. Inst. & Meth. in Phys. Res. A*, 389:81–86, 1997.

[94] CERN Application Software Group. `GEANT` Detector Description and Simulation Tool, Version 3.21, 1994.

[95] F. James. Fitting tracks in wire chambers using the Chebyshev norm instead of least squares. *Nucl. Instrum. Methods Phys. Res.*, 211(1):145 – 152, 1983.

[96] R. Veenhof. Garfield - simulation of gaseous detectors. CERN Program Library W5050.

[97] P. Depommier. TN49: Muon polarisation in electromagnetic fields. Technical report, TWIST collaboration, TRIUMF, 2001.

[98] P. Gumplinger, TRIUMF, Vancouver, Canada. Private communication, 2009.

[99] R. Bayes. *Measurement of ρ (in preparation) (TWIST).* PhD thesis, University of Victoria, 2010.

[100] G.M. Marshall, TRIUMF, Vancouver, Canada. Private communication, 2009.

[101] R. Armenta. TN68: TWIST Magnet: Field Map Study. Technical report, TWIST collaboration, TRIUMF, 2002.

[102] M. Simard. Study of the Magnetic Field Map for TWIST at TRIUMF. Technical report, TWIST collaboration, TRIUMF, 2003.

[103] C. Boucher-Veronneau. The TWIST Solenoid Magnetic Field. Technical report, TWIST collaboration, TRIUMF, 2004.

[104] S. Fotooshi. Modeling the TWIST Magnet. Technical report, TWIST collaboration, TRIUMF, 2005.

[105] A. Grossheim, TRIUMF, Vancouver, Canada. Private communication, 2009.

[106] G. Marshall. OPERA map corrections from Hall and NMR maps at 2.0 T. Internal TWIST report, April 2009.

[107] M. Czakon, J. Gluza and M. Zralek. Low energy physics and left-right symmetry: bounds on the model parameters. *Phys. Lett. B*, 458(2-3):355–360, 1999.

[108] COMSOL. COMSOL Multiphysics. http://www.comsol.com.

[109] V. Selivanov. TN65: GARFIELD Simulation of the TEC. Technical report, TWIST collaboration, TRIUMF, 2002.

[110] V. Selivanov. TN65.4: One dimension TEC at $B = 2\,\mathrm{T}$. Technical report, TWIST collaboration, TRIUMF, 2004.

[111] P. Depommier. TN96: Muon polarization in the $\pi^+ \to \mu^+\nu\gamma$ decay. Technical report, TWIST collaboration, TRIUMF, 2005.

[112] A. Grossheim *et al.* Decay of negative muons bound in ^{27}Al. *Phys. Rev. D*, 80(5):052012, Sep 2009.

[113] H. Yukawa. On the Interaction of Elementary Particles. *I, Proc. Phys.-Math. Soc. Jpn.*, 17:48, 1935.

[114] C.D. Anderson and S.H. Neddermeyer. Cloud Chamber Observations of Cosmic Rays at 4300 Meters Elevation and Near Sea-Level. *Phys. Rev.*, 50:263–271, 1936.

[115] P. Kunze. Untersuchung der Ultrastrahlung in der Wilsonkammer. *Z. Phys. A*, 83(1):1–18, 1933.

[116] S.H. Neddermeyer and C.D. Anderson. Cosmic-Ray Particles of Intermediate Mass. *Phys. Rev.*, 54:88–89, 1938.

[117] S.H. Neddermeyer and C.D. Anderson. Nature of Cosmic-Ray Particles. *Rev. Mod. Phys.*, 11:191–207, 1939.

[118] M. Riordan. *The Hunting of the Quark*. Sage Publications, 1987.

[119] M. Conversi, E. Pancini and O. Piccioni. On the Disintegration of Negative Mesons. *Phys. Rev*, 71:209–210, 1947.

[120] E. Fermi, E. Teller and V. Weisskopf. The Decay of Negative Mesotrons in Matter. *Phys. Rev.*, 71:314–315, 1947.

[121] C.M.G. Lattes, G.P.S. Occhialini and C.F. Powell. Observations on the Tracks of Slow Mesons in Photographic Emulsions Part 1. *Nature*, 160:453–456, 1947.

[122] C.M.G. Lattes, G.P.S. Occhialini and C.F. Powell. Observations on the Tracks of Slow Mesons in Photographic Emulsions Part 2. *Nature*, 160:486–492, 1947.

[123] T.D. Lee. A Brief History of the Muon. *Hyperfine Interact.*, 86:439–453, 1994.

[124] C.S. Wu and V.W. Hughes. *Muon Physics: Introduction and History.* Academic Press Inc., 1977.

[125] R.A. Millikan. Mesotron as the Name of the New Particle. *Phys. Rev.*, 1:105, December 1938.

[126] C.D. Anderson and S.H. Neddermeyer. Mesotron (Intermediate Particle) as a Name for the New Particles of Intermediate Mass. *Nature*, 142:878, November 1938.

[127] H. Lyle. Anderson, Carl. Interview by Harriett Lyle. Oral History Project, California Institute of Technology. Retrieved April 2008 from the World Wide Web: `http://resolver.caltech.edu/CaltechOH:OH_Anderson_C`, 1979.

[128] A.H. Compton. Foreword: Symposium on Cosmic Rays. *Rev. Mod. Phys.*, 11:122, 1939.

[129] H.J. Bhabha. The Fundamental Length Introduced by the Theory of the Mesotron (Meson). *Nature*, 143:276–277, Feburary 1939.

[130] J. L. Beveridge *et al.* A spin rotator for surface μ^+ beams on the new M20 muon channel at TRIUMF. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 240:316–322, 1985.

[131] Triumf CMMS. Helios superconducting solenoid. Website: `http://cmms.triumf.ca/equip/helios.html`.

[132] G. Morris, TRIUMF, Vancouver, Canada. Private communication, 2009.

[133] S.R. Dunsiger *et al.* Magnetic field dependence of muon spin relaxation in geometrically frustrated $Gd_2Ti_2O_7$. *Phys. Rev. B*, 73(172418), 2006.

[134] F. James. MINUIT – Function Minimization and Error Analysis, CERN Program Library Entry D 506.

[135] J.F. Ziegler. The stopping and range of ions in matter. http://www.srim.org/.

[136] TRIUMF CS. TRIUMF Computing Services description of SRIM. http://it-services.triumf.ca/scientific-computing/software/detector-engineering/srim-1.

[137] D.M. Garner. *Application of the muonium spin rotation technique to a study of the gas phase chemical kinetic of muonium reactions with the halogens and hydrogen halides.* PhD thesis, University of British Columbia, June 1979.

[138] J. Doornbos, TRIUMF, Vancouver, Canada. Private communication, 2009.

Appendix A

Personal contributions

The TWIST collaboration has about 30 members. For this final measurement there were three graduate students writing theses, including myself. I joined the experiment in September 2005, at which time the previous $P_\mu^\pi \xi$ measurement was being completed. I participated in an engineering run from October 2005 to December 2005, and helped in the design and construction of a new downstream trigger (see Appendix E). The data for this thesis were accumulated from October 2006 to December 2006, and June/July of 2007. As well as taking shifts during these periods, I was the run coordinator for October/November of 2006, and in July of 2007.

I was then part of the final review of the analysis code, which included the reconstruction and detector simulation software. I focussed on the aspects that were most relevant to $P_\mu^\pi \xi$. The reports that accompanied these studies are listed in Table A.1, and some of these were edited for inclusion into this thesis. These reports were distributed within the TWIST collaboration, after detailed discussions with Glen Marshall and Richard Mischke. The most significant projects (in addition to the $P_\mu^\pi \xi$ systematic uncertainties) were thorough quality of data checks and a review of the analysis code for the time expansion chambers. I also completed a review of the theory behind the depolarisation of muons in high purity metals.

From October 2008 until June 2009 I was the coordinator of submissions to the WestGrid computing facility, where the experiment had a continuous allocation of 150 CPUs. During this period I prepared and submitted most of the simulations that constituted our first round of statistics.

I also helped in a subsidiary μ^+SR experiment in December 2006, and was then largely responsible for the analysis of the data. This is described in Appendix H.

Table A.1: Internal analysis reports produced by the author. The reports that were used in this thesis are indicated.

Date	Title	Pages	Reference
Mar 2006	B2 scan with target stops	8	-
Jun 2006	B2 scan with target stops using Monte Carlo	10	-
Jun 2006	Quad Steering: combining results	5	Section 2.4
Oct 2006	Muon stopping position for aluminium and silver	5	Section 2.11
Apr 2007	Documentation: MSOR6 procedure	7	Section 4.3
Sep 2007	Muon momentum dependence on TEC position	14	Section 4.3
Nov 2007	Systematic uncertainty due to target depolarisation in 2004	9	Section 6.2.10
Dec 2007	Time dependent depolarisation in Jodidio's paper	7	-
Dec 2007	Comparing beam profiles between 2004 passes 1 and 2	6	-
Mar 2008	μ^+SR measurements in high purity metal targets	66	Appendix H
Mar 2008	TEC code review	54	Appendix G
Apr 2008	TEC pattern-rec improvements	15	Appendix G
Apr 2008	TEC short width hits	12	Appendix G
Apr 2008	TEC miscellaneous studies	8	Appendix G
May 2008	Quality of data checks for 2006/7 Part I	26	Section 5.4
Sep 2008	Lower momentum data and the systematic uncertainty	9	Section 6.2.11
Sep 2008	$(g-2)$ and the M13 beam line	3	Appendix I
Sep 2008	Preparing g501 and onwards	7	Chapter 4
Sep 2008	Muon spin in `GEANT`	30	Section 1.6
Oct 2008	Depolarisation in silver	12	Section 6.2.10
Oct 2008	Quality of data checks for 2006/7 Part II	64	Section 5.4
Nov 2008	Burying the missing material	20	Section 2.11
Dec 2008	Initial muon momentum in `GEANT`	19	Section 4.3
Dec 2008	Fringe field shape systematic uncertainty: First estimate	11	Section 6.2.1
Jan 2009	Depolarisation in aluminium and silver: Theory and literature review	26	Section 1.6
Feb 2009	Fringe field studies: Measures to compare P_μ	7	Section 6.2.1
Feb 2009	Asymmetry analysis: short and long times	5	Section 6.2.10
Apr 2009	Fringe field studies: Magnetic field maps	20	Section 6.2.1
Apr 2009	Stability of the TECs and muon beam	54	Section 6.2.1

Appendix B

History of the muon

In 1935, Yukawa postulated the existence of a new field to explain the binding between nucleons[113], with the force mediated by a heavy quantum. The electromagnetic field was already known to mediate force with its own light quantum, namely the photon. Yukawa predicted the nuclear force quantum would have unit charge, mass "2×10^2 times as large as the electron", and substantial interactions with matter. The new field would be very strong, but only over a small range due to the quantum's mass.

In 1936, Anderson and Neddermeyer published cloud chamber observations of cosmic rays[114], and noted "About one percent of the exposures ... reveal the presence of strongly ionizing particles which in most cases seem to be protons ... and usually arise from a type of nuclear disintegration not heretofore observed." A similar observation of a single track had been made in 1933[115]. In the following years, Anderson, Neddermeyer and other independent experimenters published new results on these particles, with the mass "ranging from 120 to about 400 electron masses"[116], "about 240 electron-masses"[116] and eventually "in the neighbourhood of 200 electron masses"[117]. Many people believed that Yukawa's postulated particle had been discovered.

World War II interrupted pure physics research. Shortly after the war, physicists were keen to understand the muon better, and identify whether it was really the Yukawa particle. The particle appeared to have the correct charge and mass, but Yukawa had predicted a particle that interacted strongly with matter, with "its flight halting abruptly the very moment it passed an atomic nucleus"[118]. In 1946, Conversi, Pancini and Piccioni demonstrated that positive and negative muons stopping in graphite produced a similar number of decay electrons[119], in contrast to the theoretical prediction that negative muons should produce far less decay electrons due to nuclear capture. Fermi, Teller and Weisskopf subsequently concluded "the interaction of mesotrons with nucleons according to the conventional schemes is many orders of magnitude weaker than usually assumed"[120]. The muon no longer appeared to be the Yukawa particle.

The issue was resolved when Powell's collaboration discovered the real Yukawa particle, the pion[121, 122]. They observed "two types of mesons exist, of different mass, which we refer to as $\pi-$ and $\mu-$ mesons"[122]. The pion was short lived and therefore had been harder to detect. The discovery of the pion meant the muon was not predicted at all, leading the theorist I.I. Rabi to allegedly exclaim "Who ordered that?"[123]. The mass of the muon is now determined as 105.7 MeV (211 electron masses), and the charged pions have mass 139.6 MeV (280 electron masses). The true carriers of the strong force are also identified as the gluons, not the pion.

The interested reader can find more detailed accounts of the muon's history in Refs. [4, 118, 123, 124].

Appendix C

Naming of the muon

Historically the muon has also gone by the name "mesotron" and "mu-meson". The occurrence of the various names in abstracts of the American Physics Society journals is shown in Fig. C.1. The history of the particle's name has involved several Nobel prize winning Physicists, and will now be summarised using first hand evidence wherever possible.

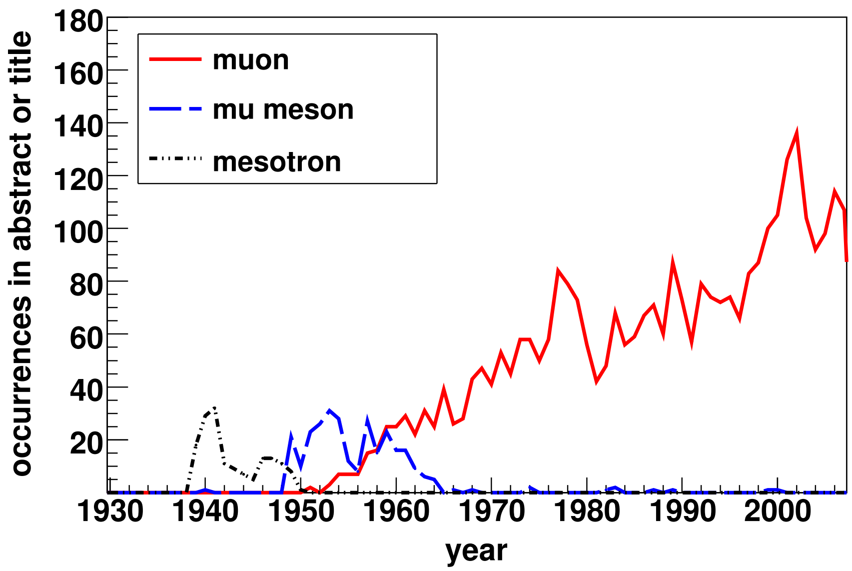

Figure C.1: Frequency of "muon", "mu-meson" or "mesotron" in the American Physical Society abstracts.

The initial name suggested by N. Bohr[48] was "yucon"[125], since it was initially believed to be the strong force particle predicted by Yukawa in 1935. The name "mesotron" was first suggested in a 1938 Nature article by C.D. Anderson and S.H. Neddermeyer[126]. Prior to this article, the particle was known be a variety of names such as "Yukon for

[48]N. Bohr was awarded the 1922 Nobel Prize for Physics "for his services in the investigation of the structure of atoms and of the radiation emanating from them".

Yukawa ... X-particle ... heavy electron ... baryon"[127], and additionally "dynatron, penetron, barytron"[126]. The Nature article itself has a colourful history, revealed in a 1979 interview[127] with C. Anderson:

> "So Seth and I – Millikan was away – wrote a little note, one paragraph, to Nature suggesting that the name of it be mesoton – "meso" meaning intermediate, like mezzanine in a building ... When Millikan[49] came back, I told him about this and showed him the letter, and he hit the ceiling. He said "That's not a good word. It should be mesotron. There should be an "r" in there." And he said, "Look, there's electron, there's a neutron. And I said, "There's proton." Well, the end and issue of it was that Seth and I cabled that "r" to Nature, and it came out mesotron, a word which I didn't like – nobody liked it"

The Nature article argues that "it does appear quite certain that the mass, whether unique or not, is greater than that of an electron and less than that of a proton", and hence the particle was given a name indicating its mass was in between the two well established particles. Amusingly, the Nature articles finishes with, "It appears quite likely that the appropriateness of this name will not be lost, whatever new facts concerning these particles may be learned in the future".

A month later, R.A. Millikan wrote a short note in Physical Review, where he quotes a letter from Bohr:

> "I take pleasure in telling you that every one at a small conference on cosmic-ray problems, including Auger, Blackett, Fermi[50], Heisenberg[51], and Rossi ... was in complete agreement with Anderson's proposal of the name 'mesotron' for the penetrating cosmic-ray particles."

C. Anderson's Nature article and R. Millikan's Physical Review article apparently failed to win popular consensus, as made apparent in a foreword by A.H. Compton[52] for a 1939 "Symposium on Cosmic Rays"[128]:

[49] C. Anderson was supervised by R.A. Millikan, who was awarded the 1923 Nobel Prize for Physics for his measurement of the electron's charge.

[50] E. Fermi was awarded the 1938 Nobel Prize in Physics for "his demonstrations of the existence of new radioactive elements by neutron irradiation, and for his related discovery of nuclear reactions brought about by slow neutrinos".

[51] W.K. Heisenberg was awarded the 1932 Nobel Prize in Physics for "the creation of quantum mechanics".

[52] A.H. Compton was awarded the 1927 Nobel Prize in Physics for his work on the change in x-ray wavelengths upon scattering.

> "An editorial problem has arisen with regard to the designation of the particle of mass intermediate between the electron and the proton ... A vote indicated about equal choice between *meson* and *mesotron* with no considerable support for *mesoton*, *barytron*, *yukon* or *heavy electron*. Except where the authors have indicated a distinct preference to the contrary, we have chosen to use the term *mesotron*."

The transition from "mesotron" to "meson" appears to have been initiated by a footnote in a 1939 Nature article[129] by H.J. Bhabha:

> "It is felt that the 'tr' in this word is redundant, since it does not belong to the Greek root 'meso' for middle; the 'tr' in neutron and electron belong, of course, to the roots "neutr" and "electra" ... It would therefore be more logical and also shorter to call the new particle a meson instead of a mesotron."

C. Anderson described other objections to the term "mesotron", since "tron" is usually reserved for instruments, such as "audiotron or cyclotron or synchrotron"[127].

The pion was discovered in 1947 by C.F. Powell. According to Ref. [118], Powell named the new heavier particles "pi-mesons" and the older particles "mu-mesons". Additionally, C. Anderson credits Powell with the use of greek letters: "And (Powell) used, I guess for his own bookkeeping, Greek letters – pi and mu". Lee[123] suggests that Powell "introduced the symbol π which stands for 'primary"', but credits Fermi as being the first to use the symbol μ.

Appendix D

Magnetic field mapping and simulation

Prior to data acquisition, the magnetic field was measured at $B = 1.96\,\mathrm{T}$, $B = 2.00\,\mathrm{T}$ and $B = 2.04\,\mathrm{T}$ using the custom built apparatus shown in Figs. D.1. and D.2. Hall probes were placed along an arm at intervals of 4.13 cm, and the apparatus was rotated by 15° increments and moved in steps of 5 cm or 2.5 cm longitudinally[53] (z). The Hall probes only measured the longitudinal component of the field, and were calibrated using an NMR probe that measured the total field. The Hall probe measurements had a precision of 0.1 mT. In 2002 measurements were taken of the tracking region using seven probes on the rotating arm. In 2003 measurements of the fringe field and the upstream part of the tracking region were taken. For the 2003 measurements only five Hall probes were used since the assembly needed to pass through the hole in the doors.

A simulation was required to smooth out the field and generate the x and y components. The `Opera` software package was used for this purpose[87]. The measurements taken with the Hall probe mapper were not given to `Opera`. Instead the pieces shown in Fig. D.3 were modelled, and their properties[54] were adjusted to minimise the differences between `Opera` and the measurements. The magnetic field from the TRIUMF cyclotron was not included since it is only $\approx 4\,\mu\mathrm{T}$ at the TWIST detector. The agreement between `Opera` and the measured map was better than 0.2 mT within the 2 T region. Further upstream, the match becomes worse as demonstrated in Fig. D.4.

[53]The smaller steps were taken at the edges of the tracking region.

[54]Specifically, the coil position, current density, coil radii, material B-H curves and door position were adjusted.

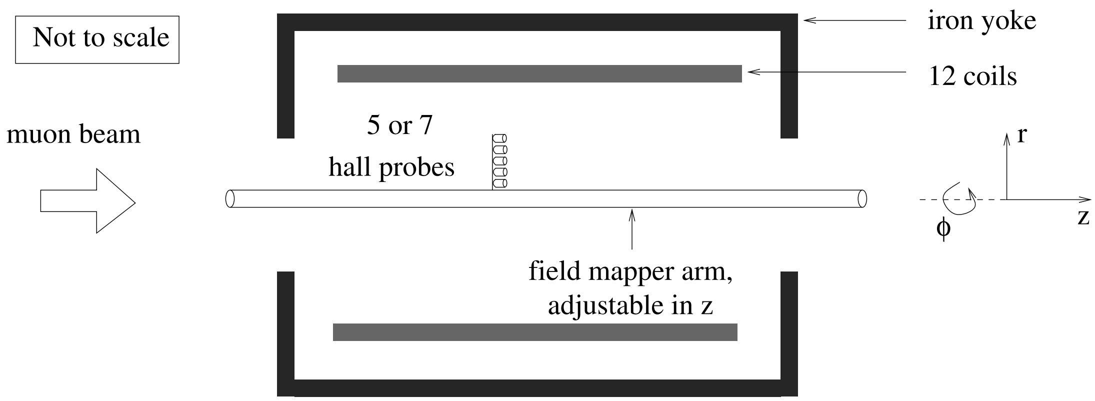

Figure D.1: Schematic of the field mapper device. Hall probes were attached at radial intervals of 4.13 cm, and the arm was rotated in 15° increments. The whole assembly was inserted from the downstream end.

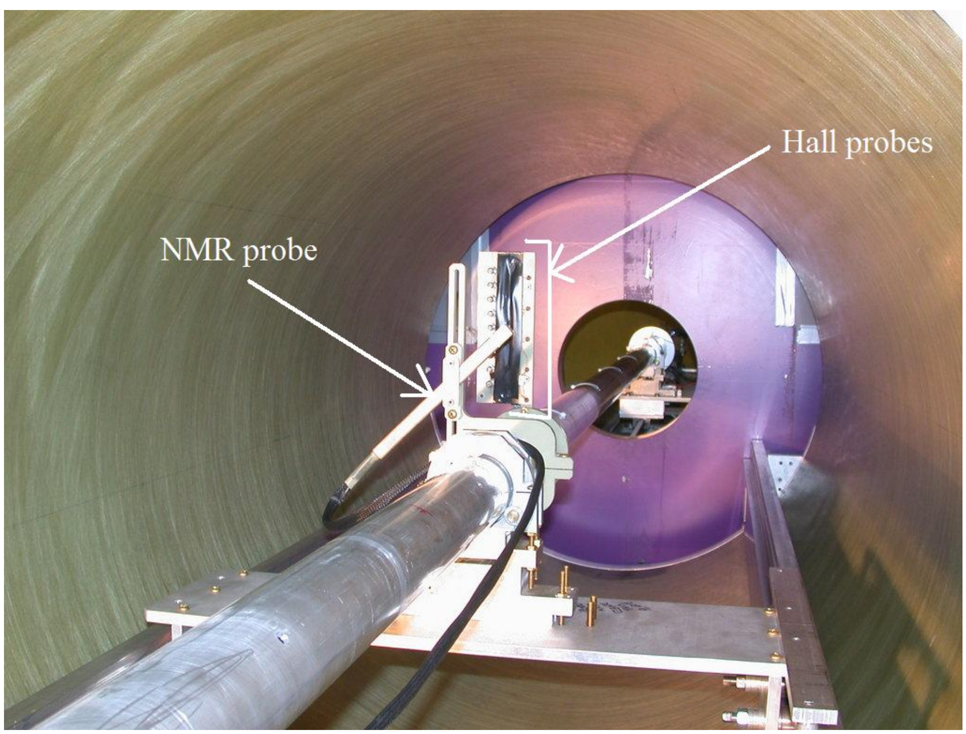

Figure D.2: Photograph of the field mapper, taken from inside the solenoid. The bore of the magnet has 1 m diameter. The device is constructed from non-magnetic materials (aluminium, brass).

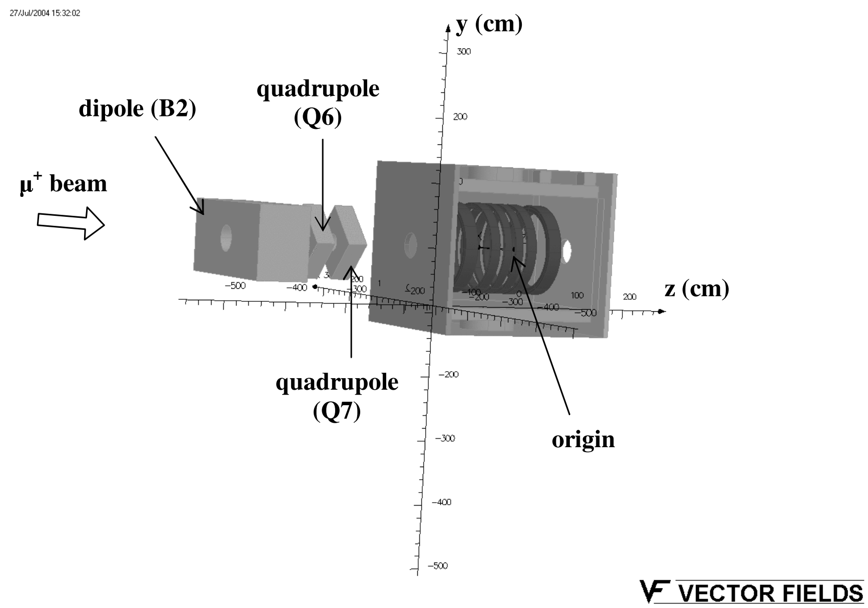

Figure D.3: Components included in the `Opera` simulation of the magnetic field.

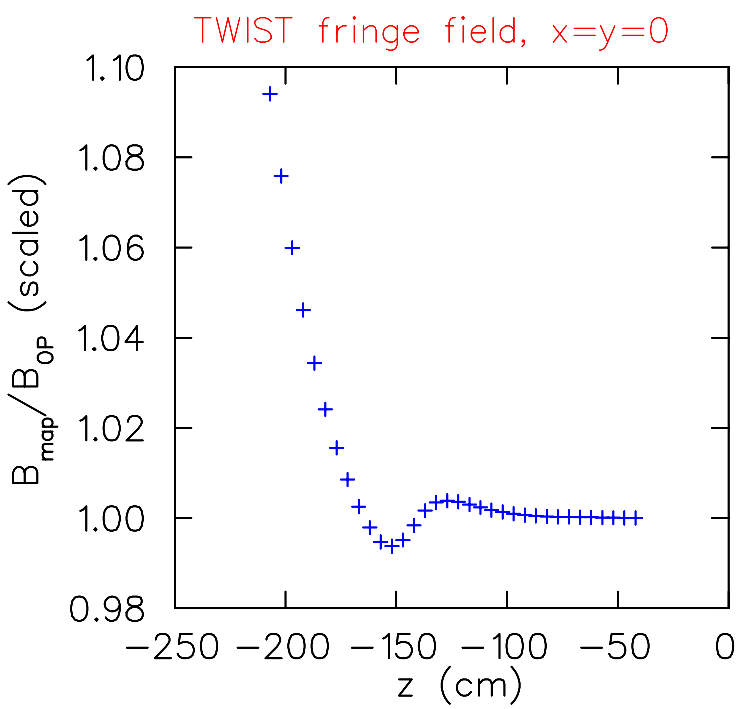

Figure D.4: For $x = y = 0$, the ratio of the measured field map to the `Opera` model[100].

Appendix E

New downstream trigger

The determination of the wire timing offsets required a reliable downstream trigger, even when the experiment's magnetic field was switched on (see Section 3.2.1). The author's first project was to help in the design and construction of such a trigger. The final design is photographed in Fig. E.1, where two scintillators of thickness 0.318 cm, separated by 0.318 cm, are read out by light guides that are matched to long (~ 2 m) light pipes. These transported the collected light to photomultiplier tubes located in a relatively weak field region. The front of the arrangement is shown in Fig. E.2, where light guides are connected to the top and bottom of each scintillator, resulting in a four-fold coincidence that gives a reliable timing measurement.

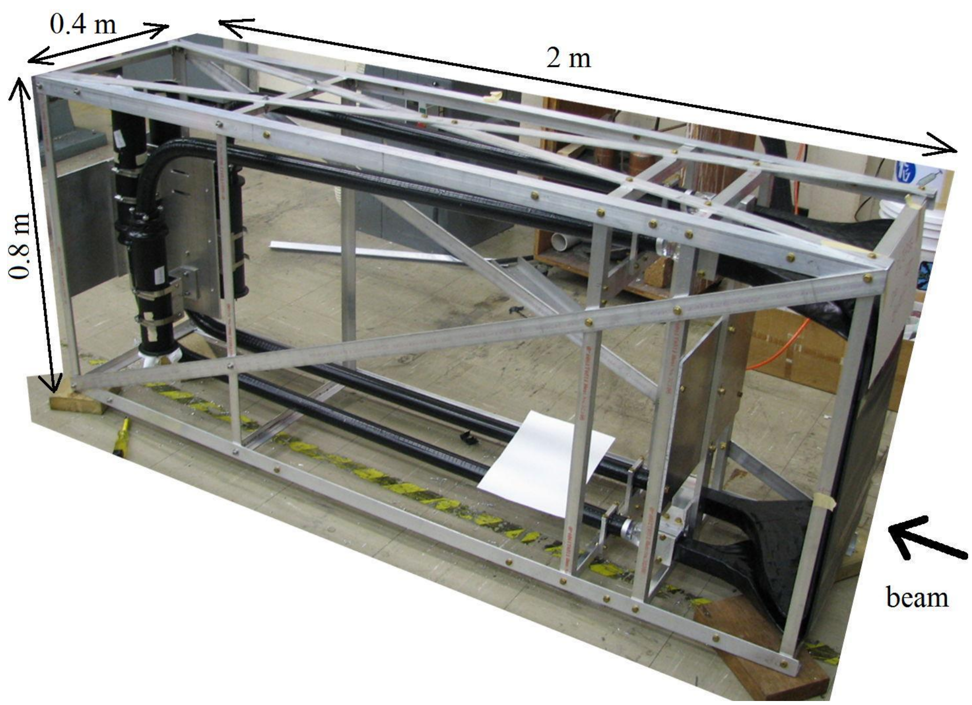

Figure E.1: Downstream trigger: photograph of completed frame, light guides, photomultiplier tubes and scintillators.

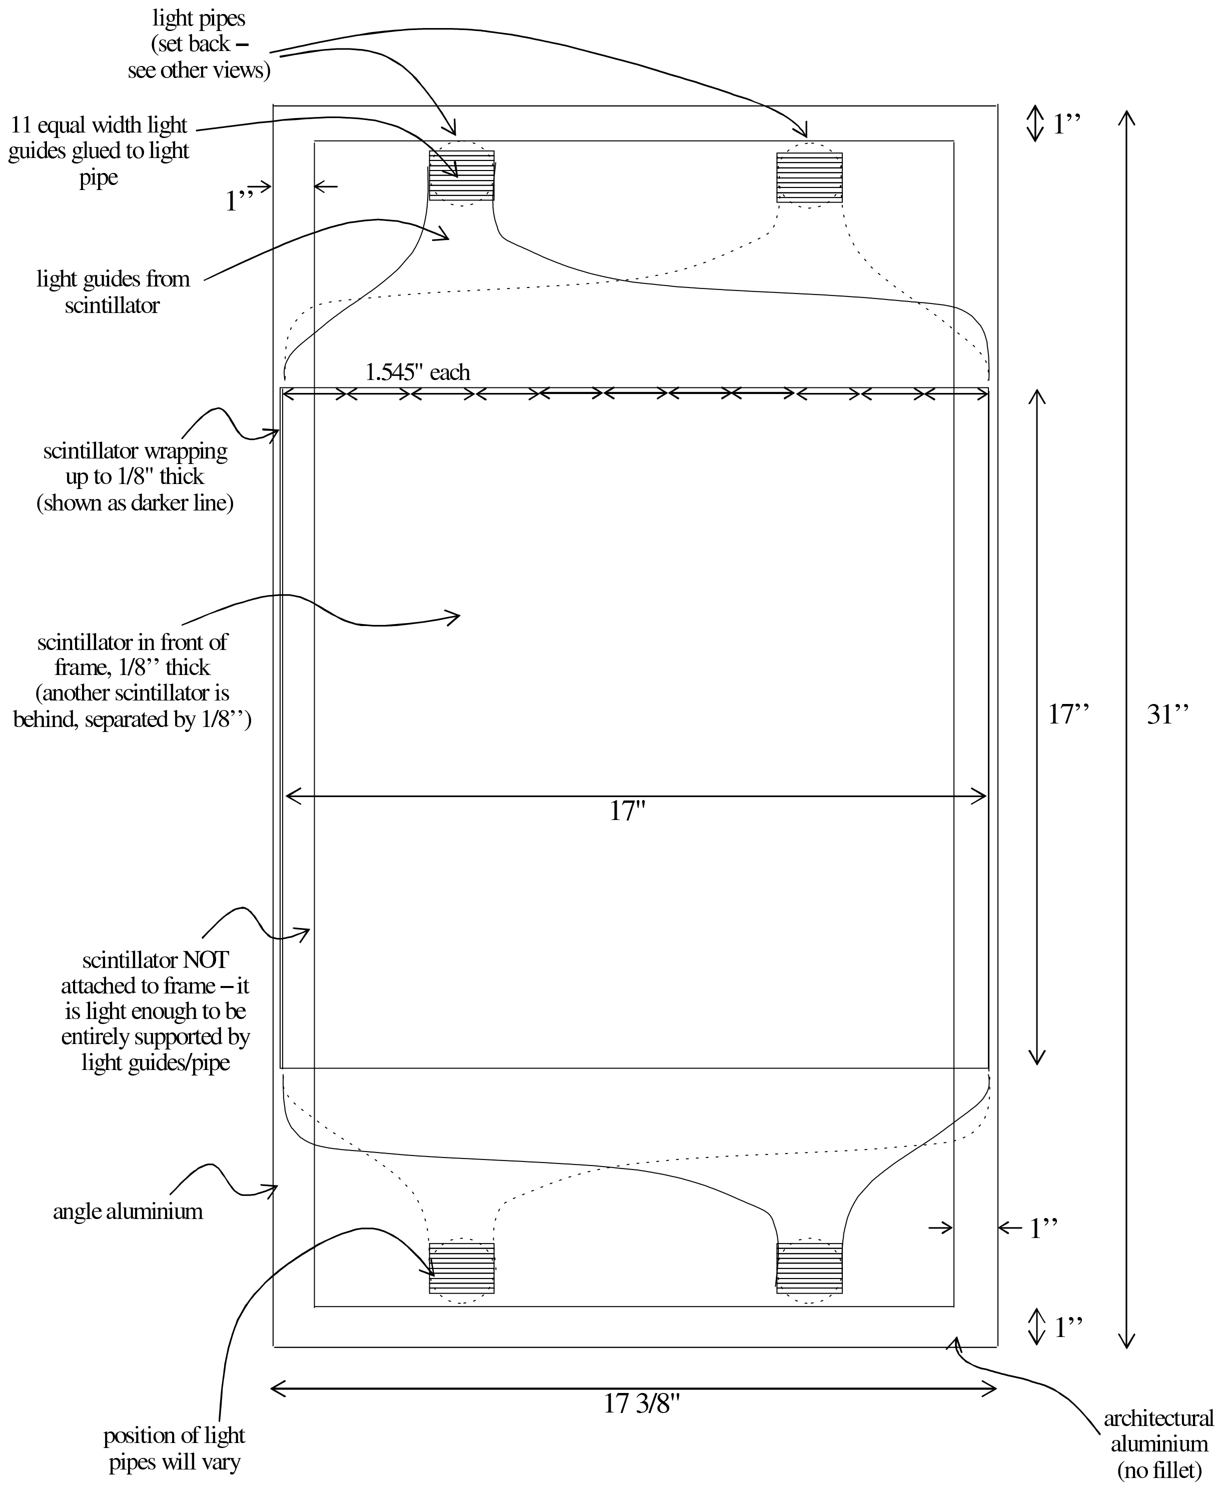

Figure E.2: Front view (looking downstream) of new downstream trigger. Measurements are in inches.

Appendix F

Bias in $P_\mu^\pi \xi$ extraction technique

The $P_\mu^\pi \xi$ result is quoted at the time of muon production, but the data and simulation spectra are constructed with P_μ at the time of decay, after all depolarisation processes have taken place. Algebraically, the difference in the product $P_\mu \xi$ between the data and simulation spectra can be written as

$$\Delta \left[P_\mu \xi \right] = P_\mu^{D,data} \xi^{data} - P_\mu^{D,sim} \xi^{sim}, \tag{F.1}$$

where the superscript D in P_μ^D implies that P_μ is at the time of decay, and *sim* refers to the simulation's values. The quantity $\Delta \left[P_\mu \xi \right]$ is assigned uncertainties. When the experiment reveals the hidden value of ξ used in the simulation, the result is then stated as,

$$P_\mu^{\pi,data} \xi^{data} = P_\mu^{\pi,sim} \xi^{sim} + \Delta \left[P_\mu \xi \right]. \tag{F.2}$$

where the superscript π in P_μ^π implies that P_μ is at the time of muon production. The experiment claims to accurately simulate the depolarisation between production and decay, so that

$$P_\mu^{\pi,data} - P_\mu^{D,data} = P_\mu^{\pi,sim} - P_\mu^{D,sim}. \tag{F.3}$$

Equations (F.1) and (F.3) can be re-arranged to yield

$$\begin{aligned} \Delta \left[P_\mu \xi \right] &= P_\mu^{\pi,data} \xi^{data} - P_\mu^{\pi,sim} \xi^{sim} \\ &\quad + \left(P_\mu^{D,sim} - P_\mu^{\pi,sim} \right) \left(\xi^{data} - \xi^{sim} \right). \end{aligned} \tag{F.4}$$

Now $\left(P_\mu^{D,sim} - P_\mu^{\pi,sim} \right)$ is the amount of depolarisation in the simulation, which is of order 10^{-3}. The term $\left(\xi^{data} - \xi^{sim} \right)$ depends on the tolerance of the hidden values, which is at most 10^{-2}. Therefore this term is of order 10^{-5}, and can be safely neglected. In other words, even though spectra at the time of decay are being compared, the result for $P_\mu^\pi \xi$ at the time of muon production can be extracted with sufficient accuracy for this measurement.

Appendix G

Time expansion chambers analysis

G.1 Introduction

An accurate measurement of the muon beam was critical in determining the muons' polarisation at the time of decay. This was achieved by determining the position and angle of individual muons using the time expansion chambers (TECs) that were described in Section 2.3. The analysis that converted signals from the TEC electronics into an unbiased measurement of the muon beam will be described.

The TEC analysis code was extensively reviewed for this measurement. The track reconstruction software did not limit the precision of the $P_\mu^\pi \xi$ measurement; instead the reproducibility of the TEC hardware and noise from the electronics meant that further improvements in the TEC analysis code were unwarranted.

G.2 Typical raw events

The TECs were in a region of weak magnetic field, so that muon trajectories were well approximated by straight lines. The majority of events were like Fig. G.1(a), with a single clean track through both the x and y-module, and most of the wires producing at least one distinct signal pulse ("hit"). Some of the wires had multiple hits due to the break-up of ionisation into clusters, and crosstalk from the electronics, and this prevented a simple straight line fit from being applied to each hit's leading edge. A minority of events suffered from electronics noise, sparks, beam positrons and multiple muons, and examples of these are shown in Figs. G.1(b) to G.1(d).

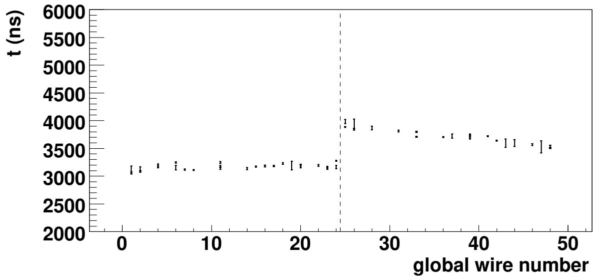

(a) The majority of events had one clean muon track in each module.

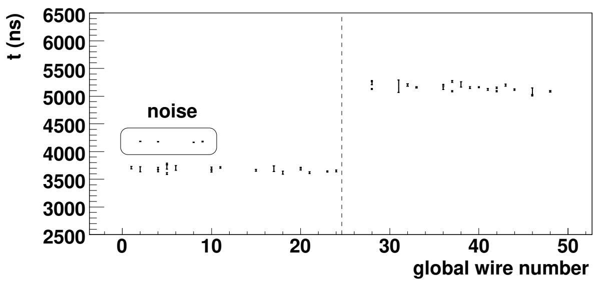

(b) x-module noise that must be removed. This was localised to seven wires of the x-module.

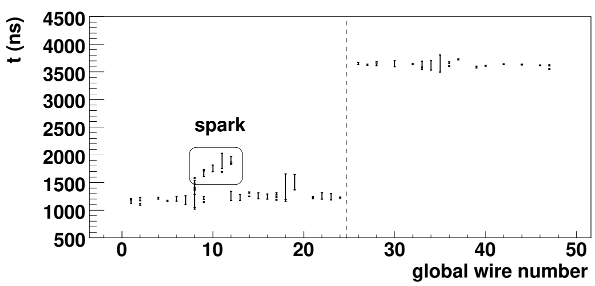

(c) x-module sparking.

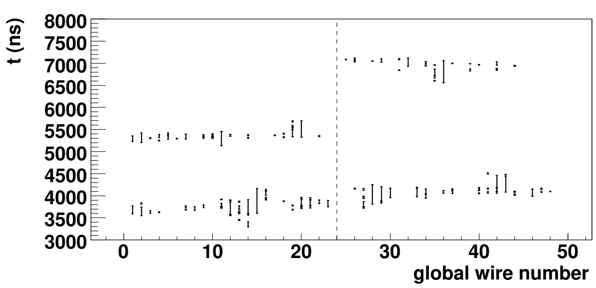

(d) Multiple muon event. Ordinate scale is changed.

Figure G.1: Examples of TEC events from data. Wires 1 to 24 are the x-module, and 25 to 48 are the y-module.

G.3 Analysis

G.3.1 Unpacking and wire time offsets

A TEC event was analysed if a particle triggered the muon scintillator, and a positron was then detected more than 700 ns from the trigger time. This corresponded to $\approx 80\%$ of trigger particles, although including the other 20% was found to introduce no measurable bias. The TEC analysis removed events where the trigger particle was not a surface muon, as determined by the the particle identification from Section 3.3.1.

The leading edge times and hit widths were first "unpacked" from the data files, using the analysis code described in Section 3.2.1. Hits with an error code from the TDC were discarded, and wire time offsets were applied to the good hits.

When the TEC sense planes were new, on average there were 17 out of 24 wires with at least one hit. This was lower than the original design specification, but was adequate for this measurement. The hit efficiency decreased as the TECs aged, requiring the sense planes to be replaced when an average of 10 wires had at least one hit.

Crosstalk and electronics noise produced hits that were all less than 20 ns in width. Unfortunately these could not be removed by a simple pulse width cut, since hits from the muon ionisation could be as narrow as 8 ns. Placing a pulse width cut to remove noise would have removed real muon hits, thus reducing the single hit efficiency and degrading the resolution. The noise was instead removed by improving the pattern recognition, which will be described later.

G.3.2 Reject multiple trigger events

Events with multiple muons were rejected since they produced two tracks in each TEC module, and these could not be matched up between the modules (see Fig. G.1(d)). If two muons were detected by the spectrometer's proportional chambers, then the event was removed from the TEC analysis. This is the only place where information from inside the detector was used.

The number of scintillator hits was also available, but was not used by default. Tests confirmed that adding a stricter multiple trigger rejection based on the number of scintillator hits had no measurable effect on the analysis.

G.3.3 Discriminator amplitude walk

For a discriminator with a fixed threshold, there exists a correlation between signal amplitude and leading edge time. At this stage in the analysis, hits with a width less than 120 ns were corrected according to

$$\Delta t = \begin{cases} (100.0 A_i)\,\mathrm{ns} & (0.0 < t_{\mathrm{width}} < 20.0)\,\mathrm{ns} \\ (120.0\,\mathrm{ns} - t_{\mathrm{width}})\,A_i & (20.0 < t_{\mathrm{width}} < 120.0)\,\mathrm{ns} \end{cases} \tag{G.1}$$

where A_i (no units) is determined for each wire individually as part of the calibration.

In the previous $P_\mu^\pi \xi$ analysis[59], the calibration of A_i for each wire was not iterated until convergence. As a result, the $P_\mu^\pi \xi$ analysis was sensitive to whether the discriminator amplitude walk correction was made. The current analysis found that disabling the correction altogether had a negligible effect on the muon beam measurement.

G.3.4 Track candidates

Track candidates were constructed according to the algorithm that is outlined in Fig. G.2, which has the following steps:

1. Select a single hit.
2. Draw a corridor ± 300 ns from this hit.
3. On each wire, pick up the single hit that is within the corridor, and closest to the hit in (1).
4. If there are more than four hits, then keep this selection as a track candidate.
5. Reject track candidate if the hits span less than 2.7 cm (14 out of 24 wires).
6. Repeat (1) to (5) for every hit.

This algorithm gave every hit a chance of being a good hit, and only allowed one hit per wire. The rejection of track candidates based on their hit span ensured good angular resolution.

The track candidate with the most number of hits (n_{max}) was then found. The other track candidates were only kept if they had enough hits[55]. The leading edge times in the

[55]The criteria was

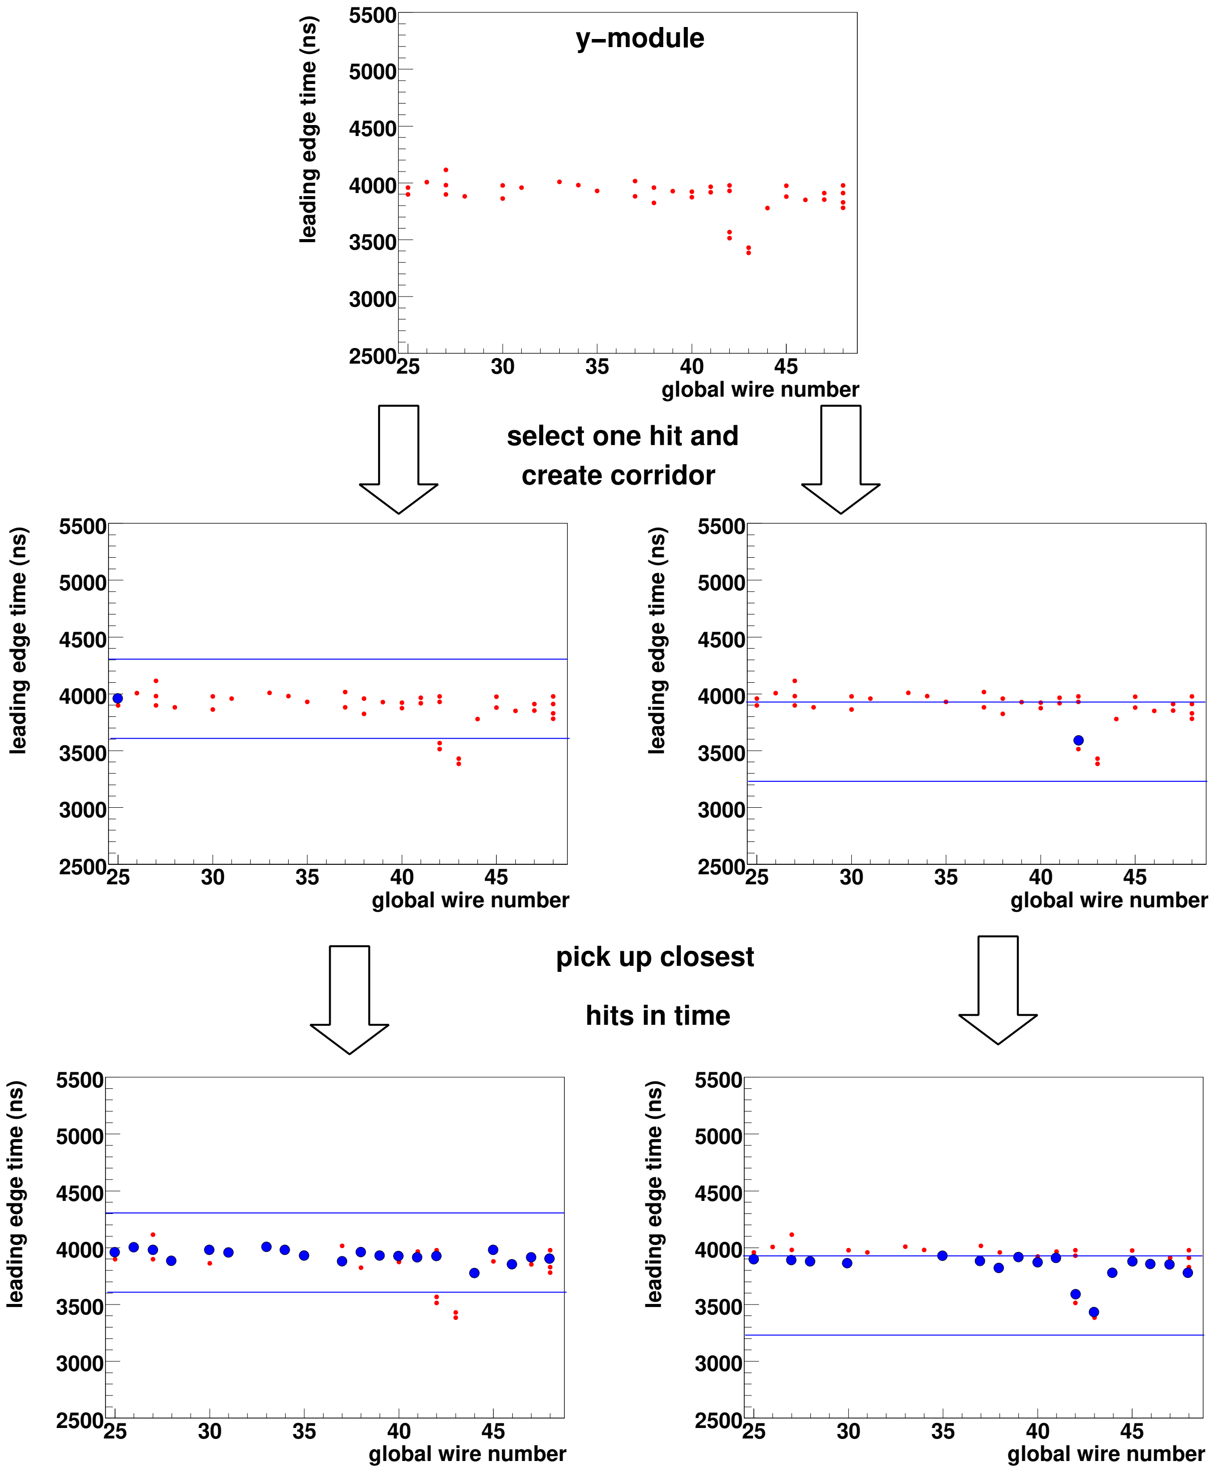

Figure G.2: Construction of two possible track candidates. See the text for more details.

track candidates were then converted to distance, using a space-time-relationship (STR) described by a cubic function,

$$s = p_0 + p_1 t + p_2 t^2 + p_3 t^3, \tag{G.2}$$

where s is x or y, depending on the module, and p_i are determined by a calibration technique that will be described later. A least-squares straight line fit was then carried out to the positions, and the track candidate with the smallest residuals was declared as the successful muon trajectory.

Finally, to avoid a suspected bias towards low angles from the fixed corridor, hits were recovered within 0.15 cm of the fitted muon track and the straight line fit was repeated. This process was iterated until convergence. Unfortunately it was not possible to definitively say whether recovering hits and refitting offered a genuine improvement; this was judged by examining individual event displays. The average position and angle were insensitive to whether an iteration was made, but the RMS (root mean square) of the reconstructed angles changed by about 4 mrad. This introduced a $P_\mu^\pi \xi$ systematic uncertainty, as described in Section 6.2.8.

G.3.5 Final multiple track removal

Even though events with two muons were reliably removed using the earlier event classification cut, there is still the possibility of multiple tracks due to noise and sparks (see Figs. G.1(b) and G.1(c)), and beam positrons, which were detected with low efficiency but could produce enough hits to make a track. In a module, if the best-fit track and another track candidate were too close, then the event is rejected.

The separation time was tuned for the current analysis. The old value of 1 μs was conservatively large, and a setting of 0.3 μs was found to be better at rejecting multiple track events, or those confused by noise.

- For $n_{\max} > 12$, require $n \geq (n_{\max} - 2)$
- For $8 < n_{\max} \leq 12$, require $n \geq (n_{\max} - 1)$.
- For $n_{\max} \leq 8$, require $n = n_{\max}$.

The selection is necessary to prevent the analysis code selecting junk tracks, such as the "spark" in Fig. G.1(c).

G.3.6 Other algorithms

Alternative reconstruction algorithms were tested, such as a Hough Transform, Kalman filter and iterative straight line fitting with outliers rejected at each stage. However they were found to need careful tuning to become robust to the noise in the data, and could not readily determine the RMS of angles better than the algorithm described here.

G.4 Sense plane calibration

Special TEC calibration data were acquired, where each end of the box containing the TECs had an (11×11) hole collimator attached (see Fig. G.3(a)). The beam line was

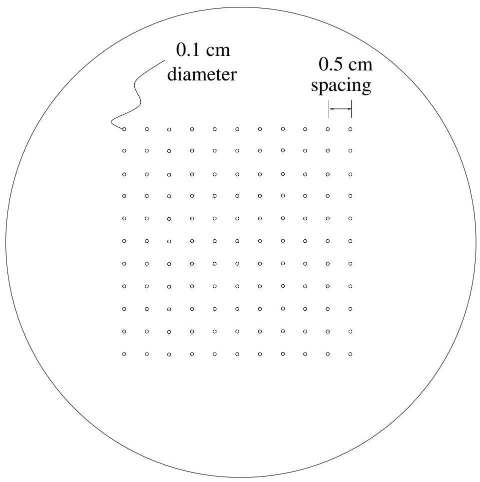

(a) 121 hole collimator.

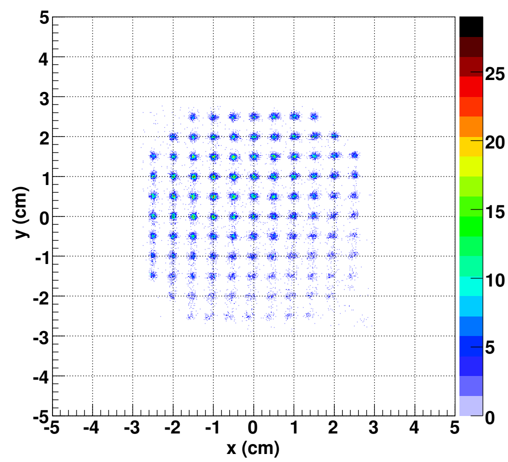

(b) Reconstructed position of muon tracks, after correcting the STRs.

Figure G.3: TEC drift cell space-time relationships (STRs) were determined by placing a 121 hole collimator on each end of the TEC box.

tuned to illuminate the entire collimator, and an angle cut was made in the analysis to select tracks that passed through opposing holes. These tracks allowed determination of the wire time offsets, discriminator amplitude walk correction, and space-time-relationships for each drift cell. The calibration technique and the results from four independent calibrations (labelled 2006H, 2007A, 2007B, 2007C) will be described. Further detail can be found elsewhere[81].

G.4.1 Wire time offsets

The time of TEC hits were measured with respect to the muon trigger scintillator. An offset had to be determined for each TEC wire due to differences in cable lengths and TDC channels. These were determined by selecting tracks from the central hole of the collimator, and histogramming the times for each wire. The offset required to match-up

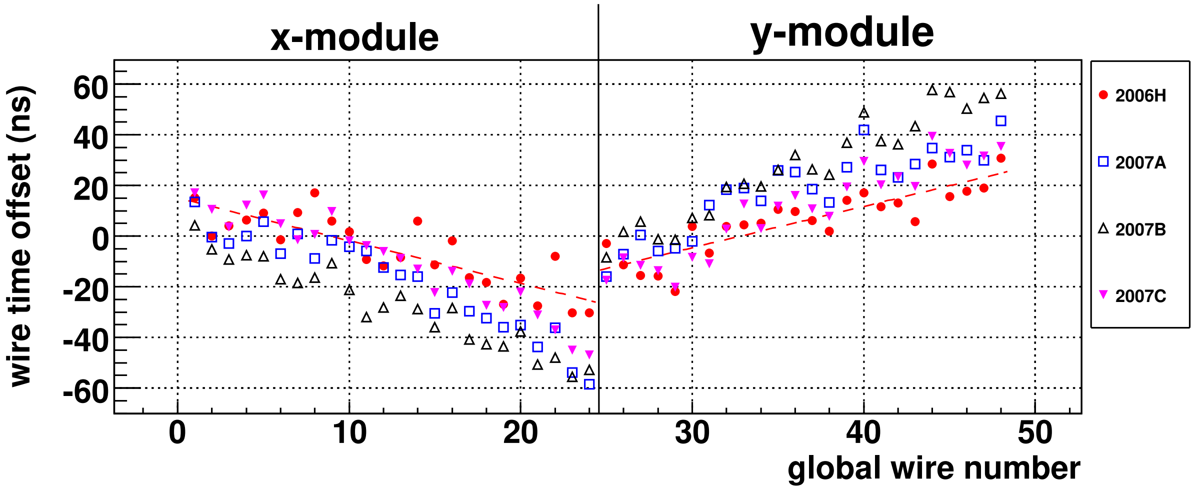

Figure G.4: Wire time offsets for the calibrations in 2006 and 2007. The gradient corresponds to a rotation of the TECs within their box of between $7 - 12$ mrad. A fit line is included for 2006H only.

the peaks of these histograms was then determined. This calibration also found the global time offset of the TEC wires relative to the muon scintillator.

The wire time offsets are compared in Fig. G.4, which shows that there is a slope corresponding to a rotation of $7-12$ mrad in each module[56]. The collimators were aligned to better than 500 μm, which corresponded to 3 mrad over their separation distance of 16 cm; therefore the wire time offset calibration compensated for a rotational misalignment of the individual TEC modules, rather than a collimator misalignment. The previous analysis[59] found similar features that suggested a $2-4$ mrad rotation, but the statistical precision was significantly worse. The possible time variation over a 0.1 cm diameter hole was < 1 ns, so that the remaining features in Fig. G.4 are differences in cable lengths and electronics.

G.4.2 Discriminator amplitude walk

The parameters characterising the discriminator amplitude walk correction (A_i in Eq. (G.1)) are compared for each wire in Fig. G.5. There is no evidence of wire-to-wire differences outside of statistics (if differences existed, this would indicate that the thresh-

[56]The gradient for the 2006H planes is between -1.7 ns and -2.6 ns per wire in the x-module, and between 1.6 ns and 2.8 ns per wire in the y-module. 1 ns is approximately 9 μm, and the wires are 0.2 cm apart, so that the rotation is $8 - 11$ mrad in the x-module, and $7 - 12$ mrad in the y-module.

olds in the electronics were not set to consistent values). The parameters are remarkably consistent between calibrations.

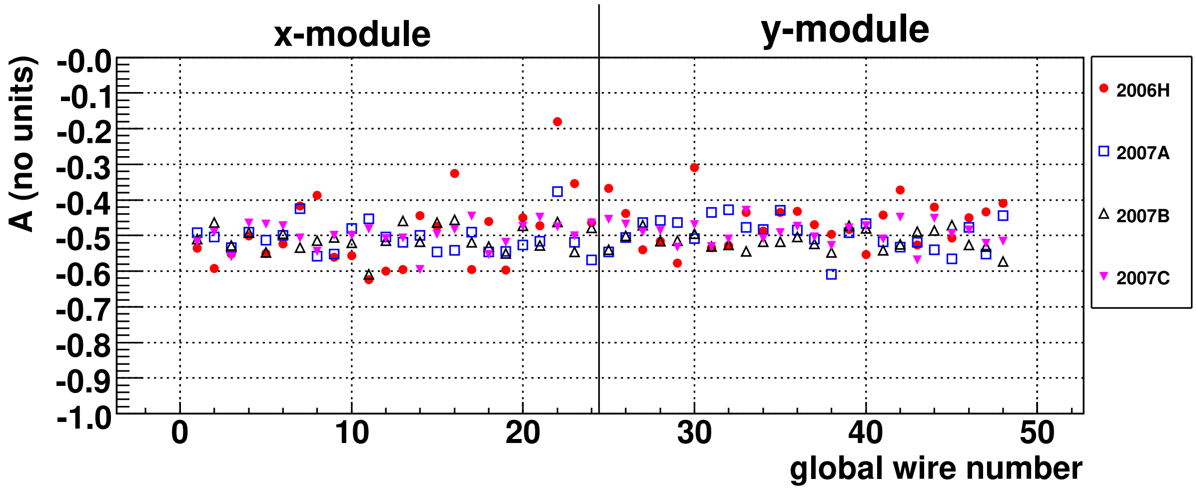

Figure G.5: Discriminator amplitude walk parameters for the calibrations in 2006 and 2007. The correction is independent of wire and sense planes.

G.4.3 Space-time relationship

The space-time-relationship (STR) for each drift cell was initially generated using the `GARFIELD` software[96]. This STR is an approximation since only one TEC module was modelled, which neglected the electric field interference between the two TEC modules. In addition, the `GARFIELD` simulation relied on accurate inputs for the voltage, geometry and gas density. The collimator data were used to adjust the STRs, by altering the parameters in Eq. (G.2) to place the collimator holes at their known positions; Fig. G.3(b) shows an example of the calibration data after the STRs have been corrected.

The gas inside the TECs was held at a constant pressure, so that temperature variations altered the gas density and hence the space-time relationship. A change of $\pm 3°$C altered the mean positions by $< 500\,\mu$m, and the average angles by $< 0.4\,$mrad, which were not significant variations for this analysis. No correction was made due to temperature since each data set had a beam measurement within 1.5°C of the calibration temperature.

G.5 Performance

G.5.1 Resolution

The single hit resolution for each wire was determined as follows:

1. For the track candidate with lowest χ^2, re-fit with a wire excluded.

2. For the excluded wire, histogram the residual distance between the hit and the fit.

3. Find the RMS of the residual histogram. This is a measure of the single hit resolution for the wire.

The resolution was between 275 μm and 375 μm for all wires, as shown in Fig. G.6, which demonstrates that the resolution did not change significantly as the planes aged. The resolution also depended on the distance from the sense plane; an earlier study determined that it varied between 150 μm and 350 μm[78].

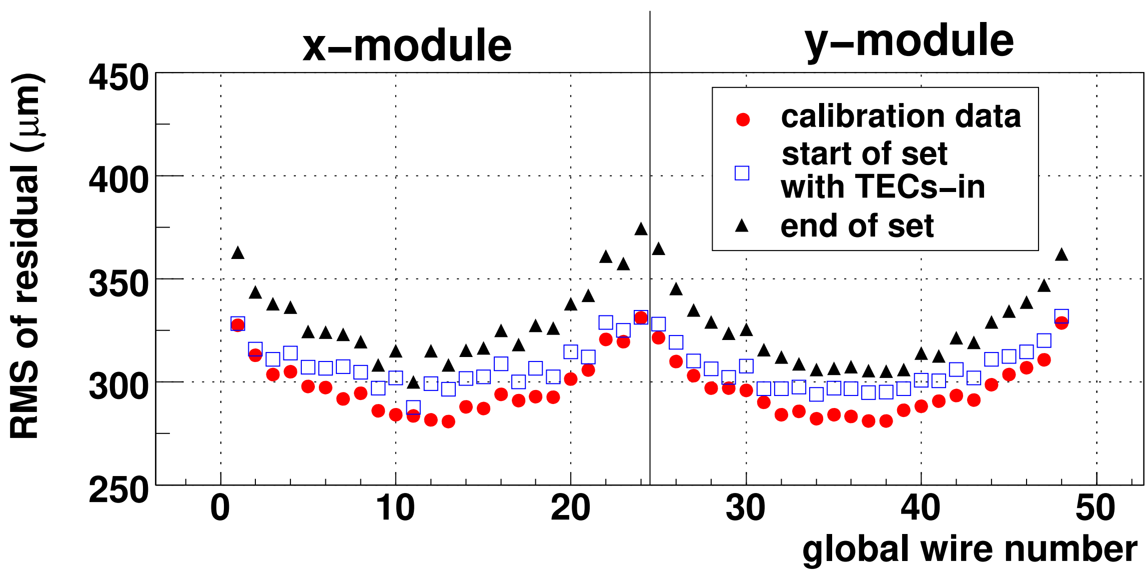

Figure G.6: Resolution (RMS of residual distribution, when wire is excluded from fit). The calibration data corresponds to newly built sense planes. The beginning and end of the set with TECs-in corresponds to a time period of one week, over which the planes aged significantly.

The resolution of the track position and angle could be determined using simulation. With wire inefficiency simulated, and no multiple scattering, the angle resolution for a single track was ~ 3 mrad, and the position resolution was $\sim 150\,\mu$m[59]. However, this is misleading since the resolution in real data will be worse due to the electronics noise and beam positrons.

G.5.2 Number of hits per event

The average number of hits in the final track was about 15. This average depended on the track's distance from the sense plane, since as an ionisation cluster drifts its electrons separate in time, so that a longer drift time resulted in a lower probability of a hit exceeding the electronics threshold. At 2.0 cm (5.0 cm), the average number of hits was 12 (17). There were no asymmetries observed in the beam profile, indicating that even far from the sense plane there were enough hits available.

The number of hits per event decreased as the sense planes aged. This was quantified by accumulating an entire set with the TECs in place, and is shown in Fig. G.7(a). The mean positions and angles were highly stable during this set (see Section 5.4.7), but with less hits available the angular resolution decreased, and the RMS of the reconstructed angles therefore increased. Figure G.7 shows this effect, where the RMS is seen to change by 2 mrad in each module. This information is used in Section 6.2.8 to assign a systematic uncertainty to $P_\mu^\pi \xi$. There was also a negligible increase in the RMS of the positions over the entire set (0.04 cm).

G.6 Summary

As a result of the electronics noise and beam positron signals that were present in the data, the TEC analysis code could not determine the RMS of the reconstructed angle distribution to better than 4 mrad. This will be used to establish a systematic uncertainty. The single hit resolution was better than 375 μm for all drift cells, and the available calibrations for wire time offsets, discriminator amplitude walk and space-time relationships are remarkably stable. The aging of the sense planes was measured, and this will be included in a systematic uncertainty estimate.

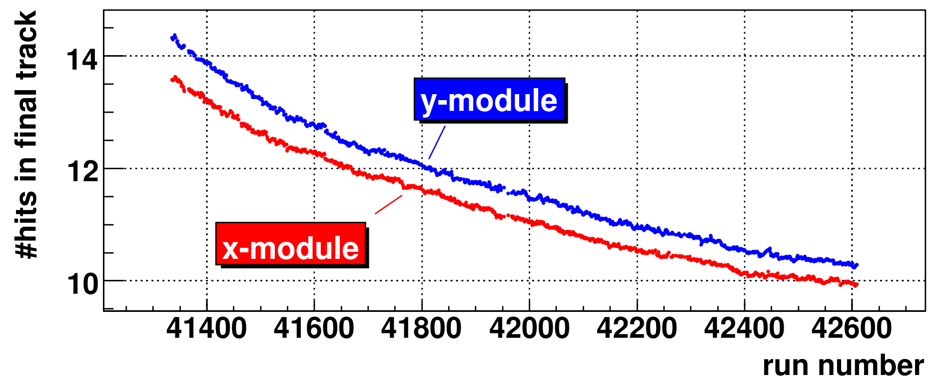

(a) Number of hits used for the successful track.

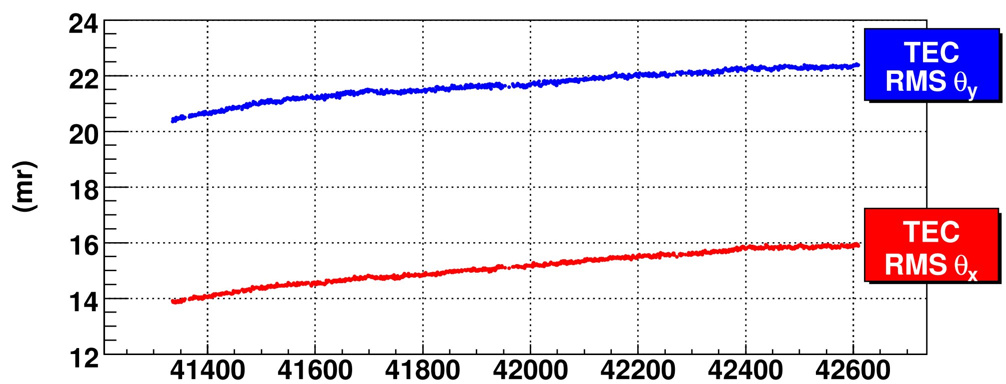

(b) RMS of reconstructed angles.

Figure G.7: Sense plane aging, for one data set (approximately one week). The muon beam was stable during this time period. As the number of hits per event decreases, the angle is less well determined, and the RMS of reconstructed angles increases.

Appendix H

E1111: Subsidiary μ^+SR experiment

H.1 Introduction

In the TWIST experiment, if a muon decayed within $1.05\,\mu$s of the trigger time, the resulting positron was not included in the normal analysis since its hits could overlap with the muon ionisation in the drift chambers. A special asymmetry analysis allowed positrons within $0.46\,\mu$s of the trigger time to be included, but this was not early enough to rule out a fast time component to the depolarisation within the muon stopping target. A muon spin relaxation (μ^+SR) experiment[33] was undertaken to look for relaxation at early times, in collaboration with CMMS[57]. μ^+SR can measure asymmetry as early as about 5 ns, since it uses fast scintillators to count decay positrons. A lack of additional depolarisation components down to 5 ns would improve confidence in the main experiment's results. At the time of writing, this is probably the shortest time scale that can be experimentally observed, and the author is unaware of any credible models for muon depolarisation within the first 5 ns in nonmagnetic metals.

μ^+SR experiments run at a higher rate ($20 - 40$ kHz compared to $2 - 5$ kHz), which should have allowed a high statistics measurement of the depolarisation rate comparable to the main TWIST experiment. However, this goal was not achieved due to the length of time required to calibrate the apparatus, and the dominant systematic uncertainty due to the fraction of muons stopping in the trigger scintillator, which will be described in detail later. Regardless, this is one of the highest precision experiments ever carried out using a μ^+SR apparatus.

H.2 Theory

Assuming standard model values of the muon decay parameters, and neglecting the electron mass and radiative corrections, Eqs. (1.14)-(1.19) can be used the write the muon

[57]TRIUMF Centre for Molecular & Materials Science

differential decay rate as

$$\frac{d^2\Gamma}{dx d\cos\theta} = \frac{G_F^2 m_\mu^5}{192\pi^3}\left[3 - 2x + P_\mu \cos\theta(2x-1)\right] x^2, \tag{H.1}$$

where θ is the angle between the positron momentum and muon spin, P_μ is the degree of muon polarisation, and $x \equiv E_e/m_\mu$. If decay positrons are detected for energy $(p < x < q)$ with 100% efficiency, then integration of Eq. (H.1) gives

$$\begin{aligned}\frac{d\Gamma}{d\cos\theta} = \int_p^q dx \frac{d^2\Gamma}{dx d\cos\theta} &= \frac{G_F^2 m_\mu^5}{192\pi^3}\left\{\left[\left(q^3 - p^3\right) - \frac{1}{2}\left(q^4 - p^4\right)\right]\right. \\ &\left. + P_\mu \cos\theta \left[\frac{1}{2}\left(q^4 - p^4\right) - \frac{1}{3}\left(q^3 - p^3\right)\right]\right\}, \end{aligned} \tag{H.2}$$

and if the angular fiducial is $(a < \cos\theta < b)$, then the normalised number of counts is given by

$$N = \frac{G_F^2 m_\mu^5}{192\pi^3}(b-a)\left[\left(q^3 - p^3\right) - \frac{1}{2}\left(q^4 - p^4\right)\right](1 + AP_\mu), \tag{H.3}$$

$$A = \tfrac{1}{2}(b+a)\frac{\tfrac{1}{2}\left(q^4 - p^4\right) - \tfrac{1}{3}\left(q^3 - p^3\right)}{\left(q^3 - p^3\right) - \tfrac{1}{2}\left(q^4 - p^4\right)}. \tag{H.4}$$

Assuming that there are counters in the forward (defined as the range $-1 < \cos\theta < 0$) and backward ($0 < \cos\theta < 1$) directions, and that each of these counters have their own values of $\{a, b, p, q\}$ and an associated background, then Eq. (H.3) can be used to show that

$$n_b(t) = b_b + N_0 e^{-t/\tau_\mu}\left[1 + A_b P_\mu(t)\right] \tag{H.5}$$

$$n_f(t) = b_f + rN_0 e^{-t/\tau_\mu}\left[1 + A_f P_\mu(t)\right], \tag{H.6}$$

where n_b and n_f are the number of backward and forward counts, b_b and b_f are the backgrounds in each counter, τ_μ is the muon lifetime, and N_0 is the number of muons at $t = 0$ multiplied by the factors in Eq. (H.3). An additional factor, r, is introduced to account for differences between the counters' energy and angular acceptance. The values of the empirical asymmetries A_b and A_f are expected to depend only on the counter geometry and the material that a positron passes through, and should remain constant if neither counters nor sample are moved, and the thickness of target material is unchanged.

If $P_\mu = P_\mu(0)f(t)$, where $P_\mu(0)$ is the polarisation at $t = 0$ and $f(t)$ is the time-dependent part, then Eqs. (H.5) and (H.6) cannot separate $P_\mu(0)$ from the A's unless r, A_b and A_f are known by other means, such as a Monte Carlo simulation. Therefore μ^+SR cannot measure $P_\mu(0)$ with high precision, and should only be used to determine $f(t)$, the time-dependent part of the polarisation.

H.3 μ^+SR Apparatus

H.3.1 M20 Beam line

Muons were delivered to the stopping target by the M20 beam line, which is a dedicated surface muon source also located at TRIUMF, in the same experimental hall as the main experiment's M13 beam line. The channel has been described elsewhere[130], and is similar to M13, with the addition of a DC separator element, which selects particles by velocity to remove most positrons originating from the production target. This is important for μSR since there is no event information, and a scattered beam positron can trigger the annular forward or backward counters in the same way as a decay positron. The DC separator had the disadvantage of introducing a significant transverse (vertical) polarisation component. M20 had a momentum acceptance that is larger than M13's and was not determined for this experiment. The beam spot was also larger than for M13, making collimation necessary, which will lead to a background from lower momentum slit-scattered particles.

H.3.2 Detector

The detector is shown schematically in Fig. H.1, where muons entered from the left and could stop in the 99.99% purity silver mask[58], the sample under investigation, or one of the trigger scintillators. This arrangement allowed the experiment to measure the sample and the silver mask simultaneously. The results for the silver mask were expected to remain constant for the experiment, allowing changes in the running conditions to be detected. Note that the forward and backward counters were approximately 0.6 cm thick. The backward counter was a disc of about 8.0 cm diameter, with a 2.5 cm hole for the

[58]This is less pure than the TWIST stopping target, which is $> 99.999\%$ pure.

muon beam. The target module was placed in the μSR HELIOS superconducting solenoid [131].

The data acquisition system (DAQ) recorded the positron time of arrival (relative to that of the muon) in two positron counter arrays, using 19200 channels of width 0.78125 ns. The experiment had an observable background made up of muon decays in or near the collimator, scattered positrons, and muon decays between the muon detector and the target. The background from muon pile-up (where a second muon enters the sample region during the same data gate as the original muon) was removed by the electronics, within 14 μs of either side of the event[59]. For this experiment, the positron gate was open for 1 μs *before* the muon trigger, so that events were recorded that could not possibly be caused by decay of the detected muon. The aim was to fix the background to the average "$t < 0$" value.

[59]'The removal of the background from muon pile-up is not perfect. With sufficiently high statistics, there is expected to be a visible small time dependence to the background.

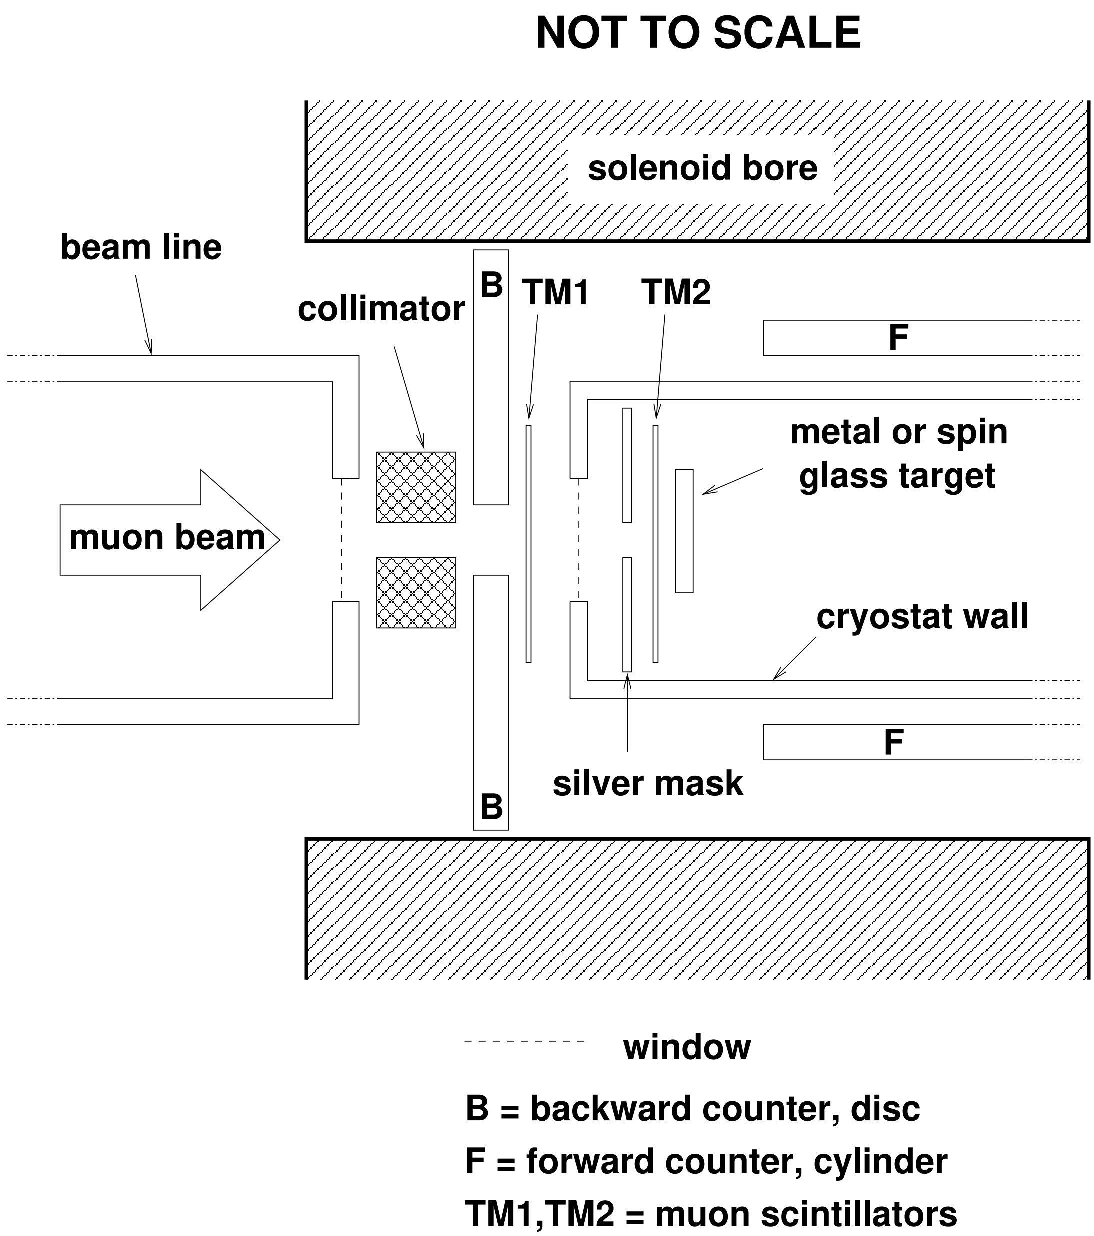

Figure H.1: Schematic for the experimental setup[36, 132]. The muons enter from the left and will trigger an event if they pass through the first scintillator (TM1) and stop in the silver mask, or pass through both scintillators (TM1 and TM2) and stop in the target. The backward positron counter (B) is protected from the muon beam by a collimator. Materials such as the vacuum tubes and light guides are not shown.

H.3.3 Sample preparation

Muons were stopped in a calibration sample, and then later in aluminium and silver samples. The calibration sample was a spin glass disc[60] of approximately 3.0 cm diameter and mass 4.5 g. The spin glass was known to rapidly and completely depolarise muons within several microseconds with a single exponential form; its analysis allowed the empirical asymmetries of the detectors to be determined. The aluminium and silver samples were extracted from the same foil as the TWIST targets, and they were stacked several layers thick[61], and clamped together with 3.6 μm thick Mylar. The layers were chosen such that muons had a negligible probability of passing through the foils. Note that there was no grease or glue used.

H.4 Data

The experiment collected data from 6 December 2006 to 20 December 2006, in runs that lasted between five and fourteen hours each. The conditions of TWIST were matched: the sample was immersed in a longitudinal 2 T magnetic field, and the temperature was not controlled[62]. The experiment started with runs that used a spin glass target. During these initial runs, the beam line was adjusted for rate and beam quality. A few days into the experiment, it was found necessary to keep the amount of material traversed by the decay positrons constant: a target of spin glass and metal foil was used back-to-back, and the target was then reversed to determine the metal depolarisation.

The data have been divided into sets, which are summarised in chronological order in Table H.1. There were unscheduled interruptions during data acquisition; the runs immediately after these interruptions were low rate and have been rejected. The bad runs have already been removed from the totals in the table.

The limited geometry of the decay positron counters (see Fig. H.1) resulted in a total positron rate that was about 10% of the muon rate. Approximately half of the beam time was spent on the nominal sets (E to H). The rate changed between the unscheduled interruptions, which are indicated by horizontal lines in the table. Set D was intended to

[60] The spin glass used in this experiment was a geometrically frustrated antiferromagnet, $Gd_2Ti_2O_7$. Condensed matter experimenters have studied the muon depolarisation in this material in order to determine the local magnetic fields [133], but have not undertaken precision tests.

[61] The aluminium target was ten layers thick, and the silver had six layers.

[62] Initially, there was a plan to test the validity of the Korringa relaxation model. This model predicts a relaxation rate, λ, which should be independent of magnetic field, but show an increase with temperature. However, due to the limited beam time this test was abandoned in favour of the primary TWIST aims.

have only a high DC separator setting, but it was also low rate.

A surface muon edge scan was not carried out, so the exact beam line momentum and resolution were undetermined.

Table H.1: The table shows the division of the data into sets, where a "set" is defined as a sum of individual runs. The horizontal lines indicate when data taking was suspended.

Data set	Target	Description	Duration (hours)	μ^+ rate Ag mask (kHz)	μ^+ rate sample (kHz)	μ^+ counts Ag mask ($\times 10^9$)	μ^+ counts sample ($\times 10^9$)
A	SG[a]	calibration	14.1	24.4	15.2	1.2	0.8
B	Al	nominal Al	8.0	25.0	15.6	0.7	0.5
C	Al	$\frac{1}{2}$-rate	15.7	12.4	9.6	0.7	0.5
D	Al	DC sep. high	5.7	16.6	10.7	0.3	0.2
E	SG+Al	nominal SG	24.0	19.4	13.3	1.7	1.2
F	Al+SG	nominal Al	40.1	19.9	13.7	2.9	2.0
G	Ag+SG	nominal Ag	32.4	21.8	14.6	2.5	1.7
H	SG+Ag	nominal SG	22.8	21.6	14.5	1.8	1.2
I	SG+Ag	thick scint.[b]	18.0	20.7	14.0	1.3	0.9
J	SG+Ag	thin scint.[c]	22.2	18.3	12.5	1.5	1.0
K	Ag+SG	thin scint.	13.3	20.1	13.8	1.0	0.7

[a] SG = Spin Glass.
[b] The thick scintillator was 508 μm, which is double the nominal thickness.
[c] The thin scintillator was 127 μm, which is half the nominal thickness.

H.5 Analysis

The aim of the analysis was to fit Eq. (H.5) to the backward counts and Eq. (H.6) to the forward counts simultaneously. The spin glass calibration sample was fit with $P_\mu(t) = \exp(-\lambda t)$, which allowed A_b and A_f to be determined, and these were fixed for the metal samples. All fitting was carried out using `MINUIT`[134].

H.5.1 Rebinning

The fit functions assumed that the background had no time dependence, and the muon beam had no transverse polarisation. In reality, protons arrived at the muon production target in bunches separated by the cyclotron period (43 ns), causing a periodic time dependence to the background. In addition, the transverse polarisation component precessed at the Larmor frequency ($\omega = g_\mu B = 272\,\mathrm{MHz}$). Since these effects were not included in the fit functions, the histograms of counts had to be carefully rebinned to minimise their effect.

The DAQ software performed a Fourier transform that included background subtraction, and found peaks at 23.06 MHz (43.37 ns) and 271.67 MHz (3.68 ns). This suggested that combining 113 channels (bin width 88.28 ns, about 2 cyclotron periods, 24 precession periods) was optimal. In practice, the rebinning was chosen by fitting a spin glass run with between 10 and 500 channels combined, and choosing a stable region. The final choice combined 115 bins (89.84 ns). The spin glass fit parameters were sufficiently robust to the choice of binning; this is demonstrated for the empirical backward asymmetry in Fig. H.2.

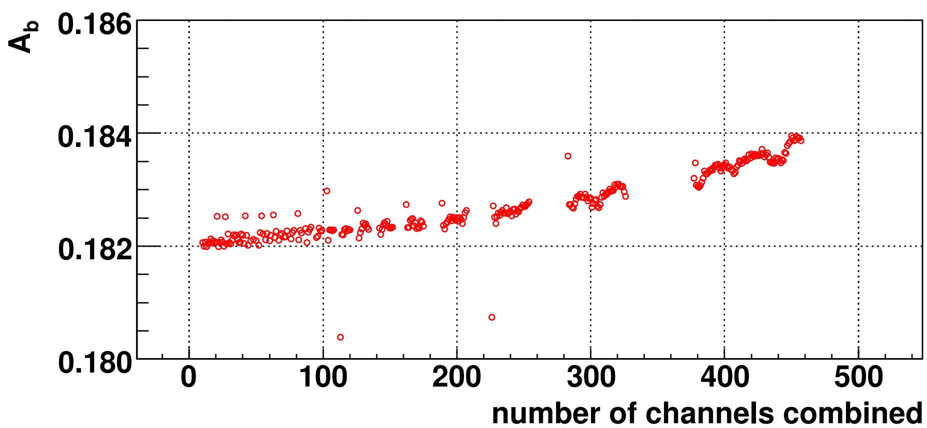

Figure H.2: Backward counter empirical asymmetry, A_b, from a single spin glass calibration run, with between 10 and 500 channels of width 0.78125 ns combined.

H.5.2 Calibration using spin glass

The spin glass (SG) data sets were fit assuming an exponential depolarisation, $P_\mu(t) = P_\mu(0)\exp(-\lambda_{\mathrm{SG}}t)$, and the results are shown in Table H.2. The A_b and A_f parameters depended on the target, which confirmed the importance of calibrating with a target of spin glass backed with metal, rather than spin glass alone. The r parameters varied by significant amounts between sets, but were consistent for the runs within sets. This suggested that the parameter was extremely sensitive to the geometry (*e.g.* target placement), and had to be left free when the metal samples were fit.

The results for a single exponential fit to the double-thickness scintillator were anomalous. In this case, a significant fraction of muons must be stopping and depolarising in the scintillator, and a single exponential model is inappropriate. The fit's confidence level was lower (14%), but not unreasonable, which indicated that the data had poor sensitivity to an extra depolarisation term for the scintillator. The results for the thin scintillator were consistent with the nominal case, suggesting the fraction of muons stopping in the nominal scintillator was small.

The background in each counter was investigated. When the backgrounds were free fit parameters, the values had a significant statistical discrepancy ($> 3\sigma$) with the result using the $1\,\mu$s of $t < 0$ counts. Since the backgrounds could be determined using the $t > 0$ decay data, and had a weak correlation with the other fit parameters, the $t < 0$ data were subsequently ignored.

Table H.2: Fit parameters for a single exponential fit to the spin glass sets. The bracketed number indicates the statistical uncertainty in the final digit.

Set	A	E	H	I	J
Sample	SG	SG+Al	SG+Ag	SG+Ag	SG+Ag
Scintillator thickness (μm)	254	254	254	508	127
χ^2/ndof	316.9/302 = 1.05	285.5/302 = 0.95	268.5/302 = 0.89	328.3/302 = 1.09	316.1/302 = 1.05
Confidence	0.27	0.74	0.92	0.14	0.28
r	0.7070 (5)	0.7019 (4)	0.6811 (3)	0.4764 (5)	0.6827 (4)
A_b	0.1858 (6)	0.1822 (5)	0.1833 (5)	0.0276 (5)	0.1836 (5)
A_f	-0.2409 (6)	-0.2365 (4)	-0.2404 (4)	-0.0522 (7)	-0.2415 (5)
$\lambda_{\mathrm{SG}}(\mu\mathrm{s}^{-1})$	0.878 (5)	0.880 (4)	0.878 (4)	0.55 (2)	0.883 (4)

Alternative models for $P_\mu(t)$ in the spin glass were tried, even though the reduced χ^2 values suggested there would be minimal sensitivity. A variable power law, $P_\mu(t) = \exp(-at^p)$, found $p = (1.005 \pm 0.007)$ for set E (SG+Al), with no change in the reduced χ^2. A sum of single exponentials was also tried, but the minimisation using `MINUIT` was unsuccessful for the nominal spin glass sets.

An effort was made to determine the depolarisation within the trigger scintillator, using the set taken with the double thickness ($508\,\mu$m) scintillator, where a significant fraction of muons stopped in the trigger. The first approach was iterative:

1. Fit $127\,\mu$m set assuming 100% muons stop in the sample, to determine λ_{SG}.

2. Fit 508 μm set allowing for stops in both the scintillator and sample using

$$P_\mu(t) = \epsilon \exp(-\lambda_{\text{scint}} t) + (1-\epsilon)\exp(-\lambda_{\text{SG}} t), \tag{H.7}$$

   where ϵ is the fraction of muons stopping in the trigger scintillator, λ_{scint} is the relaxation rate in the scintillator, and λ_{SG} was fixed to the 127 μm value from (1).

3. Fit 127 μm set with Eq. (H.7), with ϵ and λ_{scint} fixed to the values from (2).

4. Iterate (2)-(3) until convergence.

Only two iterations were required, and the results for the 127 μm and 508 μm sets are shown in Table H.3. There was a reduction in χ^2 when including a second exponential, and the high confidence level suggested that a single exponential model was valid for depolarisation in the scintillator. The addition of a scintillator depolarisation term helped r, A_b and A_f from Set I become more consistent with the other spin glass results. For the thin scintillator, the χ^2 was changed minimally, and the fraction of muons stopping in the scintillator was consistent with zero.

Table H.3: Comparison of fit parameters for a single exponential and a sum of exponentials, where the latter allows for separate depolarisation in the spin glass and scintillator. The bracketed number indicates the statistical uncertainty in the final digit.

Set	Set I: SG+Ag, 508 μm scint.		Set J: SG+Ag, 127 μm scint.	
Fit	Single λ	Converged	Single λ	Converged
χ^2/ndof	328.3/302 = 1.09	284.5/301 = 0.95	316.1/302 = 1.05	315.9/301 = 1.05
Confidence	0.14	0.74	0.28	0.27
r	0.4764 (5)	0.648 (1)	0.6827 (4)	0.69 (3)
A_b	0.0276 (5)	0.157 (1)	0.1836 (5)	0.19 (2)
A_f	-0.0522 (7)	-0.218 (2)	-0.2415 (5)	-0.25 (2)
λ_{SG} (μs^{-1})	0.55 (2)	0.886	0.883 (4)	0.886 (8)
λ_{scint} (μs^{-1})	-	0.0132 (8)	-	0.00132
ϵ (%)	-	86.0 (3)	-	4 (9)

The other spin glass sets were fit with an extra scintillator depolarisation term, with $\lambda_{\text{scint}} = 0.0132\,\mu$s fixed. The fit to Set A (SG only) would not succeed. For sets E (SG+Al) and H (SG+Ag), the fraction of muons stopping in the scintillator were determined as $\epsilon = (0 \pm 10)\%$ and $\epsilon = (8 \pm 8)\%$ respectively. Combined with the 127 μm result of

$\epsilon = (4 \pm 9)\%$, this suggests that somewhere between 0% and 10% of muons stopped in the nominal (254 μm) trigger scintillator. The effect of including a scintillator depolarisation term on $\lambda_{\rm SG}$, A_b and A_f is demonstrated in Fig. H.3, where the statistical precision becomes worse, but the central values do not systematically change.

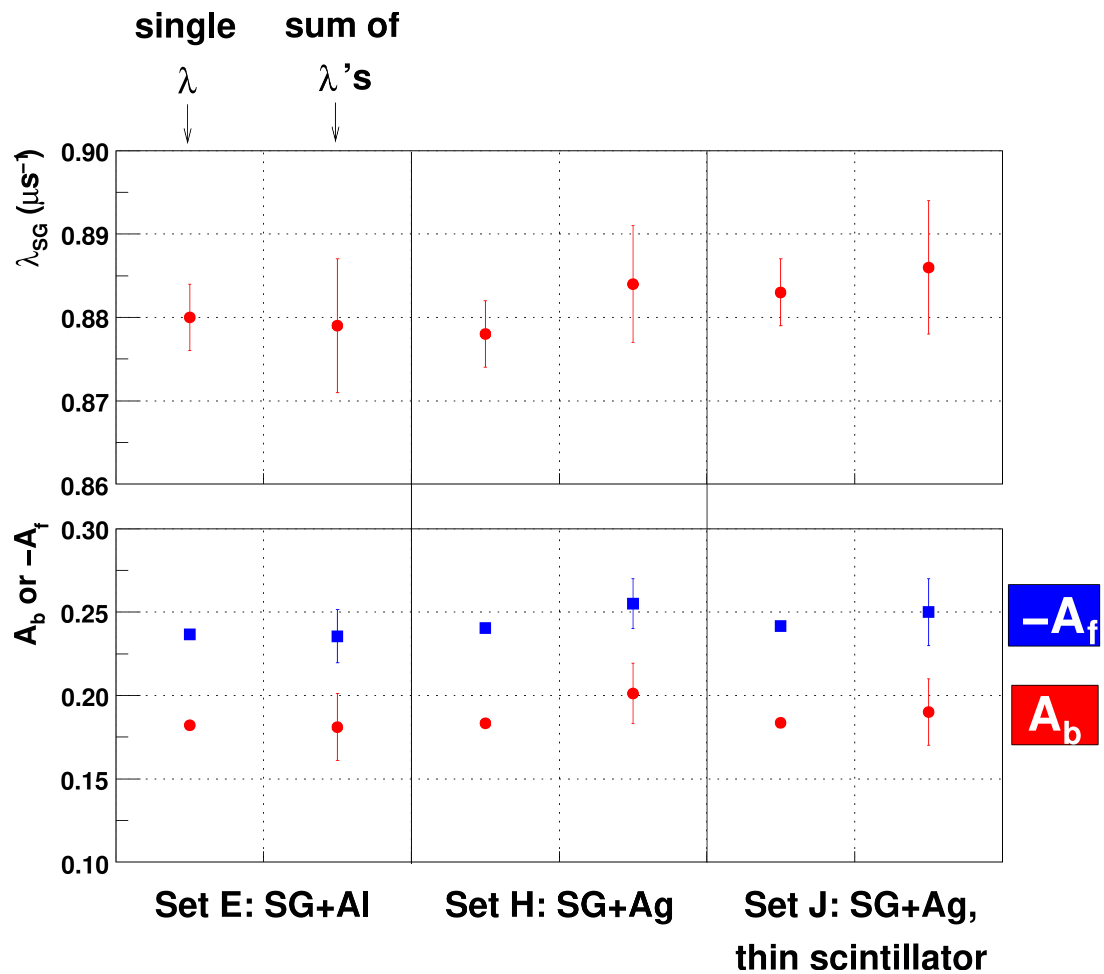

Figure H.3: Effect on $\lambda_{\rm SG}$, A_b and A_f when a second exponential scintillator depolarisation with $\lambda_{\rm scint} = 0.0132\,\mu$s is included.

In another approach, a simultaneous fit was made to all sets, with $\lambda_{\rm SG}$ and $\lambda_{\rm scint}$ as common parameters. The minimiser was able to converge on a solution, but the uncertainties on the fit parameters could not be reliably determined, and the `MINUIT` global correlations[63] were close to 1.0. The relaxation rates from this approach were

$$\lambda_{\mathrm{SG}} = (0.887 \pm 0.002)\,\mu\mathrm{s}^{-1}, \tag{H.8}$$

$$\lambda_{\mathrm{scint}} = (0.0123 \pm 0.0007)\,\mu\mathrm{s}^{-1}, \tag{H.9}$$

which are consistent with the earlier approach. This approach also determined the fraction of scintillator stops for the nominal case as somewhere between 0% and 10%.

In summary, the data suggest that about 86% of muons stop in the thick (508 μm) scintillator, but cannot precisely determine the fraction of muons stopping in the nominal scintillator. Later a simulation will be used to better determine the fraction.

H.5.3 Metal samples

The silver mask and metal foil data were analysed with the empirical asymmetries fixed[64] to $A_b = 0.185$ and $A_f = -0.238$, under the assumption that no muons stopped in the trigger scintillators. A run-by-run analysis of the silver mask was used as a consistency check, since Fig. H.1 showed that the mask is upstream of the regular target, and should be unaffected by any downstream changes. The r values and relaxation rates are shown in Fig. H.4, where r has clear systematic steps whenever the target or running conditions were altered. The relaxation rates are visually inconsistent between runs, and a weighted average of the results finds

$$\lambda = (1.07 \pm 0.07)\,\mathrm{ms}^{-1}, \tag{H.10}$$

with a confidence level of just 0.2%. Notably one run from set A is several σ from the average, set D has a negative relaxation rate (but a reasonable fit quality), and there is suggestion that set G has a systematically smaller rate than set H.

The results for the metal samples from TWIST are shown in Fig. H.5. The confidence levels are reasonable except for the last run of set F. For this set, the relaxation rate is still consistent with the other runs, so it was not rejected. The r parameters show a similar pattern to the silver mask's values in Fig. H.4. However, the relationship was not strong enough to impose a constraint. The relaxation rates for aluminium are statistically

[63]The correlation coefficient is is a number between -1 and $+1$ that describes the degree of correlation between the estimate of two fit parameters. The `MINUIT` *global* correlation for a parameter (α) "is a number between zero and one which gives the correlation between parameter [α] and that linear combination of all other parameters which is most strongly correlated with [α]"[134]. A global correlation coefficient close to one indicates there are too many free parameters.

[64]This is an approximation for the silver mask; it will not have the same empirical asymmetry as the metal target due to slightly different geometries.

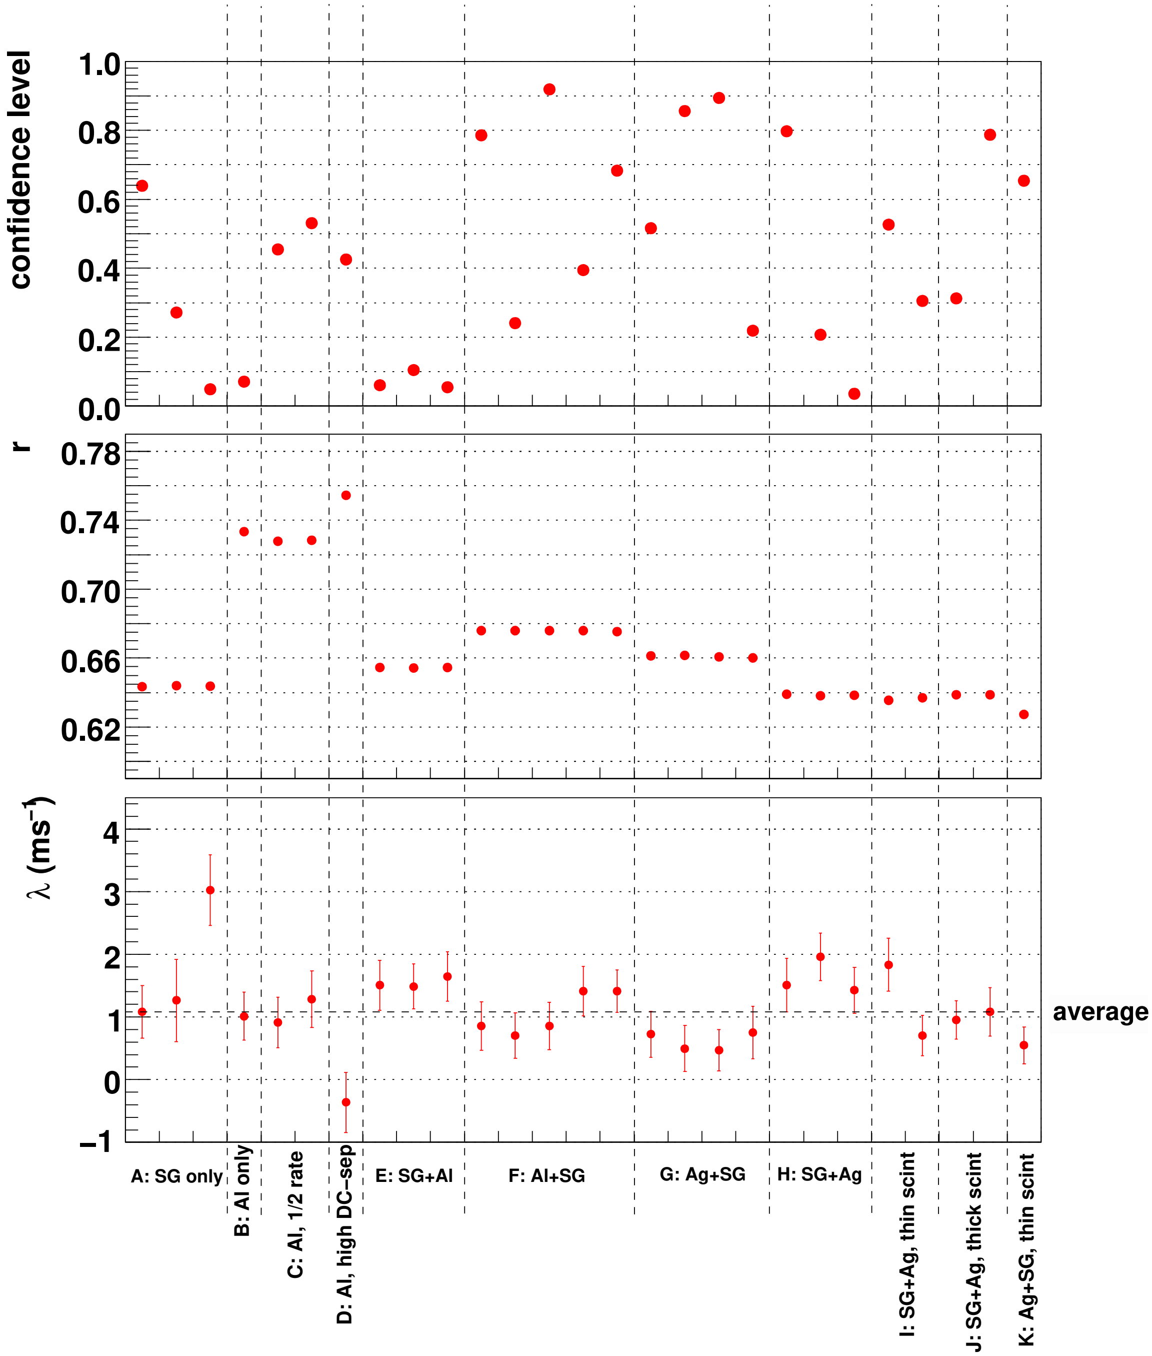

Figure H.4: Silver mask (99.99% purity) fit results. There is evidence that the relaxation values are inconsistent. The weighted average is $\lambda = (1.07 \pm 0.07)\,\mathrm{ms}^{-1}$.

consistent for sets B, C, D and F, with a weighted average of

$$\lambda_{\mathrm{Al}} = (1.7 \pm 0.2)\,\mathrm{ms}^{-1}, \tag{H.11}$$

under the assumption that all muons stop in the target. Similarly, the relaxation rates for silver are statistically consistent for sets G and K, yielding

$$\lambda_{\mathrm{Ag}} = (1.2 \pm 0.2)\,\mathrm{ms}^{-1}. \tag{H.12}$$

A single exponential fit is shown for an aluminium and silver run in Figs. H.6 and H.7. There is no evidence that anything beyond a single exponential is needed to describe the depolarisation. There is no observable fast depolarisation component below $1\,\mu$s. Note that fits were *not* made to the asymmetry; instead Eqs. (H.5) and (H.6) were used to fit the backward and forward histograms. In fact, the asymmetry figures cannot be constructed unless a model is assumed for $P_\mu(t)$.

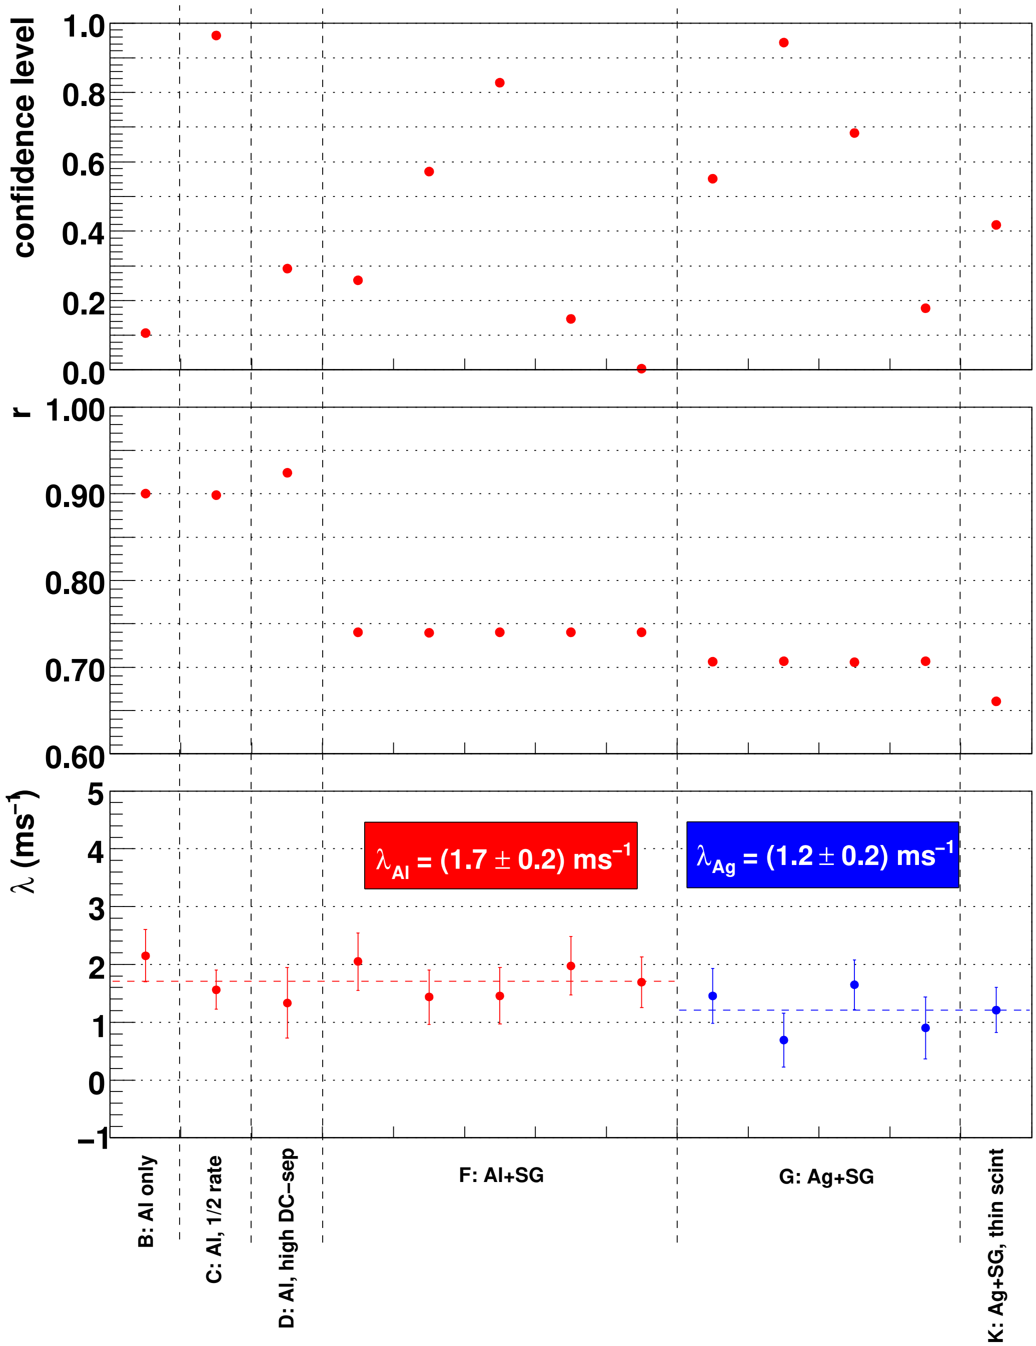

Figure H.5: TWIST metal samples (Al and Ag, both > 99.999% purity) fit results. For each foil, the sets taken under different running conditions are statistically consistent.

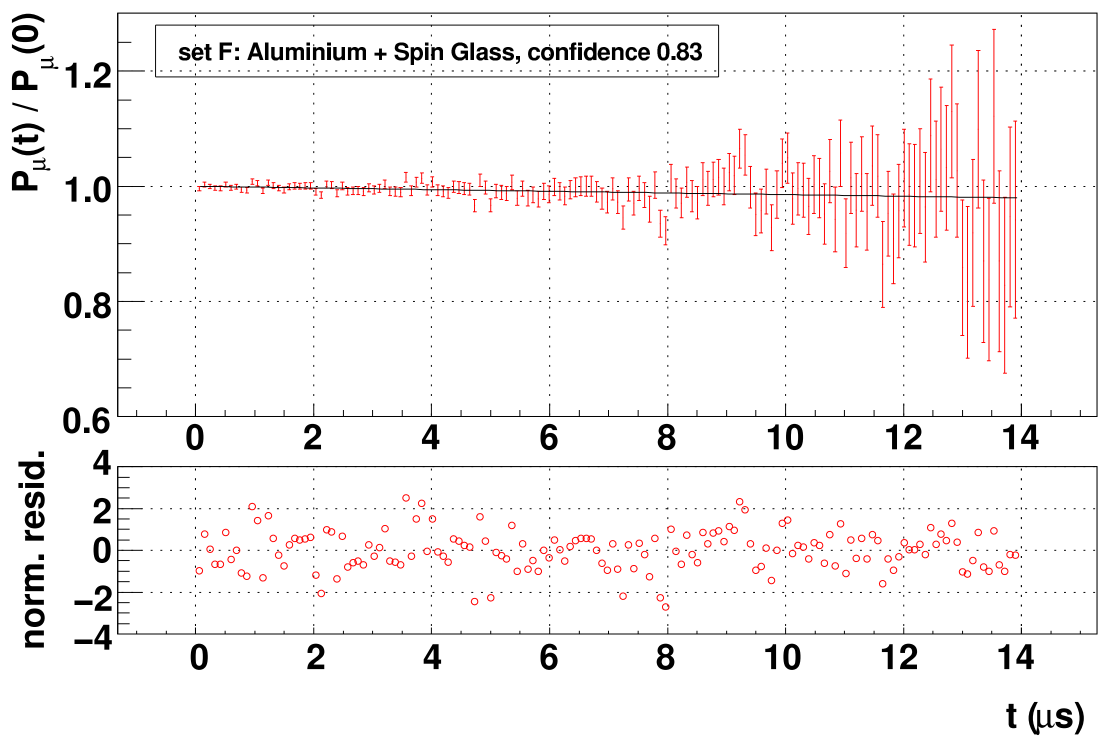

(a) Whole time range.

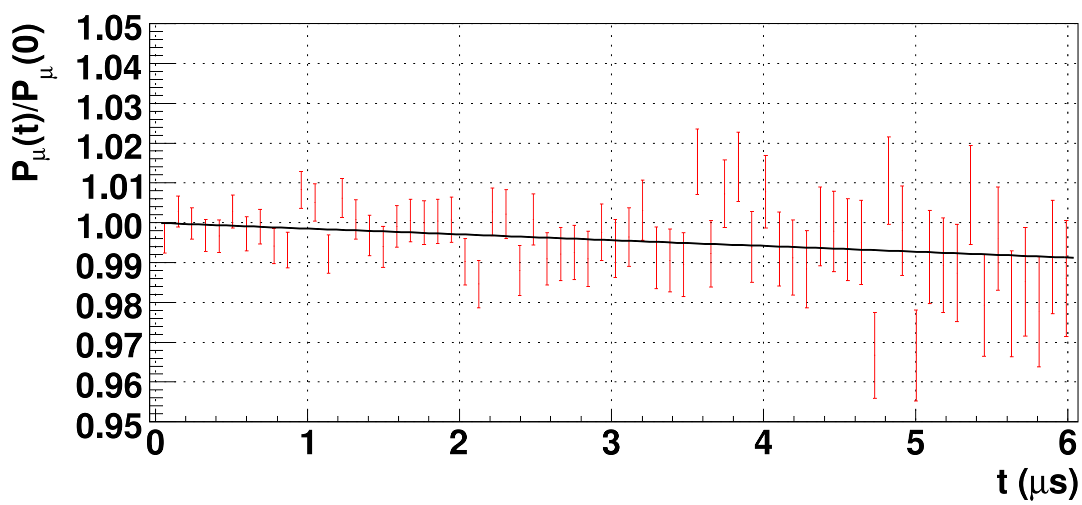

(b) Restricted to early times.

Figure H.6: Aluminium target: one run fit with $P_\mu(t) = P_\mu(0)\exp(-\lambda t)$.

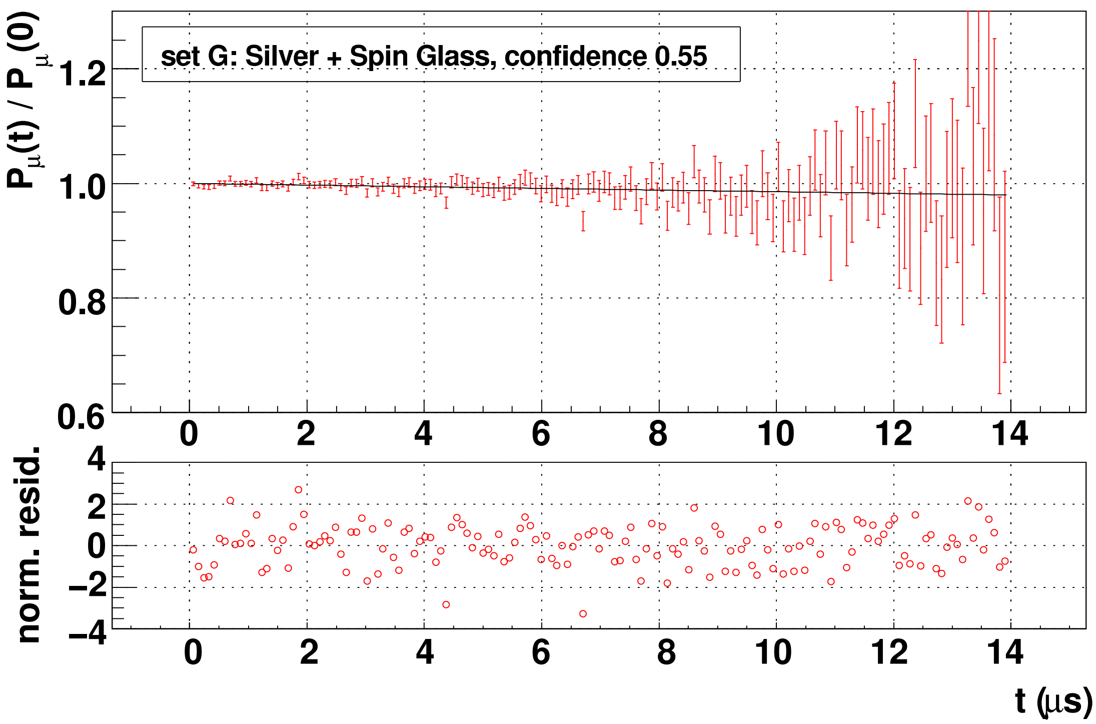

(a) Whole time range.

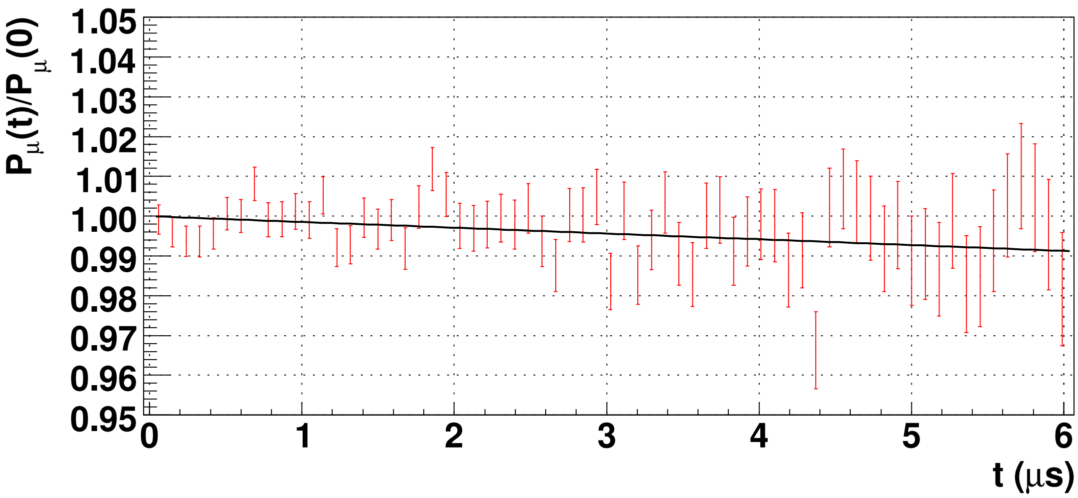

(b) Restricted to early times.

Figure H.7: Silver target: one run fit with $P_\mu(t) = P_\mu(0)\exp(-\lambda t)$.

H.6 Correction due to scintillator stops

The data indicated that between 0% and 10% of muons stopped in the nominal scintillator (see Section H.5.2). This was a problem, since the depolarisation rate in the scintillator was a factor ≈ 10 larger than the metal, allowing the possibility that all the depolarisation apparently observed in the metal samples was really due to depolarisation in the scintillator.

The fraction stopping in the scintillator was better determined using a simulation[100]. Software from the SRIM (the **S**topping and **R**ange of **I**ons in **M**atter) group of programs[135] was used, which included full quantum mechanical treatment of ion-atom collisions[136] (in this case the "ion" was the muon). The simulation required the initial momentum to be specified, but this was not measured during the data acquisition period. The M20 channel was set to maximise surface muon intensity, suggesting a minimal part of the momentum distribution lay beyond the kinematic cutoff at 29.79 MeV/c. Therefore simulations were run with a central momentum adjusted to be 2σ below the kinematic cutoff, with conservative estimates of the momentum resolution σ. The initial momentum distributions are shown in Fig. H.8(a).

Instead of running separate simulations for each scintillator, a single simulation was run at each momentum, with three pieces of scintillator in place: 127 μm, 127 μm and 254 μm. This allowed the fraction of stops in three thicknesses of scintillator (127 μm, 254 μm and 508 μm) to be determined simultaneously by integrating the number of stops appropriately. In each simulation, the stopping distribution was adjusted[65] so that the total fraction of scintillator stops was about 86%; this matched the precise result from Table H.3 that 86.0% of muons will stop in 508 μm of scintillator.

The stopping distributions from the simulation are shown in Fig. H.8(b). The number of muons stopping in each scintillator thickness are integrated in Table H.4, which shows that it is possible to obtain anything from 0.5% to 5.9% stopping in the nominal scintillator; this disfavours the data's suggestion of up to 9%. Clearly the initial momentum distribution is important in such a simulation, and it would have been valuable to make a measurement while acquiring data. In the thin (127 μm) scintillator, the simulation found that no more than 0.2% of muons stopped in the trigger.

A systematic uncertainty can be estimated by repeating the analysis with a fraction of the muons forced to stop and depolarise in the scintillator. The results for between 0% and 10% of the muons stopping in scintillator are shown in Fig. H.9. If the central value

[65] The simulation included a piece of variable thickness Mylar upstream of the trigger scintillators.

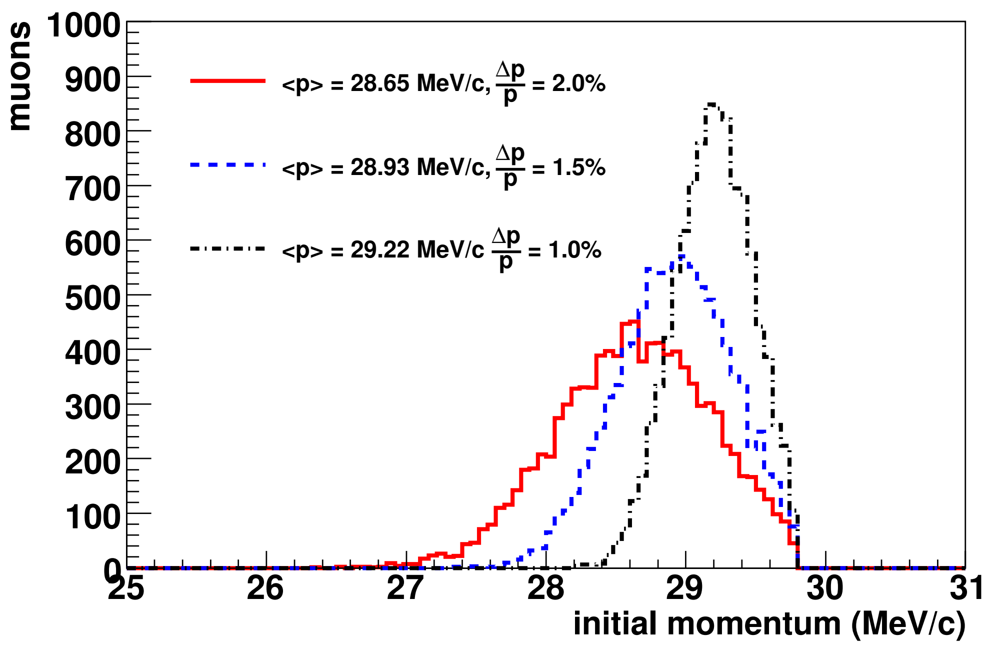

(a) Initial momentum used in the simulations.

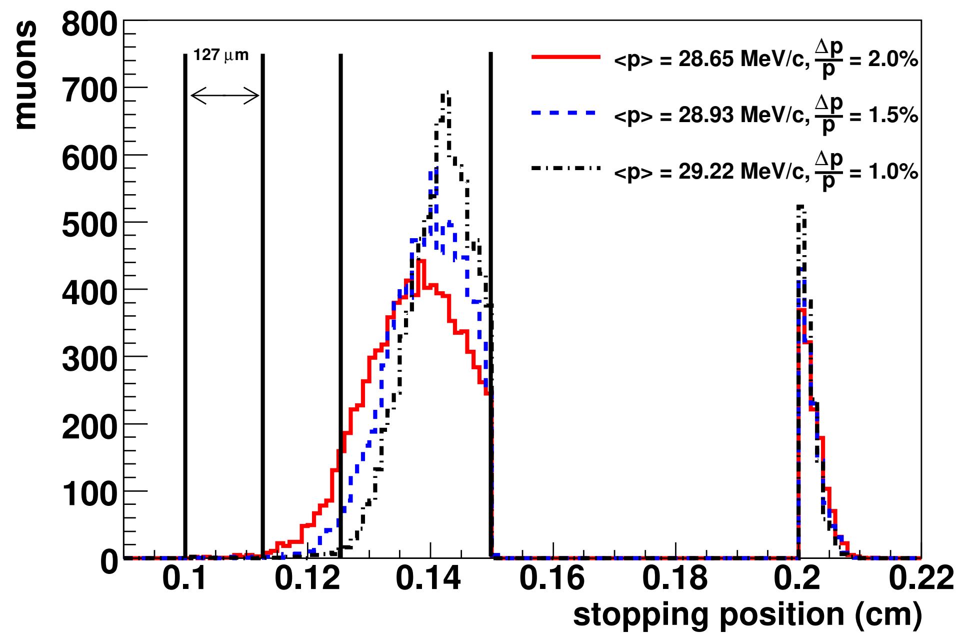

(b) Stopping distributions for the different initial momentums.

Figure H.8: Initial momentum distributions used in the SRIM simulation, and the resulting stopping distributions.

Table H.4: The fraction of muons stopping in each scintillator, for three different initial momentum distributions.

$< p >$ (MeV/c)	$\Delta p/p$ (%)	Fraction of stops (%) thin (127 μm)	nominal (254 μm)	thick (508 μm)
29.22	1.0	0.07 ± 0.03	0.5 ± 0.1	85.8 ± 0.9
28.93	1.5	0.09 ± 0.03	1.9 ± 0.1	86.8 ± 0.9
28.65	2.0	0.22 ± 0.05	5.9 ± 0.2	86.7 ± 0.9

of λ is corrected so that $\epsilon = 3.5\%$, with an uncertainty that allows for between 1% and 6% scintillator stops, then the nominal scintillator results become

$$\lambda_{\mathrm{Al}} = (1.3 \pm 0.2\,(\mathrm{stat.}) \pm 0.3\,(\mathrm{syst.}))\,\mathrm{ms}^{-1} \quad \text{(H.13)}$$

$$\lambda_{\mathrm{Ag}} = (0.9 \pm 0.2\,(\mathrm{stat.}) \pm 0.2\,(\mathrm{syst.}))\,\mathrm{ms}^{-1}. \quad \text{(H.14)}$$

For the single run with the thin (127 μm) scintillator, the contribution from scintillator stops is negligible, but the statistical uncertainty dominates so that

$$\lambda_{\mathrm{Ag}} = (1.2 \pm 0.4\,(\mathrm{stat.}))\,\mathrm{ms}^{-1}. \quad \text{(H.15)}$$

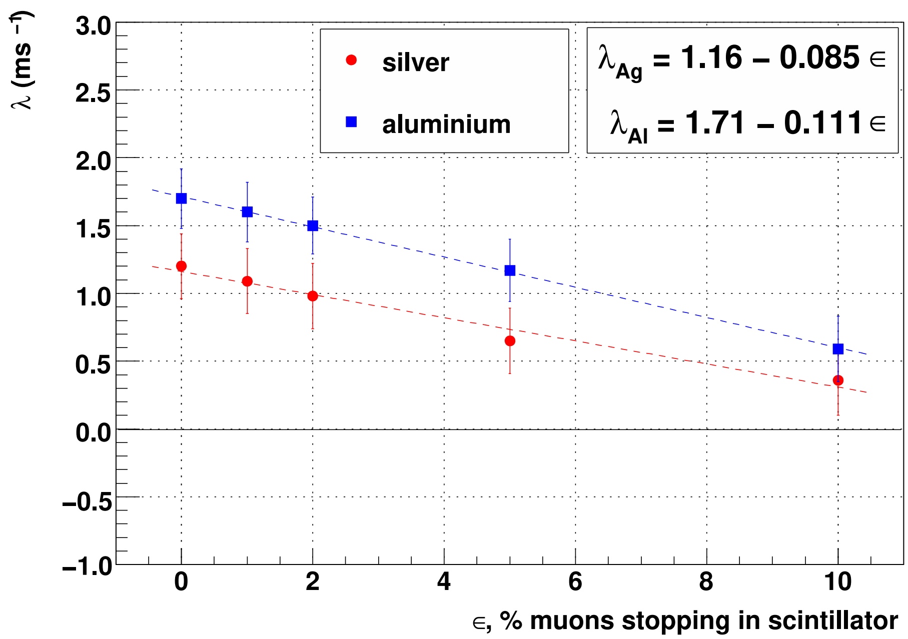

Figure H.9: Exponential depolarisation rate for silver and aluminium, for different fraction of muons stopping in the scintillator. The parameter ϵ represents the percentage of muons stopping in the scintillator.

H.7 Other systematic uncertainties

There is a systematic uncertainty from the fixed A_b and A_f values that were used in the metal fits. A 10% change in either of these parameters altered the relaxation rate by at most $0.12\,\mathrm{ms}^{-1}$. A_b and A_f were determined with statistical precision $< 0.3\%$, but were forced to be $A_b = 0.185$ and $A_f = -0.238$ for all fits, which meant they were wrong by up to 2% (see Table H.2). An estimate of the systematic uncertainty on the relaxation rate is $(2/10) \times 0.12 = 0.02\,\mathrm{ms}^{-1}$. This is negligible compared to the systematic uncertainty from the scintillator stops.

There was no systematic uncertainty assigned due to the background having a time structure. Figure H.5 showed that the confidence levels were reasonable, except for one set, which still produced a consistent relaxation rate. If the confidence levels were lower, or the relaxation rate for the low confidence level run had been anomalous, a systematic uncertainty due to the time structure of the background may have been necessary[66].

The sets with different muon rates and DC separator settings showed no evidence of an inconsistent relaxation rate (see Fig. H.5). Therefore no systematic uncertainties are assigned for these effects.

H.8 Conclusions

The μ^+SR experiment demonstrated that down to a decay time of 10 ns, a single exponential function is appropriate and sufficient to describe the depolarisation in aluminium and silver. This is an essential contribution to the TWIST result.

The analysis found the following relaxation rates for the TWIST target foils:

$$\lambda_{\mathrm{Al}} = (1.3 \pm 0.2\,(\mathrm{stat.}) \pm 0.3\,(\mathrm{syst.}))\,\mathrm{ms}^{-1} \tag{H.16}$$

$$\lambda_{\mathrm{Ag}} = (0.9 \pm 0.2\,(\mathrm{stat.}) \pm 0.2\,(\mathrm{syst.}))\,\mathrm{ms}^{-1}. \tag{H.17}$$

These are consistent with the results from the TWIST detector, but their uncertainties are not competitive. The dominant systematic uncertainty could not be reduced without accumulating more data.

[66] If the background was the leading systematic uncertainty, then a further investigation could revisit the calculations presented in Appendix 3 of Ref. [137], where time dependences due to pile-up are discussed.

H.9 Future experiments

Several recommendations can be made for a future experimenter. The fraction of muons stopping in the trigger scintillator could have been determined with greater certainty if the beam line had been momentum calibrated. A stopping target of pure scintillator would have provided a clean measurement of the depolarisation rate in the trigger. The thin (127 μm) scintillator provided signals with sufficient amplitude, and could have been used to accumulate all the nominal data.

The fit quality could have been improved by including a time dependence to the background, but this would only be appropriate if more statistics were acquired, which would take far more than two weeks of beam time at TRIUMF. We were unable to find a constraint using the silver mask's fit parameters; for example, the r parameter from the silver mask and regular metal fits were highly correlated, and it may have been possible to constrain them. A detailed simulation of the apparatus may have allowed such constraints to be investigated.

Appendix I

The M13 beam line and $(g-2)$

The muons in the simulation are started at the end of the M13 beam line, with their spin and momentum anti-parallel. A small depolarisation takes place through the beam line, since there is a difference in the deflection of the momentum and spin through the magnetic fields of the dipoles. Recall from Eq. (1.38) that the precession frequencies for the momentum (ω_p) and spin (ω_s) differ by

$$a = \frac{\omega_p - \omega_s}{\omega_p} = \frac{g-2}{2} = 1.17 \times 10^{-3}. \tag{I.1}$$

For an on-axis particle, shown in Fig. I.1, the deviation from the initial momentum direction $\vec{v_0}$ will be equal in B1 and B2, but with opposite sign in each dipole. The same is true for the spin. Therefore the momentum and spin both have a net deflection of zero, and they remain anti-parallel.

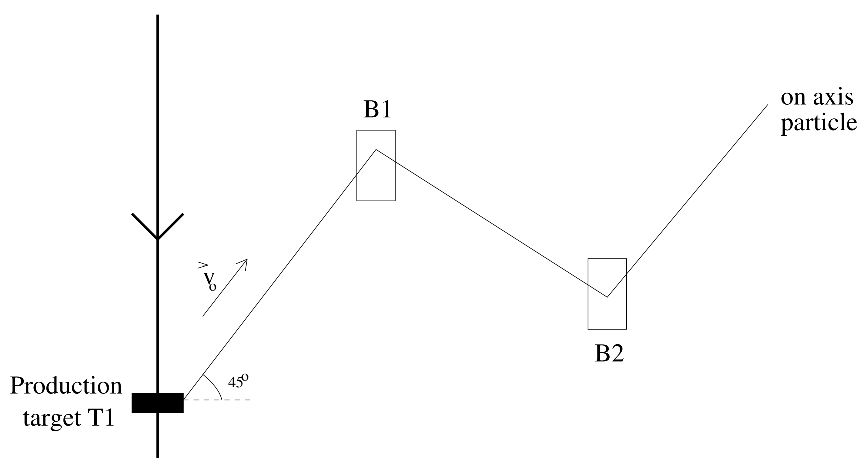

Figure I.1: An on-axis particle passing through M13. The dipole magnets are B1 and B2.

Now consider the two most extreme trajectories that can pass through M13, assuming two different dipole configurations. These are shown in Fig. I.2. The real configuration is the single crossover in Fig. I.2(a)[138]. An upper limit can be set on the total net angular

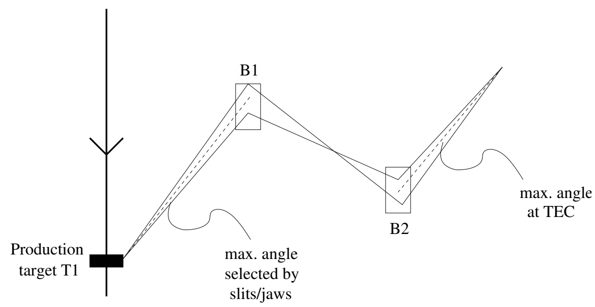

(a) Real dipole configuration for M13.

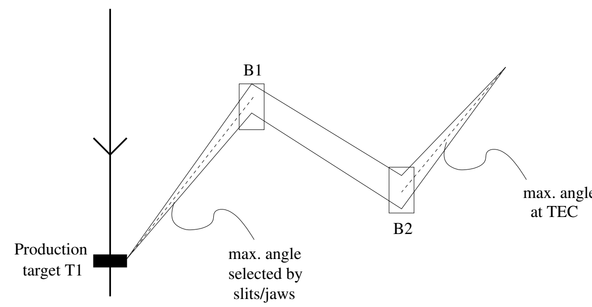

(b) Alternate dipole configuration for M13.

Figure I.2: Most extreme trajectories through M13 beam line, under two configurations.

deflection using the emittance at the production target ($\lesssim$ 250 mrad, before focusing by Q1 and Q2) and the emittance measured by the TECs ($\lesssim$ 150 mrad). Under the configuration in Fig. I.2(a), the largest possible deflection is $\sim$ 100 mrad, which from Eq. (I.1) changes the angle between the momentum and spin by $\Delta\theta \sim 1 \times 10^{-4}$. This contributes a depolarisation $\sim 1 - \cos\Delta\theta \sim 10^{-8}$. Even if the beam line were configured according to Fig. I.2(b), the depolarisation would be $\sim 10^{-7}$. Both of these depolarisation estimates are completely negligible.

Appendix J

Estimating muon stops in PC6 gas

A cut is placed on the muon's pulse width in PC5 and PC6 that removes muons stopping in the PC6 gas. The positions of this cut are shown in Fig. J.1(a); zone 1 is selected for the normal analysis. The cut is imperfect, and allows a small contamination of gas stops into zone 1. The main text uses three numbers to determine the systematic uncertainty from this contamination: the depolarisation in the gas, the fraction of total PC6 stops that are in gas, and the fraction of the gas distribution that leaks across the cut-B line. The method of determining these numbers will now be described.

A special analysis selected zone 3 in order to determine the depolarisation in the chamber gas. The intercept in the asymmetry fits was statistically consistent with zone 1, demonstrating that no measurable gas depolarisation[67] occurs before $1.05\,\mu$s. The relaxation rate is $\lambda_{\text{gas}} = (28 \pm 3)\,\text{ms}^{-1}$, which is significantly larger than the metal. Using Eq. (6.8), the average gas depolarisation is then $\langle P_\mu \rangle\,|_{\text{gas}} = 8.0\%$ $(\lambda_1 = 0, \lambda_2 = 30.0\,\text{ms}^{-1})$.

Each data set was re-analysed with *no cut* on the PC5/6 pulse width, and the resulting change in $P_\mu^\pi\,\xi$ is divided by the gas depolarisation (8.0%) to estimate the ratio of gas to metal target stops in PC6. For 2006, between 4% and 5% of the total stops are gas, with the exception of set 68 which has 7% gas stops since the stopping distribution was further upstream. For 2007, between 6.5% and 7.5% of the muons stop in the gas, depending on the set.

The fraction of the gas distribution leaking between zones 1 and 2 is estimated using the contamination of zone 3 into 4 (a clean separation of gas and metal is not possible in zone 2). Figure J.1(b) is a projection of zones 3 and 4, rotated to make the cut-B line vertical. The distribution is fit with the sum of two Gaussian distributions. The fraction of the gas distribution leaking into zone 4 was tuned to be less than 0.5%, to ensure a negligible systematic uncertainty.

Note that the simulation was unable to estimate the fraction of gas stops in zone 1

[67] However, if muons are selected that stop in PC5, PC7 and PC8, a change in the intercept of about 3% is observed, indicating a "fast" depolarisation. The simulation predicts that for these selections, 50% of muons stop in the gas and 50% in the Mylar foil. Since we do not observe a fast depolarisation when pure gas stops are selected in PC6, we conclude that the fast depolarisation is within the Mylar.

since it was not tuned to reproduce the PC chamber response.

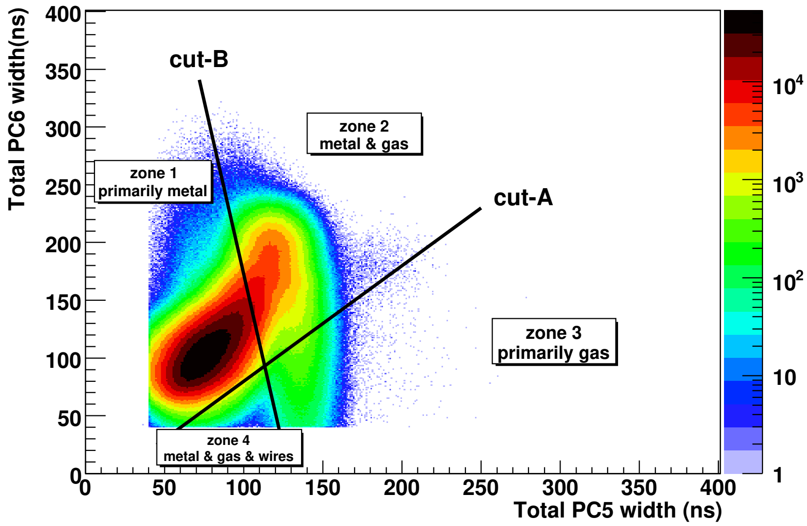

(a) Zone definitions for the muon pulse width in PC5 and PC6.

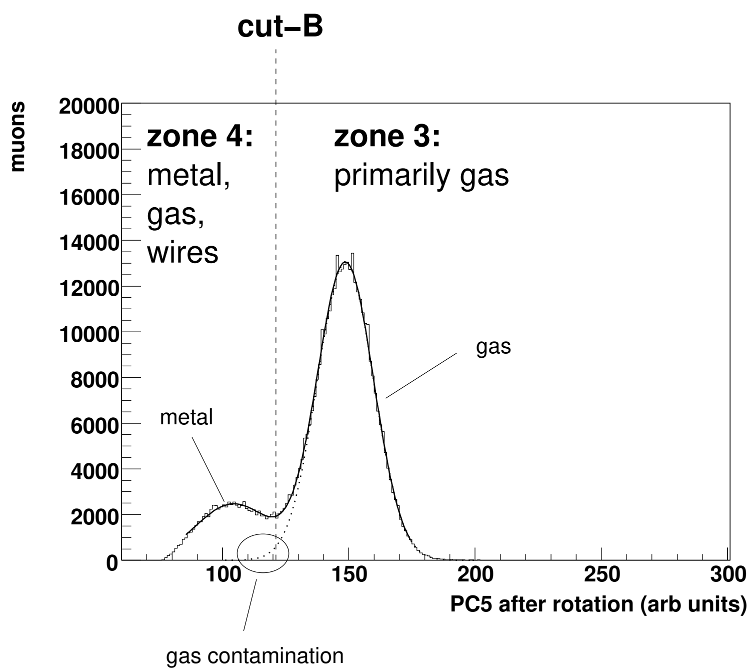

(b) Projection of zones 3 and 4, perpendicular to cut-B. A fit has been made to the sum of two Gaussian distributions.

Figure J.1: Estimation of gas contamination after a cut on the PC5 and PC6 muon pulse width.