

High Precision Track Reconstruction and First
Emittance Measurements in the MICE Step IV
Cooling Channel

Christopher J Hunt

Supervised by Jaroslaw Pasternak

PHYSICS DEPARTMENT
IMPERIAL COLLEGE LONDON

31ST AUGUST 2016

A THESIS SUBMITTED AS
REQUIRED FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY.

Acknowledgements

I was extremely fortunate to find myself in the company of the MICE Collaboration. A truly brilliant and dedicated group of people, who showed an unending determination that no research council or funding body could deter! I will forever be grateful to them for showing me the true heart of scientific research.

Special mention goes to those few people who had the greatest impact these last few years:

- My supervisor, Jaroslaw, for the unwavering support and faith in my abilities. He provided support and guidance, but allowed me to follow my own strengths, and grow from a student into a peer.
- My close colleagues, Adam, Melissa, Victoria, Ken, and the many Chris', for making sure that I always had something to work on! But also for helping when there was too much.
- And of course, my family for pushing me in the right direction. As reluctant as I may have seemed at first.

Thank you all.

Declaration of Originality

The content of Chapters 1 and 2 reviews the existing literature surrounding muon accelerators, the MICE Experiment and the theoretical framework used throughout this document.

Chapter 3, describes the Kalman Filter based track fit, as documented by the existing literature. The algorithm and its application to the MICE trackers was developed by fellow collaborators in the MICE Collaboration, the actual implementation within the analysis framework was my own work.

Chapters 4 to 7 describe the analysis tasks that were conducted personally. The raw data and reconstruction was due to the efforts of the MICE collaboration and at all stages there were discussions with fellow collaborators regarding the concepts and methods used. The analysis programs that were used to generate the results were developed personally.

The analysis in chapter 7 was conducted in parallel with an official MICE Collaboration paper. The systematics studies were conducted without assistance, however there were discussions regarding the direction and methodologies of the general analysis task. The content of the paper will differ in complexity and scope, despite sharing a similar goal.

Copyright Declaration

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Abstract

The International Muon Ionisation Cooling Experiment (MICE) is currently being commissioned at the Rutherford Appleton Laboratory (RAL) and aims to be the first experiment to demonstrate the technique of muon beam ionisation cooling, a key component in Neutrino Factory and Muon Collider designs. In order to do so, high precision muon track and beam phase-space reconstruction is required.

A Kalman Filter based track fit has been designed, implemented and tested for the reconstruction of both straight and helical muon tracks, using data from the MICE Scintillating Fibre Trackers. Detailed Monte Carlo simulation studies have been conducted in order to verify the implementation and predict the performance of the fitting algorithms when applied to real data. A combination of simulation studies and statistical data analyses have been used to commission the track reconstruction software and perform simple validation techniques.

The straight track reconstruction was applied to an alignment study, designed to measure the relative positions of the two tracking detectors within the MICE cooling channel. A translational resolution of 0.11 mm in both the x and y directions was extracted, in addition a rotational resolution of 0.05 mrad of the relative yaw, Φ_x , and pitch, Φ_y , between the two trackers was calculated.

The helical track reconstruction was successfully used in the analysis of the first helical track data acquired within the MICE Cooling Channel. A momentum window of 195 to 205 MeV/ c was applied to the default MICE muon beam, which resulted in 5049 reconstructed tracks with a normalised transverse emittance of $5.08 \pm 0.05_{\text{Stat}} \pm 0.06_{\text{Sys}} \pm 0.02_{\text{Field}}$ mm. The systematic error was calculated based on the modelled momentum residuals, and a second systematic error, due to the incomplete knowledge of the true field, was estimated through comparisons of different field maps.

Contents

Acknowledgements	2
Declarations	3
Contents	5
List of Tables	8
List of Figures	10
1 Introduction	18
1.1 Muon Acceleration	18
1.2 Beam Optics and Emittance	20
1.2.1 Phase-Space Description	21
1.2.2 Linear Beam Optics	24
1.2.3 Solenoid Optics	25
1.2.4 Liouville's Theorem	26
1.2.5 Matching a Beam	28
1.3 Ionisation Cooling	29
1.3.1 Principles of Ionisation Cooling	30
1.3.2 The Cooling Equation	31
1.3.3 Building a Muon Cooling Channel	33
2 MICE	35
2.1 Introduction	36
2.2 The Muon Beamline	36
2.3 Detector Systems	38
2.4 MICE Cooling Channel	40
2.4.1 Magnetic Layout	41
2.4.2 The SciFi Trackers	44

2.5	Conclusion	46
3	MICE Track Reconstruction	48
3.1	Introduction	48
3.2	Geometry	48
3.3	Track Selection	50
3.3.1	Digitisation	50
3.3.2	Clustering	50
3.3.3	Space Point Production	51
3.3.4	Pattern Recognition	52
3.4	Kalman Filtering	53
3.4.1	Conceptual Description	54
3.4.2	The Algorithm	55
3.4.3	The Extended Kalman Filter	58
3.4.4	MAUS Implementation	59
3.5	Conclusions	64
4	Validation of Track Fitting Routines	66
4.1	Introduction	66
4.1.1	Chi-Squared Distributions	66
4.1.2	P-Value Distributions	67
4.2	Algorithm Validation	68
4.2.1	Energy Loss	68
4.2.2	Measurement Effects	70
4.2.3	Stochastic Effects	70
4.2.4	Full Reconstruction	72
4.3	Monte Carlo Simulation Performance	72
4.3.1	Transverse Reconstruction	73
4.3.2	Momentum Reconstruction	75
4.3.3	Track Finding Efficiency	77
4.4	Conclusions	78
5	Validation of Reconstructed Data	82
5.1	Introduction	82
5.2	Beam Selection	83
5.3	Low Level Reconstruction	85
5.3.1	Digit Production	85
5.3.2	Cluster Production	87
5.3.3	Spacepoint Production	87

5.3.4	Pattern Recognition	91
5.4	High Level Reconstruction	92
5.4.1	Straight Tracks	92
5.4.2	Helical Tracks	96
5.5	Conclusions	99
6	Alignment of the MICE Trackers Using Data	101
6.1	Introduction	101
6.2	Alignment Method	102
6.2.1	Outline	102
6.2.2	Assumptions	104
6.2.3	Procedure	104
6.3	Monte Carlo Validation	107
6.4	TOF Extrapolation	109
6.5	Axial Rotation	110
6.6	Results	112
6.7	Conclusions	113
7	Measurement of the MICE Muon Beam Emittance	115
7.1	Experiment Configuration	115
7.2	Procedure	116
7.3	Error Analysis	119
7.3.1	Statistical Errors	119
7.3.2	Covariance Matrix Corrections	121
7.3.3	Field Alignment Correction	122
7.3.4	Field Uniformity Correction	125
7.3.5	Combined Correction Terms	127
7.4	Monte Carlo Analysis	129
7.5	Reconstruction Results	131
7.6	Conclusion	134
8	Conclusions	136
A	Pattern Recognition in MAUS	145
A.1	The Method of Least Squares	145
A.2	Straight Track Fitting	147
A.3	Helical Track Fitting	147
B	Reconstruction Validation Data	150

List of Tables

2.1	Description of the magnetic lattice used to perform the simple Step IV Monte Carlo Studies shown in figures 2.7 and 2.8.	42
3.1	Summary of survey measurements of the tracker stations made during the construction of the trackers. The centres of the first and last stations were used to define the coordinate system. . . .	49
3.2	Summary of geometry of each plane within any given station. z is the longitudinal position with respect to the centre of the station and θ is the rotation of the plane around the tracker's central axis. Plane 2' refers to the anomalous plane in the downstream tracker, station 5.	50
4.1	A summary of the reconstruction parameters produced from the straight track Kalman fit. The values are averaged across uniform phase-space which spans 140 to 260 MeV/c in p_z and 0 to 20 MeV/c in p_\perp	81
4.2	A summary of the reconstruction parameters produced from the helical track Kalman fit. The values are averaged across uniform phase-space which spans 140 to 260 MeV/c in p_z and 0 to 180 MeV/c in p_\perp	81
4.3	A summary of the correction functions required to optimize the results from the Kalman Track Fit.	81
5.1	List of the runs that were used in this and subsequent chapters. The momentum value refers to the nominal beam momentum at the centre of the cooling channel, and the beamline value refers to the nominal beam emittance at the entrance to the cooling channel. These two values allow for the correct beamline settings to be identified from a predefined catalogue.	83
5.2	Comparison of single-digit clusters to double digit clusters. . . .	88

6.1	A summary of the transverse alignment measurements for each run in the two Datasets that were analysed. The combined alignment is calculated from the weighted mean of all runs.	112
6.2	A summary of the relative alignment measurements for each of the trackers and TOF detectors. The upstream tracker defines the origin of the coordinate system, and hence has no relative alignment.	113
7.1	The results of the linear fits applied to the total momentum residuals for the four different field models that were analysed.	126
7.2	Reconstruction results of the key optical parameters as calculated using the official Monte Carlo simulation. 4669 muons were included in the calculation following the beam selection. A statistical sampling error due to the number of muons selected and systematic error due to the momentum bias are included. The true Monte Carlo values are included for comparison.	131
7.3	Reconstruction results of the key optical parameters as calculated from Dataset C. 5049 muons were included in the calculation following a momentum selection and p-value cut. A statistical sampling error due to the number of muons selected and systematic error due to the momentum bias are included. An additional systematic error for the uncertainty in the field map is not included.	133

List of Figures

1.1	Block diagram describing the overall construction of a Neutrino Factory (NuMAX) and a Muon Collider. A visualization including which components can be easily transferred from one facility to another is included.	19
1.2	The relations between the Twiss parameters in the context of the RMS beam envelope.	22
1.3	The effect of some small drift on the x - x' phase-space distribution for a beam in arbitrary units. The particles with larger gradient have moved further from the centre of the position axis, while the distribution of gradients has remained constant.	24
1.4	An illustration of the filamentation phenomenon and how over time it may lead to a non-linear emittance growth despite Liouville's theorem.	29
1.5	Diagrammatic representation of the process of ionisation cooling in momentum space. The green arrow represents the effect of the absorber - a net reduction in total momentum; the red arrow represents a smearing of the momentum spread due to multiple Coulomb scattering; and the blue arrow represents the final momentum following some accelerating structure. The net effect is a reduction in momentum spread, while maintaining the mean beam momentum.	30
1.6	Representation of a FOFO-type cooling lattice. The absorbers (blue) are placed in regions of low betatron function to enhance the cooling performance. The RF accelerating cavities are placed in between to replenish the longitudinal momentum components. The relative field flip between pairs of solenoids prevents the accumulation of canonical angular momentum.	34

1.7	Representation of a Super-FOFO-type cooling lattice. The absorbers (blue) are placed in regions of low betatron function to enhance the cooling performance. The RF accelerating cavities are placed in between to replenish the longitudinal momentum components. The relative field flip between pairs of solenoids prevents the accumulation of canonical angular momentum. The increase in the number of coils allows them to be smaller and provide greater focussing through the absorber.	34
2.1	Schematic of the MICE beam line from target to the upstream tracker. Q1 to Q9 are quadrupole magnets used for containment and matching, Dipole 1 and 2 are used to select the momentum of pions and muons respectively, the Decay Solenoid increases the efficiency of containing muons following pion decay and the TOF0, 1, and 2 detectors permit particle ID and global tracking. The Cerenkov detectors (Ckov) offer additional, coarse data for particle identification and the beam position monitor (BPM) may be used to determine the approximate envelope of the beam. Finally, the diffuser can be used to artificially increase the emittance of the beam before the cooling channel.	37
2.2	Left: schematic diagram of the structure of the target. Right: photograph of the target as installed upon the ISIS synchrotron.	37
2.3	Left: front view of a Cerenkov detector with the aerogel removed. Each quadrant is monitored using a single photo-multiplier tube. Right: Both Cerenkov detectors as installed in the MICE beam-line, immediately downstream of Q6 (blue in the image).	38
2.4	Left: TOF0 as installed on the end of Q6. Right: The TOF2, KL and EMR detectors as installed together. TOF2 is the first detector, just in front of the KL detector, followed by the EMR.	39
2.5	Photograph of the EMR as installed in the MICE experiment during the construction phase of MICE Step IV.	40
2.6	Layout of the MICE cooling channel, as constructed. The Diffuser, SciFi Trackers and Absorber are marked, and each magnetic coil is shown in red.	41
2.7	The mean z -component of the magnetic field throughout the cooling channel (left) and the mean transverse betatron function throughout the cooling channel (right). The red line marks the position of the absorber and the green lines mark the positions of the 10 tracker stations.	43

2.8	Behaviour of the normalised transverse emittance of a realistic muon beam with the MICE Step IV cooling channel. The centre of the liquid hydrogen absorber is marked in red and the 10 tracker stations are marked in green.	44
2.9	Schematic layout of the planes and how they are related to each other within a station. All measurements are in microns.	45
2.10	Photograph of a completed tracker. Note that to prevent discolouration of the scintillating fibres, they must be stored either in darkness or a yellow light. The 5 individual stations can be clearly seen in the image as the wavelength shifting fibres have not yet been connected to the edges of the planes.	46
3.1	Conceptual behaviour of the Kalman Filter in the prediction and filtering stages. The blue arrow is a prediction made from a previous state with an error indicated by the blue circle. The measurement (red star) and associated error is used to form the filtered state and associated error (green arrow and circle). . . .	54
3.2	Conceptual behaviour of the Kalman Filter in the smoothing stage. The black arrows represent the smoothed track, which is created by using the reverse-propagated filtered states (green arrows), compared with the measurements at each point (red stars). Owing to the gain of information, the earliest filtered states are more likely to be heavily corrected by this process than the later states.	55
4.1	The reconstructed p-value distribution for a simulation of straight (left) and helical (right) muon tracks, where the only physical effect is the mean energy loss.	69
4.2	The reconstructed p-value distribution for a simulation of straight (left) and helical (right) muon tracks, where the only physical effect is the measurement system.	70
4.3	The reconstructed P value distribution for a simulation of straight (left) and helical (right) muon tracks, where the only physical effects are multiple Coulomb scattering and energy straggling. . .	71
4.4	The reconstructed P value distribution for an simulation of straight (left) and helical (right) muon tracks, where all physical effects were modelled in the reconstruction.	72
4.5	The $x - y$ position residuals for both straight tracks (left) and helical tracks (right) at the tracker reference plane.	73

4.6	The transverse gradient residuals (left) and transverse momentum residuals (right) of the upstream tracker.	74
4.7	The longitudinal momentum resolution (left) and transverse momentum resolution (right) as a function of transverse momentum.	75
4.8	The deviation of the pattern recognition total momentum from the true total momentum at the Kalman seed position in the upstream tracker (left) and downstream tracker (right).	76
4.9	The deviation of the total track momentum, evaluated at both the upstream (left) and downstream (right) reference planes, from the true track momentum.	77
4.10	Efficiency of the upstream (left) and downstream (right) helical track reconstruction algorithms as a function of transverse momentum. The inefficiency in the lowest bin is due to the difficulties of reconstructing tracks with low transverse momentum.	78
5.1	The time-of-flight distribution between TOF1 and TOF2. The events within the TOF cut are highlighted in red.	85
5.2	Distribution of raw ADC values (left) and the estimated NPE (right) for every digit across Plane 1, Station 2 in the upstream tracker.	86
5.3	Distribution of reconstructed digits across a single tracker plane in the upstream tracker, station 2, plane 1. The shape of the distribution is due to the natural muon beam profile after leaving the final quadrupole.	87
5.4	The number of spacepoints in each station. Red is the total number of spacepoints recorded, green are only those spacepoints that were assigned to a track during pattern recognition and the blue shows only those spacepoints that were assigned to a track and were composed of three clusters. Negative station IDs refer to the upstream tracker and positive station IDs refer to the downstream tracker.	89
5.5	The distribution of recorded digits in planes 0 (left), 1 (centre) and 2 (right) of Station 5 in the downstream tracker. Note the large number of dead channels in plane 2 that detrimentally affect the reconstruction.	89

5.6	An example Kuno distribution of the clusters found in Station 1 of the upstream tracker (left) and a comparison with Monte Carlo on the region of interest (right). The well defined peak at 318.5 corresponds to the expected Kuno-peak for the number of channels used in the construction of the station. The remaining entries correspond to a collection of noise, knock on electrons and poorly formed spacepoints. The two peak structure is an artefact of the integer binning used in determining the channel number.	91
5.7	NPE per spacepoint. The red line includes all spacepoints that were reconstructed, the black line only includes spacepoints that were successfully used in a straight line, track fit.	92
5.8	The distributions of weighted trackpoint pulls at the reference frame of the upstream (left) and downstream (right) trackers, for Dataset A (top) and Dataset B (bottom). The RMS deviation from zero is approximately 0.37 for all plots.	94
5.9	Planar representation of the overlapping fibres in the centre (green) of a station. The red markers indicate the centre of the overlapping channels and can be seen to alternate in their vertical position.	94
5.10	A comparison of the reconstructed p-values, for the straight track reconstruction, from data (black) and the Monte Carlo simulated p-values (red). Top: Dataset A, Bottom: Dataset B, Left: Upstream Tracker, Right: Downstream Tracker.	95
5.11	The time-of-flight distribution between TOF0 and TOF1. The events within the TOF cut are highlighted. The secondary peak was believed to be due to a combination of poor calibration and noise, and was removed as a precaution only.	97
5.12	The distribution of weighted trackpoint pulls at the reference frame of the upstream tracker.	98
5.13	A comparison of the reconstructed p-values, for the helical track reconstruction, from data (black) and the Monte Carlo simulated p-values (red) in the upstream tracker.	98
6.1	Schematic representation of a reconstructed straight track and how a misalignment affects the global placement of the reconstruction. Note that the misalignment has been greatly exaggerated and is not to scale.	102
6.2	Geometrical layout of the alignment variables, and their relation to each other. This diagram describes some arbitrary axis, the x and y parametrisations are identical.	104

6.3	Comparison of the two Time of Flight windows used in this analysis. Dataset A, corresponding to the 200 MeV/c muon beam are in the left plot and Dataset B corresponding to the 240 MeV/c muon beam are in the right plot.	105
6.4	Position (left) and angular (right) residuals calculated by propagating the track parameters from the upstream tracker to the downstream tracker. The position residuals have x and y standard deviations of approximately 20 mm. The angular residuals have x and y standard deviations of approximately 8.3 mrad. . .	107
6.5	Position (left) and track angle (right) residuals for the ideal Monte Carlo Simulation. The x and y standard deviations of the position distribution is approximately 19 mm and the x and y standard deviations of the angular distribution is approximately 8 mrad.	108
6.6	Position (left) and track angle (right) residuals for the ideal Monte Carlo Simulation. The standard deviation of the position distribution is approximately 19 mm and the standard deviation of the angular distribution is approximately 8 mrad.	109
6.7	Upstream (left) and downstream (right) residuals calculated by propagating the track parameters to the nearest TOF detector. The upstream tracker was propagated to TOF1 and the downstream tracker was propagated to TOF2.	110
6.8	The correlations between the predicted and reconstructed transverse coordinates at TOF1, as compared to the upstream tracker. x is shown in the left histogram and y is shown in the right histogram. The black points mark describe the profile of the residuals in each TOF measurement bin. The gradient of the correlation was calculated from the linear fit drawn in red.	110
6.9	The correlations between the predicted and reconstructed transverse coordinates at TOF2, as compared to the downstream tracker. x is shown in the left histogram and y is shown in the right histogram. The black points mark describe the profile of the residuals in each TOF measurement bin. The gradient of the correlation was calculated from the linear fit drawn in red.	111
6.10	Residual between the projected track azimuthal angle and the measured azimuthal angle. A gaussian fit (shown in red) was applied to the peak in order to determine an approximate parametrisation for the distribution.	111

7.1	A schematic layout of the simulation geometry used to probe the effects of misaligning the tracker with respect to the solenoid field.	123
7.2	The mean residual between the reconstructed beam emittance and the true beam emittance, evaluated at the reference plane, as a function of the misalignment of the tracker to the field. Four difference emittances were simulated: 3, 6, 8 and 10 mm.	124
7.3	Comparison of the total momentum residual before (left) and after (right) a linear correction factor has been applied.	126
7.4	The emittance reconstruction residuals of the upstream tracker (left) and the reconstruction residual including the linear momentum correction (right), as predicted by Monte Carlo.	128
7.5	The emittance reconstruction residuals of the upstream tracker including the momentum corrections and a covariance matrix correction, calculated using the simulated 5 mm emittance beam.	129
7.6	Distribution of total momentum at the upstream tracker reference plane. The events that were included in this Monte Carlo analysis are highlighted in red.	130
7.7	Distribution of total momentum at the upstream tracker reference plane for Run 07469. The events that were included in this analysis are highlighted in red.	132
7.8	The x (left) and y (right) phase-space distributions for the reconstructed beam at the upstream reference plane for Run 07469.	133
B.1	Distributions of reconstructed digits from Dataset A for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the upstream tracker.	151
B.2	Distributions of reconstructed digits from Dataset A for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the downstream tracker.	152
B.3	Distributions of reconstructed pulls from Dataset A for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the upstream tracker.	153
B.4	Distributions of reconstructed pulls from Dataset A for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the downstream tracker.	154
B.5	Distributions of the Kuno-Sums for spacepoints used to form a track from Dataset A in stations 1 to 5 in the upstream tracker.	155
B.6	Distributions of the Kuno-Sums for spacepoints used to form a track from Dataset A in stations 1 to 5 in the downstream tracker.	155

B.7	Distributions of reconstructed digits from Dataset B for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the upstream tracker.	156
B.8	Distributions of reconstructed digits from Dataset B for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the downstream tracker.	157
B.9	Distributions of reconstructed pulls from Dataset B for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the upstream tracker.	158
B.10	Distributions of reconstructed pulls from Dataset B for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the downstream tracker.	159
B.11	Distributions of the Kuno-Sums for spacepoints used to form a track from Dataset B in stations 1 to 5 in the upstream tracker. .	160
B.12	Distributions of the Kuno-Sums for spacepoints used to form a track from Dataset B in stations 1 to 5 in the downstream tracker. Note that station 5 was constructed with an additional channel in plane 2, hence the peak is appropriately displaced.	160

Chapter 1

Introduction

1.1 Muon Acceleration

Muon acceleration has been recognised as an attractive future technology for High Energy Physics (HEP) experiments. Muon colliders would offer a precise centre-of-mass energy similar to electrons, but with vastly reduced losses from both synchrotron radiation and “beamstrahlung”,¹ as the radiative losses scale as $1/m^4$. At present the number of complex technological stages required, although believed by most to be within reach of current technology, leave substantial reservations in many supporters. The predominant difficulty is the short muon life time (approximately $2.2\mu\text{s}$) which cannot simply be overcome by exploiting time dilation. The production, collimation and acceleration of a muon beam must be conducted faster than in any other accelerator.

In order to further the development of technology, the Muon Accelerator Project (MAP) [1] developed the Muon Accelerator Staging Study [2], in which a staged construction and development project is conducted allowing individual components to be built and tested, whilst making detailed physics measurements. Initially a NuSTORM [3, 4] like facility could be constructed, providing neutrino beams from pion and muon decays for neutrino interaction studies and sterile neutrino searches. This could be later upgraded to a full Neutrino Factory facility (NuMAX) as shown in figure 1.1. The overlap between the initial stages of a Neutrino Factory and a Muon Collider would later permit further upgrades, culminating with a dedicated collider storage ring. The energy of the Muon Collider could be initially designed to probe the Higgs boson with unparalleled precision, before a final set of upgrades allows for multi-TeV

¹Beamstrahlung is the radiation originating from the interaction of a charged particle beam with the electromagnetic field generated by an oncoming beam.

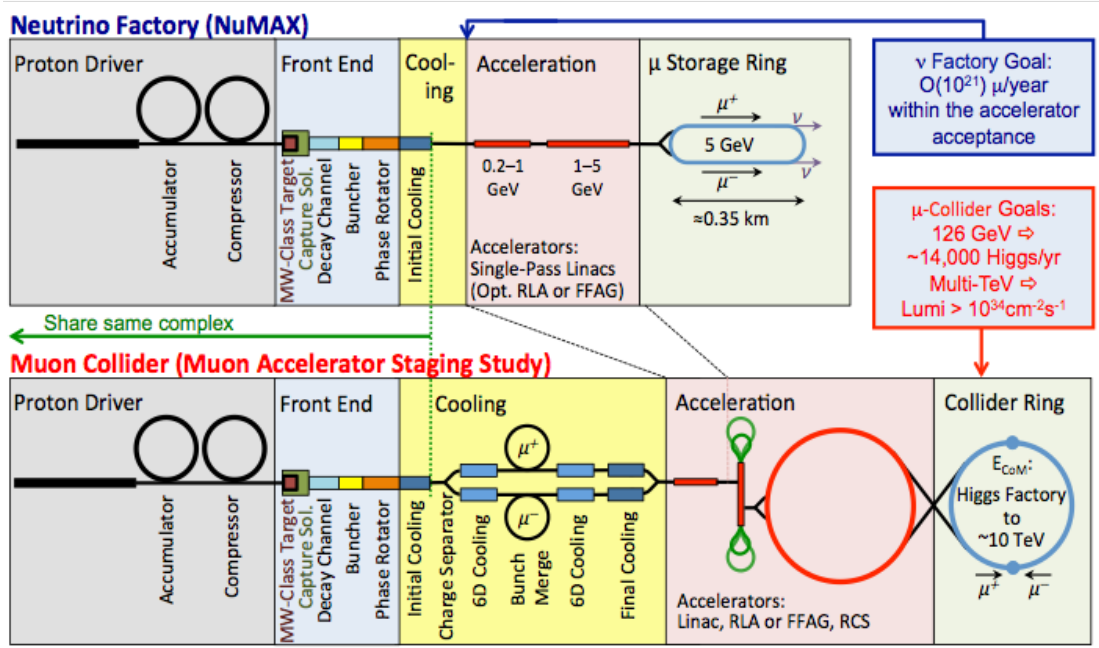


Figure 1.1: Block diagram describing the overall construction of a Neutrino Factory (NuMAX) and a Muon Collider. A visualization including which components can be easily transferred from one facility to another is included.

centre-of-mass lepton-antilepton collisions.

Before progressing towards a construction phase however, each component must be individually developed. Several of which must also demonstrate the expected design performance experimentally.

The Front End Test Stand [5] was developed to demonstrate early stage proton acceleration techniques required to generate the megawatt powered proton beam. Such a beam is necessary to generate high enough numbers of pion decays to produce high luminosity muon beams.

In order to quickly and efficiently disperse the energy of the high powered proton beam, whilst providing a sufficient target cross section for pion production, a liquid heavy metal target was the optimal choice. More traditional choices such as solid metal and graphite targets simply would not sustain the energy deposition for an extended period of time. The MERIT experiment [6], conducted at CERN, showed that the construction of liquid mercury target is feasible.

The options for rapid acceleration have also been investigated. Feasible designs [7] include Fixed Field Alternating Gradient (FFAG) accelerators and Recirculating Linear Accelerators (RLAs), as depicted in figure 1.1. At higher

energies, rapid-cycling synchrotrons become more favourable, e.g. [8].

Most stages required to accelerate a muon beam however, must support a much larger beam diameter when compared to more conventional facilities. This is due to the large muon beam emittance at source, a side effect of both the proton–target interaction and subsequent pion decays, and significantly increases the difficulty of construction.

The most essential component of a high energy muon facility is therefore the cooling stage. In order to collimate and focus a highly divergent beam with a small lifetime, a novel method to reduce the transverse (and longitudinal) phase-space must be considered. Ionisation cooling [9] is currently the only feasible process by which the beam phase-space may be reduced quickly enough to produce a high energy and high luminosity beam, without experiencing significant losses due to muon decay. In order to achieve the transverse and longitudinal beam dimensions required for a Muon Collider however, the relative cooling power must be significantly greater than that for the Neutrino Factory. This suggests that a staged construction/development project could be implemented, where the additional stages further reduce the transverse phase-space in addition to providing longitudinal cooling.

The Muon Ionisation Cooling Experiment (MICE) [10, 11] is currently under construction and aims to be the first experimental demonstration of the use of ionisation cooling to reduce the transverse phase-space of a muon beam. With the conclusion of MICE, all the major components of a Neutrino Factory will have been demonstrated successfully, opening a route for muon accelerators to be further developed.

The content of this thesis will focus on the development of the MICE experiment, with particular focus on the primary detectors, the Scintillating Fibre Trackers. They are required to measure the phase-space of the MICE muon beam, before and after a cooling channel, thereby demonstrating ionisation cooling. However, in order to measure “cooling”, the phase-space of a muon beam must be precisely parameterised.

1.2 Beam Optics and Emittance

The transverse coordinates of particle beams must be well controlled in order to confine beams within beam pipes and highly transversely focussed at interaction points, as in collider experiments. Additionally, the longitudinal coordinates must be carefully controlled to maintain a correct phase relationship with RF accelerating structures and general time-based events like interactions, injec-

tion and extraction, etc. The transverse phase-space will be the focus of this document, as the MICE experiment was primarily designed to provide precise transverse diagnostics.

1.2.1 Phase-Space Description

Every particle in a beam, travelling along the positive z axis, can be completely described by its position and momentum vectors, (t, x, y, z) and (E, p_x, p_y, p_z) , where t represents the time in the laboratory frame of reference, z is the longitudinal spatial dimension, parallel to the beam direction, and x and y are the remaining transverse coordinates where y is assumed to be the vertical coordinate; E represents the total particle energy and p_x , p_y , and p_z represent the momentum of the particle projected into the position coordinate system. In the transverse plane, particles are typically described by their geometric coordinates: (x, x', y, y', E, t) , where $x' = p_x/p_z$ and $y' = p_y/p_z$.

When sufficient numbers of particles are analysed, it is possible to describe the whole ensemble of particles by their distribution in phase-space. It is typical to consider longitudinal, (E, t) or (p_z, z) , and transverse, (x, x', y, y') , phase-spaces separately, however the full 6-dimensional phase-space (x, x', y, y', E, t) may also be considered.

A distribution of particles may typically be described by a probability density function in 6-dimensional phase-space, thus allowing the distribution to be described at first order with the mean and variance of the phase-space density in each dimension. In multiple dimensions, the scalar mean is generalised to a vector of the means in each dimension. Similarly, the variance generalises to a 6×6 covariance matrix. Higher order moments become increasingly complex higher-order objects. The mean of each dimension is typically constrained by the beam pipe or choice of coordinate system, hence it is typical to completely describe a particle beam using only the 6×6 covariance matrix, or some subset of coordinates.

Under the assumption that an ensemble of particles forms a gaussian probability distribution, modelling the ensemble by a single covariance matrix will generally provide an accurate description of the beam. In rare occasions however, non-linear transformations will cause effects that are not correctly modelled, resulting in a non-linear growth of the measured phase-space (see section 1.2.2). However, most magnetic configurations are designed to mitigate such effects.

The mean and covariance of each conjugate pair of coordinates (e.g. $x-x'$) may be computed such that a 2×2 covariance matrix, Σ , may be constructed that

describes the distribution, equation 1.1. The central 68% of the distribution forms an ellipse in phase-space and can be described by the covariance matrix,

$$\mathbf{\Sigma} = \begin{pmatrix} \text{Var}(x, x), & \text{Cov}(x, x') \\ \text{Cov}(x', x), & \text{Var}(x', x') \end{pmatrix} = \epsilon_x \begin{pmatrix} \beta_x & -\alpha_x \\ -\alpha_x & \gamma_x \end{pmatrix}. \quad (1.1)$$

where ϵ_x is the Root Mean Squared (RMS) emittance of the beam and α_x , β_x and γ_x describe the shape of the beam phase-space in the (x, x') dimensions. The emittance corresponds to the volume of phase-space occupied by the central 68% of the beam and is the most commonly considered parameter for the requirements of an accelerator. It combines the overall size of the beam with its dynamic confinement and as such provides a single measure of how well contained and focussed the beam is.

α_x , β_x and γ_x are the Twiss parameters and correspond to the physical description of the phase-space distribution. Commonly the betatron function, β_x , is discussed as it describes the width of the position distribution for a given emittance. It therefore is a measure of how well focussed the beam is and is commonly used in the design of magnetic lattices.

Using the properties of an ellipse, centred on the origin, the Twiss parameters can be related to each other as in figure 1.2, and can be used to parameterise the phase-space ellipse,

$$\epsilon_x = \gamma_x x^2 + 2\alpha_x x x' + \beta_x x'^2. \quad (1.2)$$




Figure 1.2: The relations between the Twiss parameters in the context of the RMS beam envelope.

If the particles within a beam were individually measured, the resulting phase-space distribution could be analysed and used to generate the corresponding covariance matrix, Σ . The emittance of the beam can then be directly calculated in x - x' plane using the determinant of the covariance matrix,

$$\epsilon_x = \sqrt{|\Sigma|}. \quad (1.3)$$

This formulation is easily generalised from 2 to n dimensions. For a $n \times n$ covariance matrix, the corresponding n -dimensional emittance, ϵ_{nD} is given by,

$$\epsilon_{nD} = |\Sigma|^{\frac{1}{n}}. \quad (1.4)$$

Including more dimensions in the calculation increases the number of features that will affect the emittance. In the 4 dimensional phase-space, the emittance is not uniquely defined by the individual x and y phase-spaces. Rather the presence of mechanical angular momentum, L , creates additional x - y correlations, where L is defined by,

$$L = \langle xy' - yx' \rangle. \quad (1.5)$$

For beams with no x - y correlations, which is typical of beams with no net angular momentum, the 4 dimensional emittance is related to the individual 2 dimensional emittances by,

$$\epsilon_{4D} = \sqrt{\epsilon_x \epsilon_y}. \quad (1.6)$$

The addition of angular momentum causes this equality to fail as additional terms are then included in the calculation of the 4 dimensional determinant.

Additionally, if the beam is subjected to accelerating forces or traverses some substantial medium, the mean energy of the beam will change. As the components, x' and y' are defined with respect to the momentum, they are no longer conserved in the transverse phase-space. For this reason the normalised RMS emittance, ϵ_N , which is invariant under acceleration, is the more commonly used parameter. It can be directly calculated from a covariance matrix as,

$$\epsilon_N = \frac{(\Sigma)^{\frac{1}{n}}}{\beta\gamma} \quad (1.7)$$

where β and γ are the usual relativistic functions.

The 2 dimensional α_x , β_x , and γ_x parameters are also generalised to an arbitrary covariance matrix. Using the relationships in figure 1.2, and averaging the appropriate variances/covariances, the 4 dimensional, transverse parameters

are given by,

$$\alpha_{\perp} = \frac{\langle xx' \rangle + \langle yy' \rangle}{2\epsilon_{\perp}}, \quad (1.8)$$

$$\beta_{\perp} = \frac{\langle x^2 \rangle + \langle y^2 \rangle}{2\epsilon_{\perp}}, \quad (1.9)$$

$$\gamma_{\perp} = \frac{\langle x'^2 \rangle + \langle y'^2 \rangle}{2\epsilon_{\perp}}, \quad (1.10)$$

where ϵ_{\perp} is the 4-dimensional emittance calculated from the (x, x', y, y') phase-space.

1.2.2 Linear Beam Optics

Once the beam has been appropriately described at a given position in space or time, it is possible to define functions that propagate that description to a subsequent position. The simplest example is the drift space of length l . Without any magnetic fields or materials, each particle will continue along its momentum trajectory unimpeded. The effect on the phase-space is a stretch of the distribution in the x and y components, and a constant distribution in x' and y' . As there is no x - y coupling the x and y distributions may be considered separately. Figure 1.3 shows the effect of a drift space on the distribution.




Figure 1.3: The effect of some small drift on the x - x' phase-space distribution for a beam in arbitrary units. The particles with larger gradient have moved further from the centre of the position axis, while the distribution of gradients has remained constant.

This can be realised by calculating a transfer matrix, designed to transform the beam parameters from one state into another. For the simple case of a drift

space, the transfer matrix can be written down as:

$$\mathbf{M} = \begin{pmatrix} 1 & l \\ 0 & 1 \end{pmatrix}. \quad (1.11)$$

If we apply this to a particle state vector, $(x \ x')^T$ it can be seen that the transverse position increases as a linear function of the transverse gradient, as expected. This transformation can be applied to both the particle state vector, \mathbf{x} , and covariance matrix, \mathbf{V} , in the usual fashion,

$$\begin{aligned} \mathbf{x}(l) &= \mathbf{M}\mathbf{x}(0), \\ \mathbf{V}(l) &= \mathbf{M}\mathbf{V}(0)\mathbf{M}^T. \end{aligned} \quad (1.12)$$

This simple linear approach can be easily extended to a variety of magnetic configurations, using the correct calculation of the transfer matrix. The transformations through the most common magnet types, (dipoles and quadrupoles)² are well documented, and typically considered as thin lenses within the formalism. A detailed overview of the mathematical foundation of modelling charged particle beams can be found in ref [12]. Magnetic lattices can always be subdivided into a series of short elements, which can then be modelled using a combination of thin lenses and drifts.

1.2.3 Solenoid Optics

For most beamlines, there is little or no transverse coupling. Beams are ideally produced symmetrically and few magnets induce any x - y coupling, hence the phase-space may be modelled in x and y independently. For solenoid optics however, there is an inherent coupling in the transverse plane as the Larmor rotation of the particles, through the magnetic field, sees them rotate between the x and y phase-space distributions. Hence the phase-space must be inherently modelled in at least 4 dimensions, the longitudinal components are still uncoupled to first order, but may also be included.

The Twiss parameters may similarly be defined [13] using the expected x - y symmetry in the beam. The standard 2×2 covariance matrix is therefore extended to a 4×4 covariance matrix, with components (x, x', y, y') , assuming

²Higher order magnetic structures, sextupoles, octupoles, etc., are not linear at first order and hence are generally more difficult to model in the linear optics regime.

a gaussian distribution that is symmetric in both x and y ,

$$\Sigma_{4D} = \epsilon_{\perp} \begin{pmatrix} \beta_{\perp} & -\alpha_{\perp} & 0 & -(\beta_{\perp}\kappa - \mathcal{L}) \\ -\alpha_{\perp} & \gamma_{\perp} & (\beta_{\perp}\kappa - \mathcal{L}) & 0 \\ 0 & (\beta_{\perp}\kappa - \mathcal{L}) & \beta_{\perp} & -\alpha_{\perp} \\ -(\beta_{\perp}\kappa - \mathcal{L}) & 0 & -\alpha_{\perp} & \gamma_{\perp} \end{pmatrix} \quad (1.13)$$

where $\kappa = qB_z/2p_z$, α_{\perp} , β_{\perp} and γ_{\perp} are the transverse Twiss parameters, ϵ_{\perp} is the geometric 4D RMS beam emittance and \mathcal{L} is a dimensionless parameter defined by the mechanical angular momentum. Additionally,

$$L = \langle xp_y - yp_x \rangle = -2p_z\epsilon_{\perp}(\beta_{\perp}\kappa - \mathcal{L}), \quad (1.14)$$

$$\mathcal{L} \approx \frac{\langle L_{\text{canon}} \rangle}{2p_z\epsilon}. \quad (1.15)$$

The structure resembles that of the 2 dimensional case, with off-diagonal terms as a direct result of the mean canonical angular momentum ($\langle L_{\text{canon}} \rangle$). The addition of these terms implies that the 4 dimensional emittance is greater than the sum of the two, 2 dimensional emittances (in x and y). Hence the full 4D phase-space must be reconstructed in order to precisely determine the total transverse emittance.

Note also that this formulation assumes some symmetry in the initial beam distribution. If the x and y components are strongly correlated, the equations of motion based on this parameterisation will not completely describe the evolution of the phase-space. Hence the distributions in each transverse component are still of use during analyses.

This parameterisation is predominantly useful in the design and modelling of solenoid beam optics as it allows for the derivation of transfer matrices, similar to those used with the Twiss parameters. Indeed the parameter, κ^2 , can be shown to be proportional to the focussing strength of a thin solenoid ‘‘lens’’.

1.2.4 Liouville’s Theorem

Liouville’s theorem is a key component in the discussion of beam optics and the propagation of beam ellipses, and may be derived with relatively few simple assumptions.

Consider a closed system of non-interacting particles with no external forces.³ The example beam phase-space can then be modelled with some unspecified cen-

³Interacting particles may also be modelled where the space charge induced electric and magnetic fields, between a particle and its nearest neighbour, are negligible in comparison to the collective fields.

tral distribution, with a number density given by, $n(t, x, y, z, E, p_x, p_y, p_z)$. It is assumed that the canonical position and momentum coordinates can be given by q_i and p_i respectively, hence a velocity vector may be constructed, $\mathbf{v} = \{\dot{q}_i, \dot{p}_i\}$, for each particle in the phase-space. Therefore, in order to conserve the number of particles within the beam, the phase-space density must obey the continuity equation:

$$\begin{aligned}\nabla \cdot (n\mathbf{v}) + \frac{\partial n}{\partial t} &= 0, \\ n\nabla \cdot \mathbf{v} + \mathbf{v} \cdot \nabla n + \frac{\partial n}{\partial t} &= 0.\end{aligned}\tag{1.16}$$

Hamilton's equations for this system,

$$\dot{q}_i = \frac{\partial H}{\partial p_i}, \quad \dot{p}_i = -\frac{\partial H}{\partial q_i},\tag{1.17}$$

where H is the Hamiltonian, may be used to simplify the first term,

$$\nabla \cdot \mathbf{v} = \frac{\partial \dot{q}_i}{\partial q_i} + \frac{\partial \dot{p}_j}{\partial p_j},\tag{1.18}$$

$$\nabla \cdot \mathbf{v} = \frac{\partial^2 H}{\partial q_i \partial p_i} - \frac{\partial^2 H}{\partial p_j \partial q_j} = 0.\tag{1.19}$$

Hence the continuity equation may be simplified to,

$$\mathbf{v} \cdot \nabla n + \frac{\partial n}{\partial t} = 0,\tag{1.20}$$

$$\dot{q}_i \frac{\partial n}{\partial q_i} + \dot{p}_j \frac{\partial n}{\partial p_j} + \frac{\partial n}{\partial t} = \frac{dn}{dt} = 0;\tag{1.21}$$

that is, the phase-space density remains constant with time. Liouville's theorem states that the total occupied volume of a given number of particles in phase-space is invariant. If conservative transformations are introduced, it can be similarly shown that the total volume of phase-space is still conserved even if the distribution changes. This directly applies to the modelling and design of magnetic lattices, as it implies that a measure of the phase-space volume, e.g. the emittance, must also remain constant.

As magnetic fields induce conservative forces, magnetic lattices will conserve the beam phase-space, and by extension the beam emittance, assuming that the distribution of the phase-space only undergoes linear transformations. Section 1.2.5 describes an occasion where non-linearities cause a growth in the measured emittance, however the total phase-space volume still remains constant. This also highlights the fundamental difficulty in beam cooling, i.e. re-

ducing the phase-space volume, as non-conservative forces must be applied to the beam.

1.2.5 Matching a Beam

Linear beam optics is widely used as the simplest method to model beamlines and predict particle phase-space distributions, without the need to simulate many thousands of particles. However, as the approach is only accurate to first order, a combination of linear optics and full tracking of particles, within a magnetic field model, is generally required in order to precisely design a magnetic lattice.

During the procedure of matching, the settings for each magnet may be adjusted in order to vary the focussing strength of each individual component. The objective is to achieve the required beam optics at positions of interest within the lattice. In circular colliders/storage rings, where beams experience multiple turns through the same magnetic lattice, the Twiss parameters are generally required to start and finish with the same values, thereby preserving the machine symmetry within the beam optics and reducing losses.

In linear beamlines, such as MICE however, the matching of the beam with respect to the quasi-symmetric structures is still of great interest. Correct matching ensures that the beam can be optimally focussed through the absorber (the reasons for doing so are discussed below), and that a high transmission can be maintained. Cooling channels such as those designed for the neutrino factory, require long chains of similarly designed cells, hence it is optimal for the optics of the beam to reflect the repetitive nature of the machine.

The natural emittance of the MICE beamline is many orders of magnitude higher than that of a typical electron or proton synchrotron due to the production mechanism of the muon beam - a proton-target interaction and subsequent pion decays. The larger phase-space volume is much more susceptible to non-linear effects as more muons are further from the origin of the phase-space, where the linear approximation starts to fail. The complex magnetic configuration of MICE (see section 2.4), with several strong focussing fields further amplifies the sensitivity to non-linear effects.

Poor matching of a beamline will generally lead to a non-linear increase in the beam emittance, which cannot be predicted by linear beam optics. Figure 1.4 illustrates one such example of this non-linear emittance growth - a phenomenon known as filamentation. This is where the extremes of the distribution are in a regime which is only very approximately linear, leading to a different behaviour which is generally only notable over a large amount of time spent in the magnetic




Figure 1.4: An illustration of the filamentation phenomenon and how over time it may lead to a non-linear emittance growth despite Liouville's theorem.

lattice.

1.3 Ionisation Cooling

Many particle accelerator based experiments are reliant on high quality beams. For example, small transverse (and longitudinal) phase-spaces are required to ensure colliding bunches are precisely focussed, and beams that use complex or stochastically driven production mechanisms are often produced with large phase-spaces that must be reduced. Although this thesis is focussed on muon beam development, similar issues are present in many other beamlines, each with a different optimal solution.

The proton-antiproton collisions that occurred at CERN were heavily reliant on stochastic cooling to reduce the transverse phase-space of the antiproton beam. Transverse electrical pick-ups detect the deviation of particles from the beam centre, and use that information to construct an active electric field designed to counteract the measured effect [14].

Additionally, schemes such as electron cooling have been used, where a “colder” electron beam is combined with the primary beam, allowed to interact through Coulomb Scattering and then extracted such that the temperature of the original beam is reduced. It was originally conceived by Budker [15] and first proven experimentally at Novosibirsk [16].

Laser cooling systems, similar to the mechanism used in laser-ion traps, have been used to great effect. A laser is focussed onto a beam of ions, thereby exciting them to a higher electrical state, before an isometric emission of photons




Figure 1.5: Diagrammatic representation of the process of ionisation cooling in momentum space. The green arrow represents the effect of the absorber - a net reduction in total momentum; the red arrow represents a smearing of the momentum spread due to multiple Coulomb scattering; and the blue arrow represents the final momentum following some accelerating structure. The net effect is a reduction in momentum spread, while maintaining the mean beam momentum.

follows. This results in a reduction of the momentum spread of the beam [17].

Each of the three systems mentioned (and indeed several others) have been experimentally tested and shown to produce efficient cooling for a variety of situations. However they all require a storage ring, such that the relatively small effect that they induce can be applied many thousands of times. For a muon beam however, the small lifetime renders almost every scheme inapplicable apart from ionisation cooling.

1.3.1 Principles of Ionisation Cooling

Consider a beam incident on a material, the Bethe formula [18] for mean energy loss describes the rate of energy loss per unit of areal density, the product of penetration distance and material density ($x = s\rho$). At first order this is proportional to the ratio of atomic number (Z) to atomic mass (A), in addition to being inversely proportional to the normalised velocity squared (β^2),

$$\left\langle \frac{dE}{dx} \right\rangle \propto \frac{Z}{A} \frac{1}{\beta^2}.$$

It is therefore optimal to use a material with a high value of Z/A , and a beam with a relatively low velocity. In practice, it is ideal to minimize the effect of multiple Coulomb scattering, which naturally acts as to increase the beam emittance, hence materials such as liquid hydrogen, liquid helium and lithium hydride are the optimal choice.

As the beam traverses some volume of the selected material - the absorber - it will naturally experience a reduction in the total momentum. If the longitudinal components are then restored, the beam will have experienced a net reduction in the transverse momentum spread. This corresponds to a decrease in the transverse emittance, i.e. transverse cooling. Figure 1.5 describes this process.

1.3.2 The Cooling Equation

In practical applications, muons are produced with a large phase-space and require strongly focussing magnets in order to contain them within beamlines. As they are generally produced at relatively low energy, and ionisation cooling performs better at low energy, solenoid optics are the natural choice.

In a system with no x - y coupling and $\alpha_{\perp} = 0$, the normalised transverse emittance can be expressed as the product of the position and angular standard deviations, $\epsilon_{\perp} = \beta\gamma\sigma_{(x,y)}\sigma_{(x',y')}$ where β and γ are the typical Lorentz factors. Therefore the differential emittance follows,

$$\frac{d\epsilon_{\perp}}{\epsilon_{\perp}} = \frac{dp}{p} = \frac{dE}{E} \frac{1}{\beta^2}, \quad (1.22)$$

as $dE = \beta dp$. Hence we may deduce that the rate of change of emittance as a function of the mean energy loss per unit of areal density, the cooling term, is given by,

$$\frac{d\epsilon_{\perp}}{dx} = -\frac{\epsilon_{\perp}}{\beta^2 E} \left\langle \frac{dE}{dx} \right\rangle. \quad (1.23)$$

Similarly, we may consider the ‘‘heating’’ effect due to multiple Coulomb scattering as a separate term. The approximate differential scattering angle per unit of areal density is typically given by [19] where a scattering model based on Highland’s work, as presented in the PDG, [18] is used.

$$\frac{\partial\sigma_{\theta}^2}{\partial x} \approx \frac{1}{X_0} \left(\frac{13.6\text{MeV}/c}{p\beta} \right)^2, \quad (1.24)$$

where X_0 is the radiation length of the material. The precise value of the constant is dependent on the choice of scattering model and the materials to which it applies. Other materials and models yield different values. Often the model described by Rossi [20] is projected into a single dimension producing a value of approximately 14 MeV/c. This can be incorporated in the emittance change, by considering the differential,

$$\begin{aligned} d\epsilon_{\perp}^2 &= \beta^2\gamma^2\sigma_x^2 d\sigma_{\theta}^2, \\ 2\epsilon_{\perp} d\epsilon_{\perp} &= \beta\gamma\beta_{\perp}\epsilon_{\perp} d\sigma_{\theta}^2, \\ d\epsilon_{\perp} &= \frac{\beta_{\perp}\gamma\beta}{2} d\sigma_{\theta}^2, \\ \frac{d\epsilon_{\perp}}{dx} &= \frac{\beta_{\perp}(13.6\text{MeV}/c)^2}{2\beta^3 E m_{\mu} X_0}, \end{aligned} \quad (1.25)$$

where the identity, $\sigma_x^2 = \epsilon_{\perp}\beta_{\perp}/\beta\gamma$ has been used. Equation 1.25 is the heating

term, the rate of change of emittance due to the effects of multiple Coulomb scattering.

The differential change in normalised transverse emittance may now be constructed as a function of the material traversed as a sum of the cooling term (1.23) and the heating term (1.25),

$$\frac{d\epsilon_{\perp}}{dx} = -\frac{\epsilon_{\perp}}{\beta^2 E} \left\langle \frac{dE}{dx} \right\rangle + \frac{\beta_{\perp} (13.6 \text{ MeV}/c)^2}{2\beta^3 E m_{\mu} X_0}. \quad (1.26)$$

This describes the approximate change in normalised transverse emittance, ϵ_{\perp} , of a muon beam with betatron function, β_{\perp} , as it propagates through a material of radiation length X_0 . Several key conclusions can be noted from the structure of this equation:

- The rate of cooling is optimized by having a large initial emittance. i.e. the rate of cooling decreases as the phase-space is compressed.
- Slower beams (approximately minimum ionising particles) experience a greater rate of cooling.
- Beams with a smaller betatron function experience less emittance growth due to multiple Coulomb scattering, hence they are “easier” to cool.
- The choice of material affects both the rate of energy loss and multiple Coulomb scattering, hence its a critical decision in any cooling apparatus.

It can also be noted that there is a well defined minimum, at the equilibrium emittance ϵ_{\perp}^0 . For a fixed value of beta and a given material, the beam cannot be cooled past this point. It is given by,

$$\epsilon_{\perp}^0 = \frac{\beta_{\perp} (13.6 \text{ MeV}/c)^2}{2\beta m_{\mu} X_0} \left\langle \frac{dE}{dx} \right\rangle^{-1}. \quad (1.27)$$

The derivation of the cooling equation uses only simple treatment of the effects of multiple Coulomb scattering and energy loss. As such the theoretical performance has been shown to disagree with Monte Carlo simulations [21], however the known deviations in mean scattering angle between Monte Carlo models and the true low-Z materials will cause further disagreements when compared to data. The MICE experiment has a unique ability to test the accuracy of these models and improve on the assumptions used in the cooling equation. This is a necessary step towards the development of a muon cooling channel.

1.3.3 Building a Muon Cooling Channel

The ideal cooling channel would be composed of very strong solenoid magnets, in order to create the smallest possible betatron function, with repeated absorbers and accelerating cavities. This would ensure that the beam energy is maintained, while the emittance is iteratively reduced. However the cost and practicality of a single long solenoid, housing liquid hydrogen or lithium hydride vessels and accelerating structures is nearly impossible, hence a periodic lattice structure is preferred. This not only allows for the required services to access each component, but the focussing strength can be concentrated at the absorbers and overall increased throughout the lattice, thereby gradually and efficiently reducing the equilibrium emittance.

The simplest periodic focussing lattice, a so-called FOFO (FOcussing - FOcussing) lattice, is designed such that the absorbers are placed in the region of lowest betatron function and RF cavities are placed in regions of higher betatron function.

An additional concern is that of how the canonical angular momentum is affected by the induced energy loss [22]. Consider a beam with zero canonical angular momentum that enters a solenoidal field, the mechanical angular momentum experiences a kick due to the field. Each individual particle experiences a change in mechanical momentum of $qB_z r/2$ due to the fringe field. The canonical momentum however, is conserved throughout by definition. When the beam passes through an absorber within the solenoid field, the momentum is isometrically reduced, which results in a change of both the mechanical and canonical angular momenta. As the beam then leaves the solenoid field and experiences the opposite momentum kick, a net increase in the canonical angular momentum is found. This is experienced as a net emittance growth. For this effect to be mitigated the solenoidal fields used must periodically change polarity through the absorbers.

Figure 1.6 describes an example solenoid FOFO lattice which forms a simplified cooling channel. After optimization however, a Super-FOFO (SFOFO) lattice is typically preferred, where twice the number of coils are used and placed much closer to the absorber. In this way the focussing strength can be increased, while the peak magnetic field is in fact decreased resulting in a more efficient design. Figure 1.7 describes this structure.

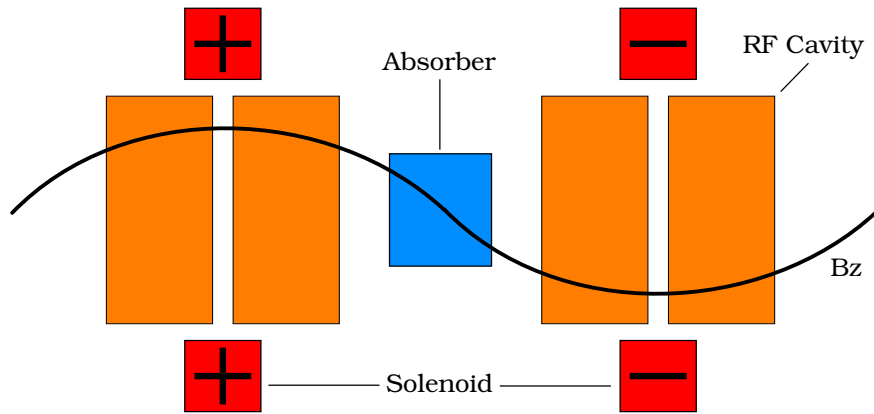


Figure 1.6: Representation of a FOFO-type cooling lattice. The absorbers (blue) are placed in regions of low betatron function to enhance the cooling performance. The RF accelerating cavities are placed in between to replenish the longitudinal momentum components. The relative field flip between pairs of solenoids prevents the accumulation of canonical angular momentum.




Figure 1.7: Representation of a Super-FOFO-type cooling lattice. The absorbers (blue) are placed in regions of low betatron function to enhance the cooling performance. The RF accelerating cavities are placed in between to replenish the longitudinal momentum components. The relative field flip between pairs of solenoids prevents the accumulation of canonical angular momentum. The increase in the number of coils allows them to be smaller and provide greater focussing through the absorber.

Chapter 2

MICE

The International Muon Ionisation Cooling Experiment (MICE), was designed to be the first experiment to demonstrate practical ionisation cooling and to provide the first measurement of normalised transverse emittance reduction in a muon beam. Since its initial conception, the experiment has been re-baselined and thoroughly scrutinised [23]. The current designs are expected to permit emittance and ionisation cooling measurements during 2016 (MICE Step IV).

A variety of detailed studies based on material physics (energy loss and multiple Coulomb scattering) and beam physics will also be conducted in parallel. MICE Step IV has simple configuration allowing both the incoming and outgoing beams to be measured to high precision, with flexible optics. The central absorber module was designed to support a range of different absorbers, allowing for the cooling performance of different materials to be measured. Hence the experimental setup is well suited to make measurements of material physics and the effects of different optical configurations, which are otherwise difficult to perform.

The programme will conclude with the completion of the final stage of development, the MICE Demonstration of Ionisation Cooling [23]. This is due in 2019 and will demonstrate emittance reduction with partial energy recovery in a realistic cooling cell, where the energy recovery is provided through two 201MHz RF cavities. Additionally, further beam-physics studies will be conducted with a detailed focus on the effects of different optical configurations on the cooling cell performance.

2.1 Introduction

The MICE experiment may be subdivided into three parts, the upstream beamline, the cooling channel, and the diagnostics and detectors; these are outlined in the sections 2.2, 2.3 and 2.4 respectively. MICE is a parasitic experiment that uses a titanium target, dipped into the halo of the ISIS proton beam [24], to produce a spray of baryons, predominately pions, within the ISIS synchrotron hall. The upstream beamline uses a combination of dipole magnets, for steering and momentum selection, and quadrupole magnets, for containment and matching to steer the pion beam into the MICE experiment hall. A superconducting solenoid is used as a strong focussing element to contain as many muons as possible from pion decay.

With the correct settings a high purity, low current (one muon per event), muon beam is produced directly into the cooling channel where the detailed emittance, material physics and optics studies are undertaken. The cooling channel contains the scintillating fibre trackers (ScFi Trackers), situated at the centre of a 4T solenoid, providing precise measurements of the muon position and momentum, in order to form a 6D phase-space for selected ensembles of particles. The remaining detectors are used for particle ID and event selection.

2.2 The Muon Beamline

The beamline is composed of some of the oldest components of the MICE experiment. Construction was completed in 2009 such that MICE Step I could be conducted [25]. An overview of the structure of beamline can be seen in figure 2.1. The MICE target [26] is composed of a titanium rod and a set of high powered linear motors. Owing to high repetition rate of ISIS, the controlled fluctuations in beam diameter and sensitivity of the ISIS machine the target had a precise specification to meet, which was achieved by ensuring an acceleration of 780ms^{-2} could be applied to the target itself.

The fine control of the target dip-depth permits very stable pion production for use in the MICE Beamline. The most upstream quadrupoles, Q1–3 were installed to capture as many of the emitted particles as possible, given constraints on construction¹. They direct the beam which is predominantly pions into the first dipole magnet, D1. The field strength in D1 permits the selection of the pion momentum and the sign of the charge, into the decay solenoid.

¹A structure similar to a magnetic horn would be a more efficient pion capture system, however the space available around the ISIS synchrotron, coupled with the low requirement on transmission, meant that this would not have been a feasible design.

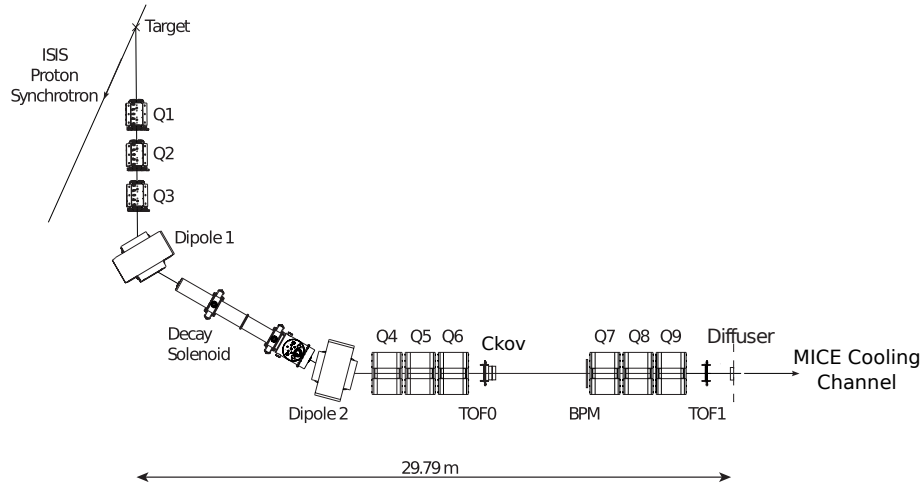


Figure 2.1: Schematic of the MICE beam line from target to the upstream tracker. Q1 to Q9 are quadrupole magnets used for containment and matching, Dipole 1 and 2 are used to select the momentum of pions and muons respectively, the Decay Solenoid increases the efficiency of containing muons following pion decay and the TOF0, 1, and 2 detectors permit particle ID and global tracking. The Cerenkov detectors (Ckov) offer additional, coarse data for particle identification and the beam position monitor (BPM) may be used to determine the approximate envelope of the beam. Finally, the diffuser can be used to artificially increase the emittance of the beam before the cooling channel.

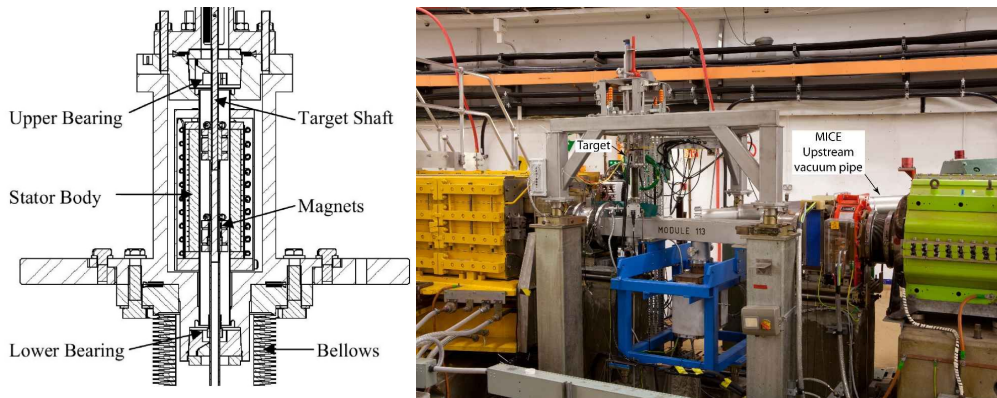


Figure 2.2: Left: schematic diagram of the structure of the target. Right: photograph of the target as installed upon the ISIS synchrotron.




Figure 2.3: Left: front view of a Cerenkov detector with the aerogel removed. Each quadrant is monitored using a single photo-multiplier tube. Right: Both Cerenkov detectors as installed in the MICE beamline, immediately downstream of Q6 (blue in the image).

The decay solenoid greatly increases the transmission of the beamline owing to the strong magnetic focussing, providing a large dynamical acceptance. A large fraction of muons are produced within the bore of the magnet, before being directed into the second dipole, D2, where the muon momentum can be effectively selected by controlling the field strength.

The remaining quadupoles serve as to contain the resulting muon beam as much as possible through the upstream section of the MICE hall. They additionally provide some magnetic degrees of freedom, permitting the beam to be matched into cooling channel as best as possible.

During operation the MICE target typically dips every 2s into the ISIS beam and produces on the order of 10-20 muons within the cooling channel, if the decay solenoid is not powered, or on the order of 100-200 muons with the decay solenoid powered. Each muon is sufficiently separated in time for individual particle ID and reconstruction routines to be applied.

2.3 Detector Systems

There is a range of detectors both up- and downstream of the cooling channel [27], designed to provide a very high precision particle identification (PID) system. The first of these are the Cerenkov detectors [28]. They are constructed from two separate types of aerogel, with carefully selected refractive indices, such that the velocity of incoming particles can be quickly and coarsely determined by analysing which aerogel slab has a high enough refractive index to invoke Cerenkov radiation.

Additionally the Time-Of-Flight (TOF) detectors are found upstream (TOF0




Figure 2.4: Left: TOF0 as installed on the end of Q6. Right: The TOF2, KL and EMR detectors as installed together. TOF2 is the first detector, just in front of the KL detector, followed by the EMR.

and TOF1) and downstream (TOF2) of the cooling channel [29]. The TOF detectors were constructed from scintillating bars and were designed to have a very fast response time to scintillation light. This permits any of the TOF detectors to act as the readout trigger for the experiment. In practice TOF1 is the most efficient trigger, as it is immediately upstream of the cooling channel. Figure 2.4 shows two of the TOF detectors, as installed within the MICE experiment. The three detectors permit two distinct measurements of individual particle time-of-flight, such that the particle species can be estimated both before and after the cooling channel. Additionally, accurate modelling of the beamline and high precision reconstruction algorithms permitted an emittance measurement of the standard MICE muon beam to be undertaken [30].

Downstream of the cooling channel, following TOF2, there are two more detectors - the Kloe-Light (KL) detector and the Electron-Muon Ranger (EMR). The KL detector was so named as it has a similar but smaller construction than that of the original Kloe calorimeter [31]. The layers of lead and scintillating fibres provide a measurement of particle penetration, especially when coupled with the EMR, as electrons can be easily tagged using the pre-shower effect of the KL. The KL and EMR were installed collectively with TOF2 to form a complete downstream calorimetry system, as shown in figure 2.4.

The EMR [32] is a fully active, scintillating-fibre based, sampling calorimeter, constructed from plastic scintillating bars. The high granularity in the longitudinal direction permits the penetration depth and track uniformity to be exploited very effectively in separating electrons from muons.




Figure 2.5: Photograph of the EMR as installed in the MICE experiment during the construction phase of MICE Step IV.

A pion contamination study was undertaken to demonstrate the accuracy of the particle ID system and the purity of the MICE muon beam [33].

2.4 MICE Cooling Channel

The MICE Cooling channel forms the main body of the MICE experiment. In the MICE Step IV configuration it is essentially composed of an upstream spectrometer (SSU), an Absorber-Focus Coil (AFC), and a downstream spectrometer (SSD). It was designed to allow the upstream and downstream spectrometers to analyse the muon beam as it enters, and again as it leaves the cooling channel, such that the evolution of the phase-space can be recorded. In order to contain and control the beam, the cooling channel contains 12 superconducting solenoid coils: 5 in each of the spectrometers and two in the absorber focus coil. Additionally, each spectrometer solenoid houses a scintillating fibre based tracker, in order to perform the particle track reconstruction. A brass-tungsten beam diffuser was installed in the upstream spectrometer, permitting the beam emittance to be artificially increased before entering the cooling channel, such that different emittance beams may be studied. It contains leaves of brass and tungsten that can be placed within the beam to artificially increase the beam emittance via multiple Coulomb scattering. Figure 2.6 describes physical layout of the cooling channel.

The MICE absorber was constructed to house liquid hydrogen within a low




Figure 2.6: Layout of the MICE cooling channel, as constructed. The Diffuser, SciFi Trackers and Absorber are marked, and each magnetic coil is shown in red.

density vacuum chamber, owing to the benefits of using hydrogen over other materials. The module is however reconfigurable, permitting different materials to be tested in addition to hydrogen. Due to the less volatile nature of lithium hydride, but very low density, it was chosen to be the primary addition to Step IV run plan.

With the magnets turned off, materials placed within the absorber module may be used to make precise measurements of multiple Coulomb scattering. With the magnets energised, the upstream and downstream beam momenta can be measured such that the mean energy-loss for different materials can be calculated. However the primary mode of operation is to measure the incoming and outgoing beam emittance and beta-function, such that the change in emittance through the chosen absorber material can be measured.

2.4.1 Magnetic Layout

A driving force behind the design of the MICE cooling channel were the existing designs for neutrino factory cooling channels that were being developed [34], section 1.3.3 outlines the basics of this development. The two spectrometer solenoids were constructed such that they may be approximated to an infinite solenoid field. This defines the points in the magnetic lattice where the beam must be matched into and out of the cooling channel. Between these two points, the “cooling cell” is defined. If the muon beam can be matched both at the beginning and the end of the cooling cell, it can be repeated to increase the

Coil	z mm	Current Density A/mm ²	Length mm	Thickness mm	Inner Radius mm
Up - E2	-3200.0	133.88	110.6	67.8	258.0
Up - C	-2450.0	144.90	1314.3	22.1	258.0
Up - E1	-1700.0	124.51	110.6	60.9	258.0
Up - M2	-1300.0	120.10	199.5	30.9	258.0
Up - M1	-861.0	129.04	201.3	46.2	258.0
Up - FC	-202.75	58.66	213.3	94.8	267.0
Down - FC	202.75	58.66	213.3	94.8	267.0
Down - M1	861.0	129.04	201.3	46.2	258.0
Down - M2	1300.0	120.10	199.5	30.9	258.0
Down - E1	1700.0	124.51	110.6	60.9	258.0
Down - C	2450.0	144.90	1314.3	22.1	258.0
Down - E2	3200.0	133.88	110.6	67.8	258.0

Table 2.1: Description of the magnetic lattice used to perform the simple Step IV Monte Carlo Studies shown in figures 2.7 and 2.8.

amount of emittance reduction.²

The original designs for the construction of MICE concluded with Step VI [35], however the experiment has since been re-baselined and will no longer continue to this final form. Designs for alternate lattices with lower financial burden and construction difficulties were investigated [23]. In both cases, the idea was to construct a component of a muon cooling channel that could represent modern designs.

The present state of MICE is the Step IV configuration, which serves to test only the AFC component of a cooling cell, containing the absorber and focussing magnets and does not include the RF cavities required to permit sustainable cooling. The AFC was designed to provide strong magnetic focussing at the absorber, to enhance the rate of emittance reduction. Figure 2.6 shows the physical layout of Step IV and figure 2.7 shows the magnetic field strength and betatron function throughout the lattice in “solenoid mode”, where there is no change in magnetic field polarity. The cooling channel can also be operated in “flip mode”, where the polarity of the fields is inverted about the AFC module. The parameters used to generate the simulation data are listed in table 2.1.

The beam was required to be correctly matched from the upstream spectrometer solenoid into the AFC module and finally into the downstream spectrometer solenoid. In order to do so, the 2 AFC coils and 2 match coils in each spectrometer were individually varied in order to find the set of magnetic currents that provide the optimal settings. The magnitude of the current in the

²The repetition however would not be indefinite due to the equilibrium emittance, defined in equation 1.27. Rather the focussing strength in each cell must be gradually increased as the transverse emittance is decreased.




Figure 2.7: The mean z -component of the magnetic field throughout the cooling channel (left) and the mean transverse betatron function throughout the cooling channel (right). The red line marks the position of the absorber and the green lines mark the positions of the 10 tracker stations.

AFC coils is the same, but may be reversed for “flip mode” operation.

Despite an iterative sequence of improvements to the magnetic configuration, there still exists some mismatch in the downstream spectrometer solenoid. This can be seen in figure 2.7–right, by the oscillations in the beta function within the uniform solenoid region. Further optimisations were conducted in order to improve the field uniformity within the tracker region and reduce this mismatch, however there was only a finite amount of improvement to be gained due to the momentum spread of the muon beam.

As the focussing strength of a magnet depends on the momentum of the particles experiencing the field, configurations such as MICE may only be truly matched for a single value of momentum, any additional spread in energy results some particles experiencing a different amount of focussing, which integrates through the magnetic fields and eventually results in a mismatch. Over a long and repeated cooling channel involving RF reacceleration, particles will undergo synchrotron oscillations, periodically moving from matched to unmatched momenta.

Although this effect is small, repeatedly inducing these so-called “chromatic” effects could eventually lead to a non-linear emittance growth due to effects such as filamentation. This highlights some of the difficulties in designing magnetic lattices of this type.

The overall performance of the beam emittance as it traverses the cooling channel with a Liquid Hydrogen absorber is shown in figure 2.8. It can be seen that despite the small mismatch due to the chromatic aberrations, the normalised emittance is well behaved and in fact decreases between the two




Figure 2.8: Behaviour of the normalised transverse emittance of a realistic muon beam with the MICE Step IV cooling channel. The centre of the liquid hydrogen absorber is marked in red and the 10 tracker stations are marked in green.

trackers.

There is a clear reduction in emittance due to the small beta function and absorber material at the centre of the lattice, as predicted by the cooling equation (equation 1.26). Additionally, each tracker station also provides a small amount of cooling due to the relatively low beta function in those regions and the relatively low-Z construction materials of each plane. Finally it can be noted that there exists some emittance growth through the regions of high focussing - or large changes in field gradient. The finite momentum spread leads to this chromatic effect, while the effects of energy straggling and energy loss means that despite a symmetric field, the focussing strength is not symmetric either side of the absorber. These effect could be further reduced by optimally tuning each side of the cooling cell individually based on slightly different momentum distributions.

2.4.2 The SciFi Trackers

The scintillating fibre (SciFi) trackers are the key measurement device within the MICE experiment. They permit the reconstruction of the position and momentum of individual muon tracks within a uniform solenoid field, and the reconstruction of the position and trajectory in configurations without magnetic fields. The body of the detector is constructed predominantly from carbon fibre, with the detector planes made out of scintillating fibres, laid to form a doublet layer, before glued in place and coated with Mylar. Figure 2.9 demonstrates

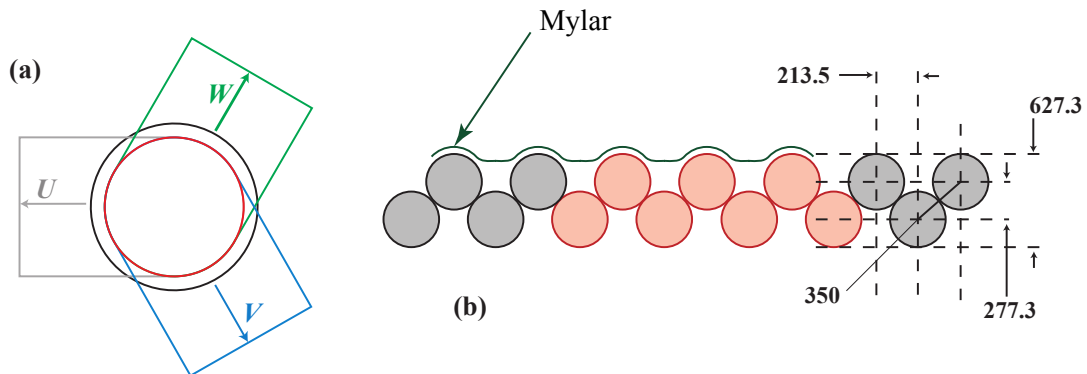


Figure 2.9: Schematic layout of the planes and how they are related to each other within a station. All measurements are in microns.

how the fibres are located with respect to each other. The scintillating fibres produce scintillation light from ionising radiation which is typically within the UV spectrum. Wavelength shifting fibres are connected to the end to transform the UV scintillation light into the visible spectrum. The wavelength shifting fibres are then connected in bunches of 7 to form a channel, and each channel is passed out of solenoid bore, via a light guide and into a Visual Light Photon Counter (VLPC).

The planes themselves are mounted in triplets (planes U, V and W), and rotated at 120 degree's to their neighbour, to form a station. The three orientations provide an x and a y component in the reconstruction. As only two planes are required to deduce an x - y position, the third planes serves as to reduce the probability that any recorded hits are due to noise by requiring a coincidence across each of the three planes. Figure 2.10 shows a completed tracker during the final stages of construction.

The read out for the detector is similar to the readout used in the D0 experiment [36] as both experiments make use of a scintillating fibre tracker and require a careful calibration system to calculate the number of photoelectrons produced by a minimum ionising particle based on the amount of charge generated in the VLPCs. Due to the relatively low numbers of photo electrons, the VLPCs must be very sensitive and are cooled by liquid Helium within specially designed cryostats. Due to the high sensitivity of the readout system, it is susceptible to noticeable electronics noise. The calibration was conducted such that during an event, on average each channel has a 2.3% probability of generating a 1PE noise digit. This makes the noise level noticeable and measurable, but is low enough that it can be distinguished from signal in the majority of cases. The electronics and calibration procedures are detailed in [37].




Figure 2.10: Photograph of a completed tracker. Note that to prevent discolouration of the scintillating fibres, they must be stored either in darkness or a yellow light. The 5 individual stations can be clearly seen in the image as the wavelength shifting fibres have not yet been connected to the edges of the planes.

2.5 Conclusion

The configuration of MICE permits a unique range of experiments to be conducted, with the goal of aiding the design of future muon accelerators and beam lines.

The material physics programme will provide data required to tune the models for mean and modal energy loss, and multiple coulomb scattering for the materials that are of greatest interest in cooling channels. This in turn will permit existing models of ionisation cooling to be improved.

The emittance change programme will provide the only experimental means to study ionisation cooling. The design of the cooling channel offers the flexibility required to study the cooling equation in detail, across a wide range of momenta and betatron functions. Such measurements are vital for future muon facilities.

The construction and commissioning of the cooling channel additionally provides a unique set of challenges. Few beamlines are composed entirely of superconducting solenoids and the procedures for installation, training and alignment are still being developed. MICE provides the first test of how feasible and complex such configurations are to install and use.

Finally, a beam physics programme is also under development. Due to the strongly-focussing solenoid magnets and high precision reconstruction devices, the SciFi Trackers, MICE is uniquely placed to be able to study the precise effect of strong focussing elements on the evolution of the beam emittance. Studies of non-linear emittance growth may be precisely conducted and analysed as a

function of momentum and field strength. Linear or quasi-linear models may then be constructed for use in future beamline models.

The outlined experimental programme and possible measurements for the MICE experiment are all heavily reliant on the performance of the MICE Scintillating Fibre Trackers. Systematic errors must be carefully controlled in order to produce the precise measurements required for the materials physics programme. Statistical errors and correlations must also be carefully modelled as they will directly bias the reconstruction of the beam emittance.

Chapter 3

MICE Track Reconstruction

3.1 Introduction

The development and testing of the track reconstruction software was a major component of the work conducted by the author. An existing codebase, MAUS [38], created by the MICE Collaboration was in place and tested, and contained the routines required for detector analysis and simulation, in addition to cooling channel simulations. The track reconstruction software was included within this framework. The main focus of the following section will be the theory and implementation of a Kalman Filter based track-fit algorithm.

The final track-fit follows on from lower-level reconstruction objects, which are created in sequence, beginning with the processing of the raw tracker DAQ data. The process of building these objects is described in section 3.3. The final track fit uses data from each plane, which is selected using a pattern recognition algorithm, before applying a Kalman Filter based fitting algorithm.

The low level reconstruction stages were implemented and tested by various collaborators and formed the required structure and processing framework on which the Kalman track fit could be implemented.

3.2 Geometry

In order to determine the parameters of the fitted tracks, in the global coordinate system of the experiment, a model of the tracker geometry was required for the reconstruction and simulation of tracker data. For the simulation, this was

Station	Upstream			Downstream		
	x [mm]	y [mm]	z [mm]	x [mm]	y [mm]	z [mm]
1	0.0	0.0	-549.9513	0.0	0.0	-549.8789
2	-0.4698	0.0052	-349.0577	-0.5709	-0.7375	-349.9143
3	-0.6717	-0.1759	-100.0523	-1.2021	-0.1657	-100.0513
4	0.1722	-0.2912	199.8771	-0.5694	-0.6040	199.9491
5	0.0	0.0	549.9513	0.0	0.0	549.8789

Table 3.1: Summary of survey measurements of the tracker stations made during the construction of the trackers. The centres of the first and last stations were used to define the coordinate system.

implemented as a specific reconstruction geometry file, written in a plain text format, and contained all the physical components of the trackers. Each fibre was individually modelled, the Mylar sheets that cover each plane were included as thin discs that stretch across the whole plane, additionally an estimate for the resin used to glue the fibres in place was also included. The ease of editing the geometry file permitted the effect of different materials, or material thicknesses to be investigated as required.

For the reconstruction however, a simplified geometry was used. In order to simplify the calculations of energy loss and multiple Coulomb scattering for each track, the scintillating fibre planes were modelled as solid planes of polystyrene. The effect of the Mylar sheets was also included, however the resin used in the construction was not. This vastly simplified the reconstruction, and provided a reconstruction rate that allowed for real time data analysis. Unfortunately, the reconstruction is sensitive to systematic errors due to discrepancies in the modelling of the precise material budget used in the tracker construction.

During the construction of each tracker, the stations were surveyed to a high precision to ensure that they could be precisely modelled. A summary is provided in table 3.1. The survey measurements were included in both the simulation and reconstruction geometries such that the tracks could be correctly reconstructed with respect to the other detectors.

The stations themselves were each identically constructed with the exception of the downstream station 5 which was constructed with an additional 3 channels. Table 3.2 summaries the relative positions of each plane within a station.

The trackers were installed within the bores of the two spectrometer solenoids such that the sensitive tracking volume was centred within the large centre coil, which provides the 4T magnetic field for use in momentum reconstruction. The magnets were carefully designed to provide the most uniform field as was cost

Plane	z [mm]	θ [Degrees]	No. Channels
0	-0.6398	-120	213
1	0.0125	+120	213
2	0.6648	0	211
2'	0.6648	0	214

Table 3.2: Summary of geometry of each plane within any given station. z is the longitudinal position with respect to the centre of the station and θ is the rotation of the plane around the tracker's central axis. Plane 2' refers to the anomalous plane in the downstream tracker, station 5.

effective (approximately uniform to 1%). However due to natural imperfections in the manufacturing process and the difficulty in maintaining a perfectly uniform field, early tests revealed non-uniformities on the order of 3% [39] which will have some systematic effect on the reconstruction parameters. More precise analyses that correctly model the thermal properties of the coils are underway in order to improve the field mapping.

3.3 Track Selection

3.3.1 Digitisation

The raw information received from the DAQ was composed of a list of all the tracker channels that registered some amount of scintillation light, which generated some non-zero Number of Photo-Electrons (NPE) in the VLPCs. Each channel was subjected to a NPE cut in order to remove electronics noise from the reconstruction. Due to the precision of the calibration, to pass the cut a channel was required to have at least 2 PE. If the channel passes the noise cut, it is saved within the reconstruction data structure as a Digit - the simplest reconstruction object. The NPE is stored, in addition to the raw signals from the front-end electronics and calibration, thereby allowing the reconstruction to access as much raw data possible when applying cuts and creating higher level objects.

3.3.2 Clustering

Clustering is the process where digits are examined in the context of the tracker plane, to see if there are any neighbouring digits, and if so, they are combined to form a cluster. Due to the geometry of the plane construction it is possible that a particle will produce a sufficient amount of light in two neighbouring channels,

it is therefore necessary to ensure both one-channel hits and two-channel hits are recorded and analysed.

The clustering algorithm is applied recursively, plane-by-plane, to a list of all the digits in that plane, ordered by their NPE. The result is that the digits with the largest signal take precedence, such that the clusters that are produced all have the largest signal possible, while only being constructed from a maximum of two digits. When more than two digits appear as neighbours, one is expected to be noise owing to the low percentage of multiple tracks recorded. The lowest signal digit would be promoted as a separate cluster.

The produced clusters, both one-channel and two-channel varieties, are then subjected to a second NPE cut if required, to further reduce the detector noise. By default this cut is not imposed, as it is statistically possible to generate low signal clusters. As the trackers are currently operating within the expected level of noise and background hits this second cut was not necessary for the processing of MICE Step IV data.

3.3.3 Space Point Production

Spacepoints are the next level in the reconstruction procedure, representing objects with 3 dimensional spatial information. Each station is considered in turn and examined to find sets of clusters from neighbouring planes that coincide. By combining the one-dimensional measurement from multiple planes, the two-dimensional x, y components of the spacepoint can be calculated using knowledge of the structure of the station. Spacepoints constructed from two clusters are referred to as “Doublets” and spacepoints constructed from three clusters are referred to as “Triplets”.

Kuno’s conjecture provides an efficient method to quickly determine if overlapping channels form a viable triplet spacepoint. It states that the sum of α (the distance of the cluster from the centre of the plane) for each of the three clusters will equate to zero if they simultaneously coincide with each other. This can be easily seen by projecting the position of any spacepoint (x, y) into the three different plane orientations as follows:

$$\begin{aligned}\alpha_1 &= \begin{pmatrix} x \\ y \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = x, \\ \alpha_2 &= \begin{pmatrix} x \\ y \end{pmatrix} \begin{pmatrix} \cos(120) \\ \sin(120) \end{pmatrix} = -\frac{1}{2}x + \frac{\sqrt{3}}{2}y, \\ \alpha_3 &= \begin{pmatrix} x \\ y \end{pmatrix} \begin{pmatrix} \cos(120) \\ -\sin(120) \end{pmatrix} = -\frac{1}{2}x - \frac{\sqrt{3}}{2}y,\end{aligned}\tag{3.1}$$

$$\alpha_1 + \alpha_2 + \alpha_3 = 0. \tag{3.2}$$

These ‘‘Kuno-Sums’’ are also useful for examining noise and ensuring correct cable mapping within a station (See section 5.3). In practice, the clusters have a channel ID, corresponding to the centre of the channel(s) that make up the cluster, rather than α . Hence it can be shown that a successful Kuno-Sum should equate half the total sum of channels.

The algorithm prioritises triplet spacepoints over doublet spacepoints as they allow a better rejection of noise-initiated clusters. The probability of noise creating a cluster is small, hence the probability of producing three coincidental noise clusters is negligible (<1 in 10^6) for most applications. The centre of the overlapping region is used as the estimate of the spacepoint position.

For the doublet-spacepoints, combinations of clusters from each unique pair of planes (0&1, 1&2 and 2&0) were analysed in turn. Each unique combination of clusters, that generate an x - y coordinate within the fiducial radius of the station, is accepted to produce a doublet-spacepoint.

Section 5.3 details the analysis of these simple reconstruction objects in Monte Carlo simulations and data.

3.3.4 Pattern Recognition

Pattern recognition is the stage where multiple spacepoints, distributed across a tracker are combined in an effort to determine if they were all formed from the same track. This is the first stage of the reconstruction where the magnetic field affect the reconstruction, as there is both a straight line pattern recognition algorithm and a helix pattern recognition algorithm, however the procedure is identical for both systems.

The spacepoints generated by the preceding stages are combinatorially selected to form candidate tracks, with precisely one spacepoint from each station, and triplet spacepoints taking priority. If there are multiple spacepoints in one or more stations, all possible combinations of spacepoints are examined in turn.

For helical tracks, there must be at least 4 spacepoints, with 5 being preferred. For straight tracks there must be at least 3 spacepoints, with 5 being preferred. The candidate tracks are then subjected to a linear least squares fit, the implementation of which is detailed in appendix A. A goodness-of-fit parameter, the χ^2 , is used to select the best fit track from all candidate tracks. Additionally a ‘‘road-cut’’, which excludes spacepoints that are perpendicularly too far from the fit, is also implemented to reduce the number of tracks containing noise hits and hard scatters.

The chi-squared requirements and road-cuts used in the track selection were determined through a series of trial and error studies using Monte Carlo simulations [40], in which the efficiency of the trackers was contrasted with the purity of the reconstructed tracks in order to obtain the optimum values.

The helical fit is prioritised over the straight line fit, when there is a field defined within the tracker geometry, due to a notable overlap in the acceptance criteria for the straight line fit and the helix fit with a low transverse momentum particle.

3.4 Kalman Filtering

The Kalman Filter [41] is an algorithm derived from signal processing applications and is an optimal, linear estimator that correctly takes into account all errors and correlations of measured data. For linear systems no fitting algorithm can perform better. For non-linear systems the Kalman filter provides only the optimal linear approximation, non-linear extensions to the Kalman may be used to improve the performance where necessary [42].

It has been used extensively in signal processing applications, and used to very good effect as an adaptive filter for GPS guidance systems, spacecraft monitoring, etc. Within the particle physics community, it has been augmented such that the usual independent variable, time, was changed to a spatial axis such that the “signal” being processed became the track parameters for a reconstructed particle [43, 44].

In order to implement a Kalman Filter, a consistent model of the system is required which must include:

- A model of the system dynamics - how one state may propagate into another state at a later time/position.
- Knowledge of how the measurement coordinate system relates to the state vectors
- Accurate estimates of the both the measurement noise, and the stochastic effects in the system dynamics.

Providing that the models used in the Kalman Filter accurately represent the system being measured, the results of the filter will include the correct correlations and uncertainties in all parameters.

3.4.1 Conceptual Description

The implementation can be broken down into three stages: Prediction, Filtering and Smoothing:

1. The Prediction stage estimates the system state (and state covariance matrix) at a subsequent position along the track, given the current state information, the system model and the expected process noise induced during the propagation.
2. The Filtering stage then makes use of the measurement at the next position to adjust the predicted state. The errors on the measurement and the prediction are used to weight the adjustment applied to the state.
3. The Smoothing stage occurs typically at the end of the fitting routine and uses the optimal estimate for the final state, which has accumulated the measurements from the previous filtering stages, and propagates in reverse to each of the preceding positions. The information acquired at all positions in the track is used to adjust each state to the optimal values given all the data obtained.

The result, following the smoothing stage, is that the state deduced at each position contains the optimal linear estimate for the system, using information extracted from all measurements. Figure 3.1 provides a cartoon demonstration of how the prediction and filtering stages may be considered in two spatial dimensions. Figure 3.2 provides a cartoon demonstration of how smoothing may appear in two spatial dimensions.




Figure 3.1: Conceptual behaviour of the Kalman Filter in the prediction and filtering stages. The blue arrow is a prediction made from a previous state with an error indicated by the blue circle. The measurement (red star) and associated error is used to form the filtered state and associated error (green arrow and circle).

In order to use the Kalman Filter as effectively as possible, the estimates for the noise, measurement and propagation functions must be as accurate as possible. The predominant difficulties during the implementation of the MAUS track fit was found to be the correct parameterisation of the sources of noise and non-linear measurement function.




Figure 3.2: Conceptual behaviour of the Kalman Filter in the smoothing stage. The black arrows represent the smoothed track, which is created by using the reverse-propagated filtered states (green arrows), compared with the measurements at each point (red stars). Owing to the gain of information, the earliest filtered states are more likely to be heavily corrected by this process than the later states.

3.4.2 The Algorithm

The algorithm of the Kalman Filter is universal to all applications, however in order to implement it, the following must be known for the system in question:

The Propagator

A matrix \mathbf{F} that linearly extrapolates the current state estimate to the next position, given the current position. It can be accompanied with a process noise vector \mathbf{w} with expectation values and covariance matrix defined by $\exp(\mathbf{w}) = 0$ and $\text{cov}(\mathbf{w}) = \mathbf{Q}$ respectively.

The Measurement

A matrix \mathbf{H} that transforms the current state estimate into the measurement space, given the current position. It can be accompanied with a measurement noise vector ε with expectation value and covariance matrix defined by $\exp(\varepsilon) = 0$ and $\text{cov}(\varepsilon) = \mathbf{V}$ respectively.

The system may now be modelled using the following equations:

$$\mathbf{x}_{k,t} = \mathbf{F}_{k-1}\mathbf{x}_{k-1,t} + \mathbf{w}_{k-1}, \quad (3.3)$$

$$\mathbf{m}_k = \mathbf{H}_k\mathbf{x}_{k,t} + \varepsilon_k, \quad (3.4)$$

$$\exp(\mathbf{w}) = 0 \quad \text{and} \quad \text{cov}(\mathbf{w}) = \mathbf{Q}, \quad (3.5)$$

$$\exp(\varepsilon) = 0 \quad \text{and} \quad \text{cov}(\varepsilon) = \mathbf{V}. \quad (3.6)$$

The true system state (denoted by the subscript t), \mathbf{x}_t at position (k) is propagated from the preceding state, ($k - 1$), in equation 3.3; and the current system state may be projected in to the measurement space using equation 3.4. This notation was used by [43] and will be used in the rest of this document owing to the conciseness and ease of viewing. Hence we similarly define:

$$\begin{aligned}
\mathbf{x}_k^n & : \text{ State estimate at position } k, \\
\mathbf{C}_k^n & : \text{ Covariance matrix of state vector,} \\
\mathbf{r}_k^n & : \text{ Measurement residual at position } k, \\
\mathbf{R}_k^n & : \text{ Covariance of measurement residual.}
\end{aligned} \tag{3.7}$$

The subscript, k , indicates the current position of the variable and the superscript n indicates the current stage of the information propagation; i.e. ($n < k$) indicates a predicted state, ($n = k$) indicates a filtered state (also written as $\mathbf{x}_k \equiv \mathbf{x}_k^k$), and ($n > k$) indicates a smoothed state.

The process of applying the three stages of the Kalman Filter is described in the following text, where the gain matrix formalism for the Kalman filter has been chosen. The alternative (the weighted means formalism) is mathematically identical, but more complex to implement (see [43] for more details).

Prediction

The current state vector (\mathbf{x}_k) and its covariance matrix (\mathbf{C}_k) are propagated to the predicted state (\mathbf{x}_{k+1}^k) and covariance matrix (\mathbf{C}_{k+1}^k) by:

$$\mathbf{x}_{k+1}^k = \mathbf{F}_k \mathbf{x}_k, \tag{3.8}$$

$$\mathbf{C}_{k+1}^k = \mathbf{F}_k \mathbf{C}_k \mathbf{F}_k^T + \mathbf{Q}_k. \tag{3.9}$$

Filtering

The Filtering stage is centred around the application of the Kalman Gain Matrix, which is calculated using the predicted covariance matrix and the estimated measurement noise. It can be considered as determining the relative weight to be applied to the measurement, given the accuracy of the prediction,

$$\mathbf{K}_k = \mathbf{C}_k^{k-1} \mathbf{H}_k^T (\mathbf{V}_k + \mathbf{H}_k \mathbf{C}_k^{k-1} \mathbf{H}_k^T)^{-1}. \tag{3.10}$$

The current state and covariance matrix can then be “filtered” by the following equations where the Kalman Gain Matrix is used to weight the adjustment applied from the measurement residual:

$$\mathbf{x}_k = \mathbf{x}_k^{k-1} + \mathbf{K}_k (\mathbf{m}_k - \mathbf{H}_k \mathbf{x}_k^{k-1}), \tag{3.11}$$

$$\mathbf{C}_k = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{C}_k^{k-1}. \tag{3.12}$$

Smoothing

The Kalman Smoother Gain Matrix (\mathbf{A}_k) determines how the information from the next state is propagated back to the current state, assuming it has already been filtered. It is calculated using both the filtered and predicted covariance matrices for the current state:

$$\mathbf{A}_k = \mathbf{C}_k \mathbf{F}_k^\top (\mathbf{C}_{k+1}^k)^{-1}. \quad (3.13)$$

The smoothed state and covariance matrix can then be calculated using the current filtered state and difference between the predicted and smoothed states at the next position. This form correctly propagates the information back to the current state:

$$\mathbf{x}_k^n = \mathbf{x}_k + \mathbf{A}_k (\mathbf{x}_{k+1}^n - \mathbf{x}_{k+1}^k), \quad (3.14)$$

$$\mathbf{C}_k^n = \mathbf{C}_k + \mathbf{A}_k (\mathbf{C}_{k+1}^n - \mathbf{C}_{k+1}^k) \mathbf{A}_k^\top. \quad (3.15)$$

At all stages in the algorithm the residual states and covariance matrices, required for the calculation of goodness-of-fit parameters (e.g. chi-squared and p-value - see chapter 4), can be calculated using:

$$\mathbf{r}_k^n = \mathbf{m}_k - \mathbf{H}_k \mathbf{x}_k^n, \quad (3.16)$$

$$\mathbf{R}_k^n = \mathbf{V}_k + \mathbf{H}_k \mathbf{C}_k^n \mathbf{H}_k^\top. \quad (3.17)$$

Weighted Pulls

Following the track fit procedure, the pull, π_k , for each reconstructed track-point may be calculated. The distribution of pulls, weighted by their respective errors, represents a measure of the whether the statistical properties of the fit are correct on the level of individual measurements. If the model correctly represents the data, the weighted pulls will be gaussian distributed with a variance of 1. This implies that the error associated with a trackpoint is correctly calculated with respect to the residual between the measurement and the estimate, and that the data is appropriately weighted within the filtering and smoothing stages.

The inverse gain matrix is first calculated in order to counter the filtering procedure, where the measurement noise (\mathbf{V}_k) is subtracted from the measured covariance matrix,

$$\mathbf{K}'_k = \mathbf{C}_k^n \mathbf{H}_k^\top (-\mathbf{V}_k + \mathbf{H}_k \mathbf{C}_k^n \mathbf{H}_k^\top)^{-1}. \quad (3.18)$$

The pull $\boldsymbol{\pi}_k$ and the associated covariance matrix $\boldsymbol{\Pi}_k$ of the current state k can then be calculated,

$$\boldsymbol{\pi}_k = \mathbf{x}_k^n + \mathbf{K}'_k(\mathbf{m}_k - \mathbf{H}_k \mathbf{x}_k^n), \quad (3.19)$$

$$\boldsymbol{\Pi}_k = (\mathbf{I} - \mathbf{K}'_k \mathbf{H}_k) \mathbf{C}_k^n. \quad (3.20)$$

The components of the weighted pull, $\boldsymbol{\rho}_k$, are then calculated using the estimated standard deviations of each component from the diagonal of the covariance matrix,

$$(\boldsymbol{\rho}_k)_i = \frac{(\boldsymbol{\pi}_k)_i}{\sqrt{(\boldsymbol{\Pi}_k)_{ii}}}. \quad (3.21)$$

3.4.3 The Extended Kalman Filter

The extended Kalman Filter differs from the standard implementation in that it features an analytic prediction of the track point, using an appropriate, non-linear propagator function. The prediction of the covariance matrix may then be performed using a first order expansion of the propagator function, about the predicted position.

If the state vector \mathbf{x} can be analytically propagated using

$$\mathbf{x}_{k+1}^k = \mathbf{f}(k, \mathbf{x}_k),$$

then the i th- j th component of the extended propagator \mathbf{F} can be calculated as the first derivative of the matrix function \mathbf{f} . This can be written as:

$$(F_k)_{ij} = \left. \frac{\partial f_i(k, \mathbf{x}_k)}{\partial (x_k)_j} \right|_{\mathbf{x}=\hat{\mathbf{x}}_k}. \quad (3.22)$$

The prediction of a subsequent track point can now be modified using the definitions of \mathbf{f} and \mathbf{F} ,

$$\mathbf{x}_{k+1}^k = \mathbf{f}(k, \mathbf{x}_k), \quad (3.23)$$

$$\mathbf{C}_{k+1}^k = \mathbf{F}_k \mathbf{C}_k \mathbf{F}_k^\top + \mathbf{Q}_k. \quad (3.24)$$

Following this step the filtering and smoothing procedures are as before, with the newly calculated \mathbf{F} matrix.

3.4.4 MAUS Implementation

Straight Tracks

The straight track system is parameterised by a 4-dimensional state vector, given by:

$$\mathbf{x} = \begin{pmatrix} x \\ m_x \\ y \\ m_y \end{pmatrix}, \quad (3.25)$$

where x and y are the spatial transverse coordinates of the track at the current position in mm, and m_x and m_y are the transverse gradients of the track at the current position.

The propagation is conducted by extrapolating the current state along a straight line using the usual Cartesian parameterisation. This is realised with the propagator matrix:

$$\mathbf{F}_k = \begin{pmatrix} 1 & \Delta z_k & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \Delta z_k \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (3.26)$$

where Δz_k represents the distance between the current tracker plane (at position k) and the next tracker plane (at position $k + 1$). Tables 3.1 and 3.2 provide the information required to calculate the values of z_k .

Helical Tracks

The helical track system is parameterised by a 5-dimensional state vector, given by:

$$\mathbf{x} = \begin{pmatrix} x \\ p_x \\ y \\ p_y \\ \kappa \end{pmatrix}, \quad (3.27)$$

where x and y are the spatial transverse coordinates of the track at the current position in mm, and p_x and p_y are the transverse momenta of the track at the current position in MeV/c and $\kappa = q/p_z$, where q is the particle charge in units of e and p_z is the longitudinal momentum in MeV/c.

The propagation is conducted by extracting the current state along the path

of a helix. The following equations of motion were used to parameterise the propagator matrix:

$$\begin{aligned}
x' &= x + \frac{p_x}{p_t} R \sin(\Delta\theta) - \frac{p_y}{p_t} R(1 - \cos(\Delta\theta)), \\
p'_x &= p_x \cos(\Delta\theta) - p_y \sin(\Delta\theta), \\
y' &= y + \frac{p_y}{p_t} R \sin(\Delta\theta) + \frac{p_x}{p_t} R(1 - \cos(\Delta\theta)), \\
p'_y &= p_y \cos(\Delta\theta) - p_x \sin(\Delta\theta), \\
p'_z &= p_z, \\
z' &= z + \Delta z,
\end{aligned} \tag{3.28}$$

where $\Delta\theta$ is the amount of rotation of the particle along the path of the helix and the prime-values of x' , p'_x , etc., mark the values at some arbitrary position along z , a distance Δz away. To simplify the implementation, the parameter u was introduced,

$$u = qcB_z = \frac{p_t}{R}, \tag{3.29}$$

where q is the particle charge in units of the electron charge, c is the speed of light in a vacuum, B_z is the solenoidal magnetic field strength and R is the radius of the helix. This simplifies the determination of $\Delta\theta$:

$$\Delta\theta = \frac{q}{p_z} \frac{p_t}{R} \Delta z = u\kappa\Delta z. \tag{3.30}$$

As the equations of motion are inherently non-linear, an extended Kalman filter must be implemented. To account for the non-linearity of the propagation, the state vectors are propagated analytically using the equations of motion (equation 3.28), ensuring the estimates are calculated correctly; while the covariance matrices must be propagated using a first order, linear approximation to the equations of motion. The extended propagator is given by:

$$\mathbf{F}_k = \begin{pmatrix} 1 & \frac{\sin(\Delta\theta)}{u} & 0 & \frac{\cos(\Delta\theta)-1}{u} & \Delta z(p_x \cos(\Delta\theta) - p_y \sin(\Delta\theta)) \\ 0 & \cos(\Delta\theta) & 0 & -\sin(\Delta\theta) & -u\Delta z(p_x \sin(\Delta\theta) + p_y \cos(\Delta\theta)) \\ 0 & \frac{1-\cos(\Delta\theta)}{u} & 1 & \frac{\sin(\Delta\theta)-1}{u} & \Delta z(p_x \sin(\Delta\theta) + p_y \cos(\Delta\theta)) \\ 0 & \sin(\Delta\theta) & 0 & \cos(\Delta\theta) & u\Delta z(p_x \cos(\Delta\theta) - p_y \sin(\Delta\theta)) \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \tag{3.31}$$

Measurement System

The measurement system describes how a state vector may be transformed into the coordinate system of an individual measurement. The measurements used are the clusters formed during the low-level reconstruction (see section 3.3), as this permits the effects of energy loss, energy straggling (stochastic variations in the rate of energy loss) and MCS to be modelled on a plane-by-plane basis, rather than station-by-station, which would theoretically improve the performance of the Kalman filter by more precisely modelling where the energy is lost.

Each measurement is a vector of length 1 and contains only the parameter α , the distance of the cluster from the centre of the plane. Hence the measurement matrix simply projects the spatial dimensions of the current state vector onto the orientation of the current plane. Although the straight track state vector has a length of 4 and the helical track state vector has a length of 5, only the x and y components are used, therefore the measurement matrices are almost identical for both cases.

$$\begin{aligned}
 \text{Plane V:} \quad \mathbf{H} &= \begin{pmatrix} \cos(120^\circ) & 0 & \sin(120^\circ) & 0 \end{pmatrix}, \\
 \text{Plane W:} \quad \mathbf{H} &= \begin{pmatrix} \cos(-120^\circ) & 0 & \sin(-120^\circ) & 0 \end{pmatrix}, \\
 \text{Plane U:} \quad \mathbf{H} &= \begin{pmatrix} 1 & 0 & 0 & 0 \end{pmatrix}.
 \end{aligned} \tag{3.32}$$

It can be noted that as Plane U has fibres that are parallel to the y axis, the projection into the measurement coordinate system is simply the extraction of x from the state vector. Planes V and W involve a rotation of $\pm 120^\circ$.

The measurement noise associated with the measurements is calculated assuming a uniform distribution of hits across the width of a channel, hence the measurement variance can be considered to be the second moment of a normalised square-function distribution, whose width is the channel width, w :

$$V = \int_{-w/2}^{w/2} \frac{x^2}{w} dx = \frac{w^2}{12}, \tag{3.33}$$

hence the measurement error as used in the reconstruction is given by: $\sqrt{V} = w/\sqrt{12}$. The nominal channel width for the MICE tracker is 1.4945 mm, hence the error on each measurement is ≈ 0.43 mm.

Energy Loss

The material within the tracker (helium between stations, scintillating fibres and Mylar sheets), must be accounted for during the propagation of a track. Two models of energy loss were considered during the development of the Kalman Filter, the Bethe formula for mean energy loss (equation 3.34) and the Landau-Vavilov formula for the most common energy loss (equation 3.35).

For an incoming particle, traversing some material, the mean energy loss per unit of areal density, $\langle dE/dx \rangle$, is given by,

$$\left\langle -\frac{dE}{dx} \right\rangle = \frac{Z}{A} \frac{K z^2}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right], \quad (3.34)$$

where K is a constant, $m_e c^2$ is the electron mass, β and γ are the usual relativistic parameters for the incoming particle, Z and A are the atomic and mass numbers of the material, z is the particle charge, W_{\max} is mean energy transferred to an electron in a single collision, I is the ionisation energy for the material and $\delta(\beta\gamma)$ is a density effect correction, normally calculated from data.

Conversely, for the same particle traversing some material, the most probable energy loss, Δ_p , is given by,

$$\Delta_p = \xi \left[\ln \frac{2mc^2 \beta^2 \gamma^2}{I} + \ln \frac{\xi}{I} + j - \beta^2 - \delta(\beta\gamma) \right], \quad (3.35)$$

where mc^2 is the mass of the particle, $j = 0.200$ derived from [45], $\xi = (K/2) (Z/A) z^2 (x/\beta^2)$ and the remaining symbols share the same meaning as for equation 3.34.

Through a sequence of tests, the Bethe formula was found to perform more accurately than the Landau-Vavilov formula. This is believed to be due to the inability of the track fit to correctly model hard scatters. In addition, due to the requirement of MICE to reconstruct statistical measures of the incoming muon beam, it is more logical to use an energy loss model better suited to the mean energy loss.

The Bethe formula was used to estimate the mean energy loss for each track depending on the material traversed during each propagation step. The calculated values were included as an additional set of coefficients during the calculation of the propagator matrix and the analytic propagation as for the extended Kalman Filter. Each material, enumerated by $i \in [1, N]$, was analysed in turn to estimate the mean energy loss, given the energy of the particle, at that point

in the propagation. The individual values were then combined as given by,

$$\Delta E = \sum_i^N \left\langle \frac{dE}{dx} \right\rangle_i \rho_i \delta_i, \quad (3.36)$$

where ΔE is the combined mean energy loss, $x_i = \delta_i \rho_i$ corresponds to areal density passed through by the particle, the product of the width of the material, δ , and the density, ρ . The fractional momentum loss can then be calculated for an assumed particle species. By default, all tracks are assumed to be muons for simplicity, hence the fractional momentum loss, μ , is given by,

$$\mu_k = \frac{\sqrt{(E_0 - \Delta E)^2 - m_\mu^2}}{p_0}, \quad (3.37)$$

where E_0 and p_0 are the initial energy and momentum respectively and m_μ is the muon mass. The helical propagator can then be updated by pre-multiplying by the matrix \mathbf{P}_k , where κ is correctly updated with the inverse of μ_k :

$$\mathbf{F}_k \rightarrow \mathbf{P}_k \mathbf{F}_k, \quad (3.38)$$

$$\mathbf{P}_k = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \mu_k & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & \mu_k & 0 \\ 0 & 0 & 0 & 0 & 1/\mu_k \end{pmatrix}. \quad (3.39)$$

Multiple Coulomb Scattering

The predominant stochastic effect is multiple Coulomb scattering, and is modelled using the results summarised by the PDG [18]. Moliere provided the initial theory, which was later incrementally improved by Highland *et al.* The best estimate for the root mean squared (RMS) angular deviation in trajectory that is projected onto a 2D plane, θ_0 , is given by,

$$\theta_0 = \frac{13.6 \text{MeV}/c}{\beta c p} z \sqrt{\frac{x}{X_0}} \left(1 + 0.038 \ln \left(\frac{x}{X_0} \right) \right), \quad (3.40)$$

where z , p and βc describe the charge, momentum and velocity of a particle respectively, travelling through a material with a thickness per radiation length of x/X_0 . The planar scattering distribution is related to the full spatial distribution by,

$$\theta_0 = \theta_{\text{plane}}^{\text{RMS}} = \frac{1}{\sqrt{2}} \theta_{\text{space}}^{\text{RMS}}. \quad (3.41)$$

The propagation noise, \mathbf{Q}_k , which is assumed to be entirely due to multiple Coulomb scattering for simplicity, can then be calculated using θ_0 . From the derivation of equation 3.40, it can be shown that the RMS trajectory deviation is given by $y = x\theta_0/\sqrt{3}$ with a correlation coefficient¹ of $\text{Corr}(y, \theta) = \sqrt{3}/2$. For a distance of d between measurement planes, where θ_0 is calculated for the mean value of x/X_0 , \mathbf{Q}_k is given by,

$$\mathbf{Q}_k = \mathbf{J}_k \begin{pmatrix} \frac{d^2\theta_0^2}{3} & \frac{d\theta_0^2}{2} & 0 & 0 \\ \frac{d\theta_0^2}{2} & \theta_0^2 & 0 & 0 \\ 0 & 0 & \frac{d^2\theta_0^2}{3} & \frac{d\theta_0^2}{2} \\ 0 & 0 & \frac{d\theta_0^2}{2} & \theta_0^2 \end{pmatrix} \mathbf{J}_k^T, \quad (3.42)$$

where \mathbf{J}_k is the required Jacobian matrix to transform from the natural (x, θ) coordinate system to either the straight track phase-space or the helical track phase-space. The straight-track phase-space is described by (x, x') , where $x' \sim \theta$ at first order, hence the straight track Jacobian is assumed to be the identity matrix.

For the helical phase-space, the Jacobian is a 5×4 matrix given by,

$$\mathbf{J}_k = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & p_z & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & p_z \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (3.43)$$

where the paraxial approximation has been used, i.e. that the transverse momentum is small compared to the longitudinal momentum, such that $p_x/p_z \sim p_y/p_z \sim \theta$.

3.5 Conclusions

The low-level reconstruction routines provided the necessary stages of noise rejection and hit-building required to optimize both the efficiency and purity of the final track fit. The rejection of noise-induced clusters and spacepoints is demonstrated in the next chapter. There are few improvements that could be made to this functionality.

A pattern recognition algorithm was implemented using linear least squares fitting routines in order to allow the process to proceed in real time, however

¹Correlation coefficients are defined by the covariance between the two variables, divided by the product of standard deviations: $\text{Corr}(x, y) = \text{Cov}(x, y)/\sigma_x\sigma_y$

the models of energy loss and multiple coulomb scattering were not included, meaning that the chi-squared and road cut parameters required carefully tuning. Further improvements based on a combined, non-linear helical fit, rather than the existing circle+linear fit, have been considered but only as a low priority.

The Kalman Filter based track fit was anticipated to work exceptionally well for straight tracks. The models for multiple coulomb scattering typically describe a Gaussian effect that agrees with the central 98% of the population. Hence the only component that is not closely approximated by a Gaussian distribution is the measurement noise. The use of additional statistical discriminators would therefore remove the outlying events, leaving a statistically accurate fitting routine.

The helical track fit was conceivably more likely to bias the fitted parameters. In addition to the non-Gaussian measurement noise, the propagation of the covariance matrix between stations is inherently non-linear, even though the state vector is propagated analytically. The assumption was made that these effects would be small and would not significantly bias the fitted parameters. Further improvements in order to reduce any possible biases would involve either further extensions involving second order propagation and measurement routines or a non-linear fitting algorithm. Both options would require a significant amount of study and are currently viewed as potential upgrades.

Different parameterisations were also considered in an attempt to improve the modelling of the longitudinal momentum. It was found that the implemented Euclidean-style parameterisation was the most efficient. Different parameterisations nearly always involved highly complex approximations in order to estimate the process noise, hence they were discarded.

The current implementation of the track reconstruction routines was believed to be the most sensible given the time for development and testing, and level of complexity. The simulated and reconstructed performance of the algorithms are discussed in detail in the following two chapters.

Chapter 4

Validation of Track Fitting Routines

Simulation studies were performed in order to provide a detailed analysis of the performance of the Kalman Filter. The resolution of the track fit in each of the reconstruction coordinates was calculated and the efficiency at which the track finding algorithm successfully identifies tracks was estimated using the MC data. Additionally, estimates for the statistical and systematic errors present in the higher-order analysis, e.g. emittance and beta function calculations, were obtained.

4.1 Introduction

The initial stage of validation was to ensure that the statistical measures of the Kalman Filter agreed with the theoretical values for a correct implementation. The fit chi-squared value provides a simple and effective test of the implementation and the associated p-value can be used to compare the global performance of the system model.

4.1.1 Chi-Squared Distributions

For a system of N independent, normally distributed, random variables (z_i), the sum of their squares is defined to be the chi-squared parameter:

$$\chi^2 = \sum_{i=1}^N z_i^2.$$

Repeated sampling of a chi-squared value will produce a distribution defined to be the chi-squared distribution, with a shape that depends only on the number of degrees of freedom within the system, N .

In practice the system in question is typically under analysis, hence the parameters that describe the system must be estimated. For a simple system with only one parameter, the mean of the distribution, the more common form of the χ^2 definition is therefore,

$$\chi^2 = \sum_{i=1}^N \frac{|\mathbf{z}_i - \bar{\mathbf{z}}|^2}{\sigma_i},$$

where $\bar{\mathbf{z}}$ is the estimate for the mean and σ_i is the error on the residual. In this case, repeated experiments would generate a chi-squared distribution corresponding to $N - 1$ degrees of freedom. A degree of freedom is lost by imposing an additional constraint on the system - the calculated value of the mean of the distribution.

In general it can be deduced that for a given system definable by M parameters, with N measurements, the Number of Degrees of freedom is given by $N - M$.

4.1.2 P-Value Distributions

A simple way to determine how close a measured distribution is to a standard chi-squared distribution is to analyse the p-value for each experiment. The p-value (P) is calculated as the probability of obtaining a fit with a chi-squared worse than the one currently obtained, if the experiment were to be repeated.

This probability can be calculated by performing the appropriate integral. If the distribution is assumed to have a chi-squared distribution with N degrees of freedom i.e. $f(x) = \chi_N^2(x)$, and a given experiment produced a chi-squared value of X , the p-value would be calculated by:

$$P = \int_{-\infty}^X f(x)dx \quad P \in (0, 1). \quad (4.1)$$

If the distributed values of X are perfectly modelled by the function $f(x)$, the resulting distribution of p-values will be flat. The qualitative ‘‘flatness’’ of the p-value distribution is generally an excellent indicator for the appropriateness of the model that was used.

4.2 Algorithm Validation

The validation of the implementation of the algorithm was conducted using multiple Monte Carlo simulations, based around a simplified geometry. In order to exclude unwanted effects due to field uniformity, misalignments and poor transmission, the geometry was greatly simplified, such that it used only perfectly aligned components within an ideal solenoid field. The beam that was used was also configurable such that the input phase-space could be changed, and the physics effects could be controlled (MCS, energy loss, energy straggling, etc). This allowed the fit performance to be examined as it was developed.

The simulations were performed using MAUS, which uses Geant4 [46] to model the particle physics. Due to the chosen implementation, additional functionality was required to test the tracker reconstruction, without energy loss being modelled. The detectors generate a “measurement” based on the amount of energy a particle deposited within a sensitive virtual volume. If energy loss is not permitted, no detector hits are recorded. In order to combat this issue, a function that bypasses the standard detector digitisation was required, which made use of the direct tracking information available from Geant4. Such a function was created and tested successfully.

The first stage of analysis was to produce distributions of the p-value found from each track, thereby allowing a simple qualitative measure of the performance of the algorithms. The track reconstruction was performed on a simulated sample of muons with a uniform distribution of both longitudinal and transverse momentum components. Both the straight and helical reconstruction algorithms were tested where different physical effects were both modelled and accounted for. This was found to highlight the components of the system that caused the design to deviate from the ideal linear model. These deviations could lead to systematic discrepancies or an increase in the resolution of the fitted parameters.

4.2.1 Energy Loss

The stochastic effects of the Monte Carlo simulations were excluded from the simulations. In addition, the standard measurement system was bypassed in order to prevent the non-gaussian nature of the measurement errors affecting the fit performance (see section 3.4.4). An algorithm that simulates the measurements with an ideal gaussian error was used instead.

The only physical effect, modelled within the simulation, was mean-energy loss. This allowed for a direct test of the basic propagation routines, in addition




Figure 4.1: The reconstructed p-value distribution for a simulation of straight (left) and helical (right) muon tracks, where the only physical effect is the mean energy loss.

to the Kalman filter algorithm. If the model correctly describes the data, and the implementation is correct and linear, the produced distributions would be perfectly flat.

Figure 4.1 shows the p-value distribution for 20000 events in a single tracker for both straight and helical tracks. The straight tracks produced a distribution which was nearly flat, however there exists some small discrepancy between model and data that has not been accounted for, it is believed that the cause is the seed of the Kalman fit. As pattern recognition is used to derive the initial seed state, the fit is biased towards an accurate track. In doing so, the p-values are biased high as the fit results are generally slightly better than expected.

The helical tracks distribution demonstrates a much stronger deviation from the ideal flatness. The causes are more difficult to define as the helix model itself is not linear. As for the straight track reconstruction, there is some bias from the seed values towards an accurate track, however the non-linearity in propagator will also affect the final fit parameters. As both the prediction and smoothing steps in the reconstruction (see section 3.4) make use of the linear approximations, the final fit results could never form the ideal distributions. It is demonstrated in section 4.3 that there is little bias in the means of the fit results due to this discrepancy, and the biases that are found can be modelled and corrected for. The predominant effect of this discrepancy will be the modelling of the errors associated with each reconstructed trackpoint.




Figure 4.2: The reconstructed p-value distribution for a simulation of straight (left) and helical (right) muon tracks, where the only physical effect is the measurement system.

4.2.2 Measurement Effects

The second stage of testing was to examine the effect that non-gaussian measurement uncertainties have on the final fit parameters. Other physical effects were still neglected. Figure 4.2 shows the resulting p-value distributions for a system where the mean energy loss and the measurements are correctly modelled, but all other physical effects (MCS and energy straggling) are neglected.

The distributions were found to deviate more strongly from the ideal distributions and include the discrepancies found previously (figure 4.1). The discretisation of the channels would prevent a “perfect” fit from appearing as the clusters are always quantised by channel number. Hence there is never a fit with a very high p-value. Additionally, as the width of the channel is finite, it is unlikely that measurements that would typically appear in the tails of the gaussian are included in the fit, hence the lower p-values are also reduced.

4.2.3 Stochastic Effects

The stochastic effects were the final component to verify. The same measurement system as described in section 4.2.1 was used to remove the effects of non-gaussian measurement errors, leaving mean energy loss, energy straggling and multiple Coulomb scattering as the only physics processes that were modelled. Figure 4.3 shows the p-value distributions for both straight and helical track reconstructions for this simulation.

The stochastic effects can be seen to have the most significant effect on the p-value distributions. This implies that the model used to estimate the amount of scattering between each plane is over estimating the effect of the material.




Figure 4.3: The reconstructed P value distribution for a simulation of straight (left) and helical (right) muon tracks, where the only physical effects are multiple Coulomb scattering and energy straggling.

The majority of tracks are biased towards high p-values which implies that the fit is performing better than expected given the noise estimates.

Geant4 uses the Lewis Theory [47] to model scattering on a particle transport level, which differs from the typically assumed Highland Formula [48]. Additionally, it is known that scattering distributions are not perfectly gaussian, as was assumed for the reconstruction, and in fact deviate by up to 11% from data [18]. It is conceivable therefore that the differences between the two models cause the p-values to be biased higher than anticipated.

There was not sufficient time to implement a detailed study of the differences between the two models, with a specific focus on the materials and muon energies assumed within the reconstruction. Following the analysis of the MICE Step 4 data, which anticipates improving the current models of both energy loss and multiple Coulomb scattering, the approximations used in the reconstruction may be updated to better reflect the physics involved. However there will still exist some discrepancies as the Kalman Filter inherently assumes a gaussian distribution for the noise, which only approximates multiple Coulomb scattering.

The first bin in the histogram, for both straight and helical tracks, has a significantly greater number of entries than the nearest neighbours. This suggests that there is some effect causing a small number of tracks to be reconstructed more poorly than expected. This can be well explained by the non gaussian nature of multiple Coulomb scattering and energy straggling, in that the longer tails are not correctly modelled by the Kalman Filter. The result is the excess of tail events are added to the lowest bin(s). If the binning is increased a smooth peak can be resolved assuming sufficient statistics are generated.




Figure 4.4: The reconstructed P value distribution for an simulation of straight (left) and helical (right) muon tracks, where all physical effects were modelled in the reconstruction.

4.2.4 Full Reconstruction

The p-value analysis was completed with a full reconstruction simulation. All stochastic physical effects were modelled and the correct digitisation routines were used. Figure 4.4 demonstrates the straight track distributions and helical track distribution. These results are a non-trivial combination of the results found for all previous stages. As such the optimal “flat” p-value distribution is no longer achievable due to the range of approximations and assumption made throughout the development of the algorithms.

The p-value is typically a powerful selection criteria, owing to the ease at which poor fit results may be objectively discarded. Typically a cut of 0.05 (corresponding to the 5% worst fit tracks) is implemented, however in a distinctly non-uniform distribution, a value of 0.05 no longer corresponds to 5%. As there is a local excess of values in the first bin due to poor modelling of tail events, it is difficult to justify this broad selection criteria. However a tighter p-value requirement may still be used to remove those tracks that are mis-reconstructed, in practice and cut on $P < 0.01$ was found to improve the samples used in analyses (see chapter 5).

4.3 Monte Carlo Simulation Performance

The entire reconstruction chain was tested using a complete simulation, including mean energy loss, stochastic effects and the non-gaussian measurement errors. This study provides the benchmarks for the track reconstruction resolution in all components, and gives a good indication of the behaviour expected




Figure 4.5: The $x - y$ position residuals for both straight tracks (left) and helical tracks (right) at the tracker reference plane.

for real data reconstruction. The Monte Carlo true track parameters were compared against the reconstructed track parameters on a track-by-track basis, such that the resulting residuals could be histogrammed and analysed.

The straight and helical track analyses were both conducted using the same geometry as described in the previous section. The beam was composed entirely of muons that were not permitted to decay with both a uniform longitudinal momentum distribution and a uniform transverse momentum distribution. For both beams, the longitudinal momentum was defined by, $p_z \in [140.0, 260.0]$ MeV/c. For the straight track analysis, the transverse component was defined by, $p_{\perp} \in [0.0, 20.0]$ MeV/c, similarly for the helical track analysis the transverse component was defined by, $p_{\perp} \in [0.0, 180.0]$ MeV/c.

The analysis of each trackpoint was conducted at the reference planes of the trackers, both upstream and downstream, as these are the nominal reconstruction planes for the experiment. The reconstruction at all other planes is performed and analysed by default, but not included for brevity.

4.3.1 Transverse Reconstruction

Figure 4.5 describes the position residuals for both straight and helical tracks. The residuals are compared at the reference frame of the upstream tracker. The position resolution of the final track can be estimated by assuming that the residuals form an approximate gaussian distribution about zero. The standard deviation of the fitted gaussian then corresponds to the position measurement resolution.

For straight tracks, the angular resolution of the trackers can be calculated in a similar fashion. Figure 4.6 demonstrates the $x - y$ angular resolution of the




Figure 4.6: The transverse gradient residuals (left) and transverse momentum residuals (right) of the upstream tracker.

upstream tracker at the reference plane. As before, the resolution is estimated using a fitted gaussian.

The transverse momentum resolution is more complex to analyse as the helix fit performance depends on the momentum of the track. At low values of transverse momentum, the radius of the helix becomes very small and is comparable to the width of the multiple Coulomb scattering distribution, coupled with the natural position resolution of the trackers. Due to this, the reconstruction of these tracks becomes much less accurate and results in an increase in resolution. This was found to have the most pronounced effect on the reconstruction of the longitudinal momentum and can be seen in figure 4.7. Both the p_z and p_\perp residuals were binned by the true value of p_\perp . Each bin was then fitted with a gaussian function and the standard deviation, representing the resolution for a specific p_\perp bin, was plotted on the y-axis. The vertical error bars represent the statistical error on the fitted gaussian, and the horizontal error bars describe the bin width.

It can be noted that the transverse reconstruction is very stable across all values. There are some fluctuations away from a uniform resolution, the causes of which have not yet been identified, however the fluctuations are small enough to not have an effect on the final reconstruction of the parameters of interest for MICE (emittance and beta function). The longitudinal distribution is however much more sensitive to the true transverse momentum. The increase in resolution at low- p_\perp will decrease the precision possible for single track analyses. However, any analyses that uses a sufficiently large ensemble of particles will still be able to make a precise momentum estimate with high enough statistics. Additionally, the inclusion of the MICE TOF and EMR detectors (see chapter 2) may be used to improve the estimate of the longitudinal momentum via




Figure 4.7: The longitudinal momentum resolution (left) and transverse momentum resolution (right) as a function of transverse momentum.

a track-matching algorithm.

4.3.2 Momentum Reconstruction

The longitudinal momentum can be seen to be the least precise reconstruction parameter for the Kalman track fit. This is due to the choice of model and the non-linearity of the p_z dependency in the propagation steps. Pattern Recognition (PR) may be discussed in contrast as the PR model performs a fit in the s - z plane (s is the helical path length), where the deviations in p_z are linear with respect to the state space, for this reason, pattern recognition may actually be used to more precisely estimate the longitudinal momentum. However, in order to do so, the helix is decoupled in the transverse and longitudinal components, such that the transverse fit may be performed first, and used during the longitudinal fit.

The Kalman fit however, does not permit the longitudinal and transverse components to be decoupled as the calculation of MCS and energy loss requires the total momentum, which inherently couples the components. This also permits a more precise level of measurement information, the clusters rather than the spacepoints, to be used. Due to this precise modelling of the stochastic effects, the p_z component is only accessible through first order corrections in the trigonometric functions of the extended propagator (see section 3.4.3 for details). This reduced sensitivity causes the increase in the longitudinal resolution of the Kalman track fit. The seed for the Kalman fit is therefore more critical, as it can have a significant bias on the reconstructed longitudinal momentum.

As pattern recognition does not model energy loss or any stochastic processes, the evaluated seed momentum will be systematically incorrect. Figure 4.8

shows how the total track momentum, as estimated by pattern recognition, deviates from the true momentum at the seed position. This is due primarily to two effects, firstly there is an amount of energy loss (downstream tracker), or energy gain (upstream tracker) between the centre of the fit and the seed position, and secondly there is a systematic effect due to using an approximate model.




Figure 4.8: The deviation of the pattern recognition total momentum from the true total momentum at the Kalman seed position in the upstream tracker (left) and downstream tracker (right).

This discrepancy was reduced through the use of two constants, the approximate momentum loss (a) and a systematic bias (b). From fits to the Monte Carlo simulation data in figure 4.8, it was found that the total momentum, as determined by pattern recognition, required a correction described by equation 4.2. The values for a and b are found when the residuals in figure 4.8 are simultaneously corrected to zero.

$$\begin{aligned}
 p^c &= p + a + b && \text{Upstream,} \\
 p^c &= p - a + b && \text{Downstream.}
 \end{aligned}
 \tag{4.2}$$

In calculating and applying this correction, the systematic influence of pattern recognition is dramatically reduced over the momentum range of interest; however, it is not completely removed. The remaining bias of pattern recognition couples with the small systematic uncertainty in the energy loss calculation, and can be seen in the total momentum residuals, shown in figure 4.9.

These residuals correspond to the most significant bias in the track reconstruction and are of critical importance for any analysis. They are most efficiently examined as a function of total momentum, rather than longitudinal momentum, as the predominant cause is the correct calculation of the mean




Figure 4.9: The deviation of the total track momentum, evaluated at both the upstream (left) and downstream (right) reference planes, from the true track momentum.

energy loss which is a function of the total momentum. The residuals in the seed, although corrected, are still a function of momentum. The ideal correction would involve some approximation of the Bethe formula, but there wasn't sufficient time to investigate this concept. Additionally, there are known discrepancies between the material model in the reconstruction, the Monte Carlo model and the true geometry, which lead to a systematic mis-estimate of the energy loss per plane.

The combined effect is that of a momentum-dependent bias in the total momentum, which can be well approximated with a linear fit. Ideally the bias would be modelled using some coefficients derived from the Bethe formula, but for simplicity this was not investigated. As 200 MeV muons are very close to minimum ionisation, the energy loss is well approximated at first order, hence a linear fit was the preferred, simpler option.

Each plot in figure 4.9 was subjected to a first order polynomial fit, such that the intercept c and gradient, g may be used to approximately parameterise the residuals,

$$\Delta p = c + gp^{\text{True}}. \quad (4.3)$$

The inverse of this equation can then be used to correct the reconstructed total momentum of each track.

4.3.3 Track Finding Efficiency

The inherent reconstruction efficiency was analysed using the datasets that were generated for the resolution analysis, discussed in the previous sections. For every simulated particle that traversed a tracker, a reconstructed track was expected. For the straight track reconstruction, the limiting factors were the




Figure 4.10: Efficiency of the upstream (left) and downstream (right) helical track reconstruction algorithms as a function of transverse momentum. The inefficiency in the lowest bin is due to the difficulties of reconstructing tracks with low transverse momentum.

selection criteria placed on the track during pattern recognition. This is a small effect leading to better than 99.99% track finding efficiency for the straight track algorithm. Only tracks that experience extremely hard scatters can fail the reconstruction.

The helical reconstruction however, was less efficient and more sensitive to the condition of the beam. A preliminary study was performed (results shown in figure 4.10), however both data-driven and Monte Carlo based analyses are still ongoing at the time of writing [49, 50].

The current results suggest that uniform reconstruction efficiency is present across most of the phase-space, with a small reduction for the lowest values of transverse momentum. There will be some small bias for absolute emittance measurements, as the centre of the phase-space will be slightly under represented in analyses. The effect on cooling measurements however, will be negligible as each muon track is required to be reconstruction in both trackers, hence any bias due to the track selection will only affect the absolute emittance that is measured, not the relative change.

4.4 Conclusions

Tables 4.1 & 4.2 describe the bias and resolution of all the reconstruction parameters available from the Kalman track fit for the straight and helical track reconstruction respectively. Table 4.3 describes the corrections to the track fitting procedure in order to optimize the helical track reconstruction.

Overall the implementation of the Kalman filter has been a success, the simplest

case performs very close to the theoretical ideal, with the ability to test various individual aspects using different configuration flags. The study of how each component affects the final fit parameters revealed the full extent of the non-linear features of the full implementation and allowed for them to be carefully studied.

Despite the non-uniform p-value distributions, they may still be used for their intended purpose - a quantitative measure of the goodness of fit for a single track - as long as the requirement is relaxed. Care must therefore be exercised when comparing the track reconstruction to other detectors and in analyses to ensure that the poor fits are correctly accounted for. Due to the tight requirements from pattern recognition for forming a track this is not anticipated to cause an issue.

The primary improvement to the Kalman track fit would be a better model of the multiple Coulomb scattering experienced by each track. This would improve the distribution of the p-values and hence allow for a more precise determination of the parameter errors, as discussed in section 3.4.4.

There would still, however be issues with the non-linear and non-gaussian components of the model. The measurement errors, for example, would require an additional correction in order to account for the non-gaussian errors. It would be possible, but not trivial, to include the kurtosis of the assumed measurement error within the filtering stage, providing a better approximation to the true top-hat distribution. Additionally, non-linear smoothing and alternative system models could also be considered and tested in case they are found to perform noticeably better than the current implementation.

Monte Carlo studies have allowed the resolution of the tracker to be modelled for both straight and helical tracks across the entire region of interest in both transverse and longitudinal momenta. The final fit parameters were found to perform within initial expectations and were in good agreement with the Monte Carlo truth. The remaining systematic effect of the total momentum reconstruction has been identified and corrected using a linear approximation. A more precise treatment of such biases would theoretically improve the overall reconstruction resolution.

The influence of pattern recognition on the final fit was found to be a significant contributor to the systematic bias of the total momentum and required careful modelling. Simple corrections for the both the PR seed and the final track fit were calculated, and found to improve the residuals within the Monte Carlo study. These studies assumed ideal conditions within the trackers and do not account for any field misalignments and non-uniformity. The determina-

tion of the momentum correction is repeated in section 7, where the emittance reconstruction requires precise modelling of the reconstruction in a non-ideal field.

In addition to the improvements outlined for the implementation of the basic algorithm, a better model of the tracker could also be included using:

- A more detailed pattern recognition correction;
- An adaptive Kalman Filter to correct the energy loss per-plane.

The pattern recognition correction would involve a momentum-dependent factor, likely a scaled implementation of the Bethe-formula, such that the momentum bias in the seed can be almost completely removed, barring some uniform systematic. This is anticipated to reduce the systematic momentum bias in the final track fit.

Finally, an Adaptive Kalman Filter (see [42]) would permit the energy loss per-plane to be fitted per-plane, hence there would be less requirements on the actual modelling of the tracker materials, as the overall momentum-dependent energy loss could be measured by the tracker, rather than assumed. This would increase the complexity of the fit and likely require a restructuring of the system model such that the p_z component is more easily accessible within the measurement procedure. There are several potential implementations and the result would permit an energy loss calculation that is predominantly data-driven. This would reduce the dependency on Monte Carlo studies for momentum corrections to the track fit algorithms.

	Upstream		Downstream		Units
	Bias	Resolution	Bias	Resolution	
x	0.0	0.247	0.0	0.224	mm
y	0.0	0.202	0.0	0.197	mm
m_x	0.0	0.004	0.0	0.003	N/A
m_y	0.0	0.004	0.0	0.003	N/A

Table 4.1: A summary of the reconstruction parameters produced from the straight track Kalman fit. The values are averaged across uniform phase-space which spans 140 to 260 MeV/c in p_z and 0 to 20 MeV/c in p_\perp .

	Upstream		Downstream		Units
	Bias	Resolution	Bias	Resolution	
x	0.0	0.213	0.0	0.217	mm
y	0.0	0.207	0.0	0.201	mm
p_x	0.0	0.884	0.0	0.874	MeV/c
p_y	0.0	0.876	0.0	0.889	MeV/c
p_z	-0.097	1.64	0.458	1.63	MeV/c

Table 4.2: A summary of the reconstruction parameters produced from the helical track Kalman fit. The values are averaged across uniform phase-space which spans 140 to 260 MeV/c in p_z and 0 to 180 MeV/c in p_\perp .

Correction	Formula	Parameters
PR Momentum	$p^c = p \pm a + b$	$a = -0.156 \text{ MeV/c}$ $b = -1.157 \text{ MeV/c}$
Upstream Total Momentum	$\Delta p = c + gp^{\text{True}}$	$c = -0.95 \pm 0.10 \text{ MeV/c}$ $g = 0.0035 \pm 0.00053$
Downstream Total Momentum		$c = 2.08 \pm 0.23 \text{ MeV/c}$ $g = -0.0088 \pm 0.0011$

Table 4.3: A summary of the correction functions required to optimize the results from the Kalman Track Fit.

Chapter 5

Validation of Reconstructed Data

5.1 Introduction

During the initial stages of the tracker commissioning, the first data taken were used to verify the installation of both the upstream and downstream trackers, and the reconstruction routines. Several runs comprising of different particle species, momenta and beamline optics were studied to ensure a good performance across the measurable phase-space. The primary goal was to verify the installation of the fibres and cables to the VLPCs, and validate the data readout was working correctly.

Following on from these initial runs, the MICE Step IV physics run plan came into effect. The primary runs selected for analysis in this and subsequent chapters were muon beam runs of 2 different momenta, 200 and 240 MeV/c, designed to examine the alignment of the trackers (no absorber material) and to measure the multiple Coulomb scattering distributions for different materials.¹ An additional run with the upstream spectrometer solenoid magnet powered was also analysed, as it was the only available run to test the helical track reconstruction. It is discussed in detail in chapter 7.

The runs that are presented below were selected following a program of improvements and corrections to both hardware and software, in an attempt to commission the trackers at full specification. These data represent only a fraction of the calibration and commissioning data sets required to fully commission

¹At the time of writing the analysis of these data is underway and expected to be published in 2016–2017 by the MICE Collaboration.

the detectors and analyse their performance.

Table 5.1 describes the run numbers of interest during the period of data taking from September 2015 to March 2016. The data taken were predominantly for the use in the multiple Coulomb scattering study, hence there were no magnetic fields. The runs that make up a control data set, where no absorber was present, were preferred, as they allowed for a direct tracker to tracker alignment to be conducted in parallel to the primary analysis. The final run, 07469, was recorded during the initial stages of the magnet commissioning and permits the first validation of the helical track fitting routines, in addition to the high-precision emittance measurement for the accepted muon beam.

Dataset	Run Number	Momentum	Beamline	Description
A	07672 07673 07681 07695	200 MeV/c	3 mm	No Absorber Scattering Study
B	07674 07682 07685 07686 07691 07693 07694	240 MeV/c	3 mm	No Absorber Scattering Study
C	07469	200 MeV/c	3 mm	Helical Reconstruction

Table 5.1: List of the runs that were used in this and subsequent chapters. The momentum value refers to the nominal beam momentum at the centre of the cooling channel, and the beamline value refers to the nominal beam emittance at the entrance to the cooling channel. These two values allow for the correct beamline settings to be identified from a predefined catalogue.

5.2 Beam Selection

To ensure that the analyses were conducted using only non-decaying, muon tracks that were fully contained within the experimental volume, several cuts were applied to the high-level reconstruction objects. These higher level cuts are indicative of the sort of selection criteria that could be determined from more detailed Particle-ID and global track fitting studies. Hence, the data analysed below describes the reconstruction performance for a dataset representative of more detailed analyses. Dataset A was selected for the low-level analysis described in this chapter, in addition to the high-level straight track validation. Dataset C was selected for the specific helical track validation in section 5.4.

A coincidence cut was required such that precisely one track was reconstructed in each of the two trackers, and precisely one TOF spacepoint was reconstructed in each of the TOF1 and TOF2 detectors. This was to reduce the number of decaying muons that were reconstructed and to ensure that no multiple track events were analysed. As the reconstruction was not constructed to correctly treat multiple muon tracks at the time of writing, this greatly simplified the reconstruction process.

A time of flight cut was imposed to ensure that the selected events were almost entirely composed of muon tracks. It was previously shown that the pion contamination of muon beam is $< 1\%$ [33], however there was still some contamination due to electrons which was easily removed. As TOF 1 and TOF 2 mark the start and end of the cooling channel, the time of flight cut was imposed on the difference between these two detectors. Figure 5.1 demonstrates the TOF seen for the 200 MeV/c muon beam, in addition to the cut that was applied,

$$28.5 < \text{TOF}_{1 \rightarrow 2} < 34.0 \quad [\text{ns}]. \quad (5.1)$$

Note that the electron peak can be seen at approximately 27.5 ns, and the tails of the muon peak were removed to reduce the contamination from slow decay electrons and to ensure that a precise central momentum was retained. No other criteria were required for the low-level validation routines. For the higher-level objects, additional cuts on p-values and fit parameters could be applied, however those are discussed in the appropriate sections.

An additional Monte Carlo dataset was also required in order to fully validate the reconstruction. A model of the inherit detector noise was included in the simulation, in addition to the effects of known dead channels and the calibration. In general some channels have a decreased efficiency of producing digits as they are either damaged, or not able to be correctly calibrated (see section 2.4.2 for more detail on the calibration procedure). Additionally, there are some channels that are more susceptible to producing noise digits. All these effects have been approximately modelled in an attempt to better simulate the reconstruction environment of the trackers [49].

More work is required in order to better tune the noise models to reproduce the distributions seen in data. Only a first approximation has so far been included, based on the measurements taken from a test on a single station [51].

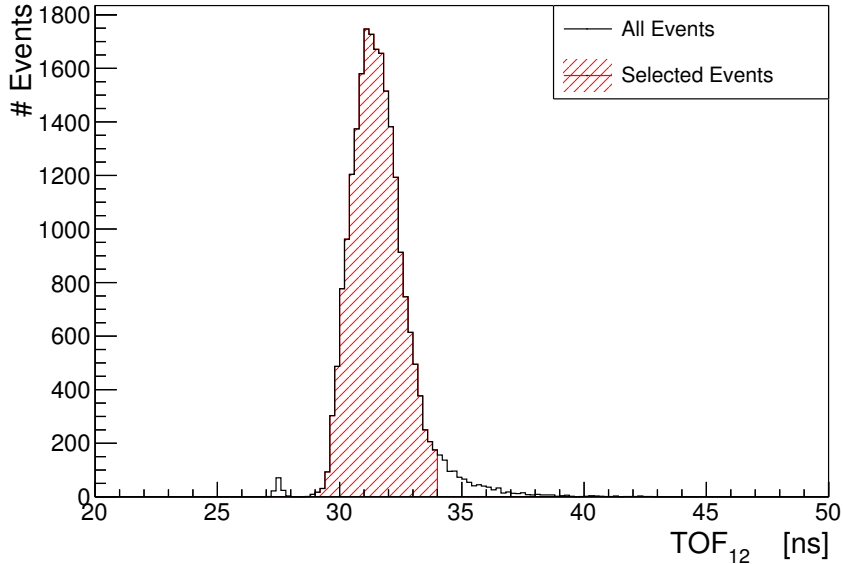


Figure 5.1: The time-of-flight distribution between TOF1 and TOF2. The events within the TOF cut are highlighted in red.

5.3 Low Level Reconstruction

5.3.1 Digit Production

The first stage of the data validation procedure was to ensure that the simplest objects in the reconstruction chain were being produced in the expected ratios, and with expected values. The data analysed were normalised by the number of events that passed the cuts described previously. This ensured that the results were not biased by the efficiency of a single detector, or by the track finding efficiency, rather they are representative of an analysable dataset with 100% transmission.

Previous studies on the material and structural composition of the trackers allowed for predictions for their performance to be calculated using only simple physics. The modal number of clusters for any track was expected to be 15. There are 5 stations, each containing 3 planes, where the average fibre has a better than 99% probability of producing a digit above background noise [52]. Additionally, clusters were expected to be predominately constructed from single digits due to the geometry of the planes, with approximately 1 two-digit cluster produced for every 11 single-digit clusters.

Distributions of the simplest reconstructed objects are shown below. For simplicity, plots from an arbitrary plane: the upstream tracker, station 2, plane




Figure 5.2: Distribution of raw ADC values (left) and the estimated NPE (right) for every digit across Plane 1, Station 2 in the upstream tracker.

1 have been included. Appendix B contains additional plots across all planes in both trackers. Data from all the planes looks extremely similar and any differences from expected distributions are highlight in the text. Figure 5.2 displays the distribution of digit raw ADC counts and NPE and figure 5.3 shows the number of digits recorded in each channel.

There are several features to note from figures 5.2 & 5.3 that are visible in multiple planes across both trackers. The combination of a very delicate construction, and large lengths of optical cabling was expected to introduce so called “dead channels” within each plane. Broken/damaged fibres and damaged connectors were expected in low numbers during the construction and can be clearly seen as channels where there are no data.

In addition, some VLPCs are known to have a poorly calibrated response to some channels, which resulted in large numbers of noise digits being produced. These channels may be masked in software at the digitisation stage, and will be indistinguishable from hardware based dead-channels. The combination of these two effects leads to some reduction in efficiency of the spacepoint reconstruction.

The steps in the distributions of figure 5.2 are due to boundaries between the electronics readout. Channels were readout in groups of 4 at a time, with every alternate 4 channels connected to a different ADC board which received a unique calibration. The boundaries between the different calibrations generate the stepped response. A program of calibration and testing was repeatedly applied in order to finalise the calibration techniques that were used in the reconstruction of the data under discussion.




Figure 5.3: Distribution of reconstructed digits across a single tracker plane in the upstream tracker, station 2, plane 1. The shape of the distribution is due to the natural muon beam profile after leaving the final quadrupole.

5.3.2 Cluster Production

The ratio of 2-digit clusters to 1-digit clusters provides a simple mechanism to verify the clustering algorithm. Table 5.2 summarises the ratio of single-digit to double-digit clusters for each plane in the upstream tracker.

It can be shown that the deviation from the ideal ratio is correlated with the amount of noise within each plane. If the noise were to be ideally removed from the analysis, it is anticipated that the ratios would be in agreement with the theoretical prediction. More detailed studies regarding the rate of noise within the trackers [53] and the effect on the reconstruction efficiencies are ongoing at the time of writing.

5.3.3 Spacepoint Production

Spacepoints are ideally constructed from 3 clusters, one from each of the three planes as this includes an extra degree of freedom when calculating the x and y transverse coordinates. The extra degree of freedom is effectively a powerful noise-rejection feature. It is likely that several noise digits will be produced in each event, however the probability of three noise digits coinciding within a station is vanishingly small. Hence, triplet-spacepoints are preferred for the track selection procedure, however this is not always possible due to dead channels or stochastically poor light yields.

The presence of dead-channels means that there are regions within several

Station	Plane	1-Digit	2-Digits	Ratio
1	0	14115	1033	13.66
	1	16451	1283	12.82
	2	16452	1146	14.36
2	0	16561	1308	12.66
	1	16493	1355	12.17
	2	16741	1195	14.01
3	0	16565	1366	12.13
	1	16916	1397	12.11
	2	16772	1213	13.83
4	0	16465	1344	12.25
	1	15779	1192	13.24
	2	16715	1200	13.93
5	0	16621	1436	11.57
	1	16771	1281	13.09
	2	16853	1133	14.87
Total		246270	18882	13.04

Table 5.2: Comparison of single-digit clusters to double digit clusters.

stations where it is physically impossible to reconstruct a triplet spacepoint. Additionally, imperfections in the fibres or deviations in the individual channels' response to scintillation light can reduce the triplet spacepoint efficiency. Figure 5.4 describes the total number of spacepoints in each station (red), including the number that were successfully used in tracks (green), and what fraction of those were triplets (blue).

The difference between the red and green curves corresponds to the amount of noise that is inherently present for the track finding algorithms. This is composed of overlapping noise digits, decay products and knock-on electrons. Theoretically the green curve, the space points that are included within tracks, should be flat in both trackers. This would demonstrate the efficiency of each station to produce a spacepoint is similar. This holds for the upstream tracker however some stations in the downstream tracker experience some reduced efficiency, most notably in station 5.

If the blue curve is also considered (the number of those spacepoints that are triplets) it can be seen that there is a significant reduction in the number of triplet spacepoints in the planes with lower efficiency. This is due to regions of the tracker where there is a high concentration of dead channels, removing the ability to actually find a cluster in each of the three planes. Additionally where dead channels from two planes overlap, the efficiency of finding a spacepoint




Figure 5.4: The number of spacepoints in each station. Red is the total number of spacepoints recorded, green are only those spacepoints that were assigned to a track during pattern recognition and the blue shows only those spacepoints that were assigned to a track and were composed of three clusters. Negative station IDs refer to the upstream tracker and positive station IDs refer to the downstream tracker.

drops to zero. Figure 5.5 shows the distribution of digits for each of the three planes in station 5 of the downstream tracker.

It can be seen that each of the planes has some number of a dead channels, especially plane 2, which corresponds to the significant drop in efficiency seen in figure 5.4. At present, possible options for improving the efficiency of downstream station 5 are currently under consideration. The option to swap out external waveguides and VLPCs for spares could be easily employed, however the exact cause of the dead channels is likely to be the internal wave guides or the fibres themselves, which would require a more involved operation. At




Figure 5.5: The distribution of recorded digits in planes 0 (left), 1 (centre) and 2 (right) of Station 5 in the downstream tracker. Note the large number of dead channels in plane 2 that detrimentally affect the reconstruction.

present there is not enough ancillary time permitted within the schedule to conduct such an operation. The algorithms however, are constructed such that they are able to correctly reconstruct non-ideal tracks, hence the efficiency will be slightly reduced but with little or no bias in the parameters of the final track fit.

Triplet spacepoints are quickly identified within the reconstruction by use of the Kuno conjecture (section 3.3.3), which requires the sum of the cluster positions to equate to a constant value determined by the number of channels in each plane. If however, all clusters across all three planes are summed together the distribution of Kuno-Sums highlights regions of high noise, in addition to verifying that the reconstruction is performing as expected.

Figure 5.6 demonstrates the distribution of Kuno-Sums across Station 1 in the upstream tracker. The expected peak at 318.5 is clearly visible (left), in addition to the effects of noise, while an enlarged scale (right) highlights alternate features of the same principle distribution.

The left hand plot sums every cluster that was present in the station per event. This allows the effects of noise to be clearly highlight. A band of noise can be seen from 320 upwards, corresponding to where there was a true triplet spacepoint, in addition to some noise clusters. Additionally there is a small excess around 220, which corresponds to a combination of noise and pairs of clusters that are missing a third in order to produce a triplet spacepoint.

A comparison with a Monte Carlo simulation is also provided, where only those spacepoints used in the track fit were included. This demonstrates the expected two-peak structure due to the integer values of 1-digit clusters and half-integer values of 2-digit clusters.

Such results could be used as an additional noise rejection procedure, where doublet spacepoints are rejected based on their Kuno-Sum for areas where there are several dead channels, thereby removing the known sources of doublet-spacepoint noise. At present, no such algorithm has been implemented, however several are under development.

A final feature to note is that the primary peak is not symmetric. Small discrepancies in the placement of the planes with respect to each other, have lead to an effective change in the kuno-number for some of the stations. This has been partially modelled in simulations, however a more precise measurement is required in order to precisely reproduce this effect.




Figure 5.6: An example Kuno distribution of the clusters found in Station 1 of the upstream tracker (left) and a comparison with Monte Carlo on the region of interest (right). The well defined peak at 318.5 corresponds to the expected Kuno-peak for the number of channels used in the construction of the station. The remaining entries correspond to a collection of noise, knock on electrons and poorly formed spacepoints. The two peak structure is an artefact of the integer binning used in determining the channel number.

5.3.4 Pattern Recognition

The final stage of track selection is pattern recognition. Spacepoints are selected combinatorially and tested to see which combination minimises the chi-squared of a track model (straight or helix) depending on the magnetic field (section 3.3.4). Spacepoints that are deduced not to belong to a track can then be removed from the subsequent track fitting and further analysis stages.

The superfluous spacepoints may be a result of several physical effects: coincidental noise; knock-on electrons, which are occasionally emitted from a plane due to the ionising effect of the primary track; decaying muons and pions; and secondary tracks due to multiple muons within the same event. However, their inclusion in the final track fitting stage would be detrimental to the analysis procedure.

The effect of the pattern recognition routine is to provide a final stage of noise rejection in the track reconstruction. This is illustrated in figure 5.7 which compares the NPE distributions for all spacepoints, and all spacepoints that are included in tracks. The peak due to single photo-electron noise has been almost completely removed with little effect on the remaining spacepoints. Knock on electrons and decays are expected at a low level, hence some spacepoints outside the 1 NPE peak were expected to be rejected also.

Detailed analyses are ongoing at the time of writing, to fully quantify the efficiency of the trackers using both MC and real data derived analyses [50]. This is the remaining component in a full validation of the track finding and fitting routines that are part of MAUS.




Figure 5.7: NPE per spacepoint. The red line includes all spacepoints that were reconstructed, the black line only includes spacepoints that were successfully used in a straight line, track fit.

5.4 High Level Reconstruction

The stages of pattern recognition were validated in detail using a Monte Carlo study (section 4.3). Additionally, the low-level reconstruction was previously validated with data. It can therefore be deduced that, assuming there are no structural differences between the trackers in simulation and as-built, the selection of spacepoints will proceed as expected. The following section will therefore concentrate on the reconstructed parameters following the Kalman Track Fit. More information regarding the performance of the specific pattern recognition stages can be found in [49, 51].

In order to improve the statistical weight of the following analyses both datasets A and B were analysed in parallel. Both datasets were subjected to the combined TOF and tracker coincidence requirements. Dataset A was subjected to a TOF cut of $28.5 \text{ ns} < \text{TOF}_{1 \rightarrow 2} < 34.0 \text{ ns}$, and Dataset B was subjected to a TOF cut of $28.5 \text{ ns} < \text{TOF}_{1 \rightarrow 2} < 32.0 \text{ ns}$.

5.4.1 Straight Tracks

The straight tracks used in this analysis were extracted from Dataset A (table 5.1). No additional cuts or requirements were placed on the tracks other than the coincidence and TOF cuts previously discussed.

In order to validate the reconstruction of the trackers, when no “true” information is present, only the more qualitative measures of the fit performance can be analysed.

In addition, a comparable Monte Carlo simulation was required. The full MICE geometry, including the up- and downstream PID detectors, and all materials, was used. To ensure that the trackers were modelled as precisely as possible, a digitisation model that mimics the NPE produced according to the current tracker calibration, in addition to modelling noise and dead channels, was included. This would theoretically produce resolutions similar to those in the real data reconstruction, hence permitting the global fit parameters to be compared between simulation and data.

Each reconstructed trackpoint has a “pull” value, defined as the distance between the fit and the measurement, when that particular data point is excluded from the fit (section 3.4.1). This is easily calculated using the Kalman filter structure by conceptually “reverse-filtering” that particular measurement, as described in [43].

The distribution of the pulls for each measurement plane can be used to demonstrate that the reconstructed track is in agreement with the raw data, and that the data is being appropriately weighted during the algorithm. Additionally, if the pulls are individually weighted by the reconstruction error, it can be seen whether the reconstruction is correctly calculating the errors associated with each trackpoint. Figure 5.8 shows the distribution of weighted pulls for both up and downstream reference planes.

It can be seen from figure 5.8 that there is good agreement between the Monte Carlo simulation and the data reconstruction. Particularly that they are both centred on zero with very similar widths. However the two-peak structure highlights a subtle effect of the geometrical layout of the channels. The nominal centre of an ideal hit - the geometric centre of three overlapping channels - is the centre of the resulting triangular overlapping region. However the centre of these triangles is not in the centre of the individual channel. Hence there are distinctly two identical distributions for each alternate channel. Figure 5.9 illustrates this effect.

This is the same asymmetry as seen in the distribution of the Kuno-sums (figure 5.6). A simple correction, was successfully applied to the upstream tracker, however the downstream tracker requires further improvements in order to correctly model this effect. At present this effect is still under study and not yet correctly modelled within the Monte Carlo geometries.

Similar plots for the other 28 planes can be found in appendix B. The RMS deviation from zero is approximately 0.37 for all distributions reconstructed at the reference plane. As the distributions are weighted by the calculated error, an RMS value of 1 is expected. These data demonstrate that the errors are




Figure 5.8: The distributions of weighted trackpoint pulls at the reference frame of the upstream (left) and downstream (right) trackers, for Dataset A (top) and Dataset B (bottom). The RMS deviation from zero is approximately 0.37 for all plots.




Figure 5.9: Planar representation of the overlapping fibres in the centre (green) of a station. The red markers indicate the centre of the overlapping channels and can be seen to alternate in their vertical position.




Figure 5.10: A comparison of the reconstructed p-values, for the straight track reconstruction, from data (black) and the Monte Carlo simulated p-values (red). Top: Dataset A, Bottom: Dataset B, Left: Upstream Tracker, Right: Downstream Tracker.

being over estimated by the fit. This is the expected behaviour given the results of section 4.2 as the multiple Coulomb scattering approximation was shown to overestimate the accuracy of the fit. If the multiple Coulomb scattering model were to be improved, the width of the distributions is expected to converge towards 1, indicating that the uncertainty due to multiple Coulomb scattering is more correctly modelled.

A more encompassing test of the implementation is available by considering the p-value distributions for each tracker. Figure 5.10 shows the comparison between the reconstructed straight track p-values and the appropriate reconstructed Monte Carlo p-values.

The plots demonstrate a strong similarity between the Monte Carlo model and the data reconstruction. As the p-value encompasses all aspects of the fit

performance, and directly compares the model of the system to the behaviour of the fit, these results strongly suggest that the reconstruction of data is behaving very similarly to that of the Monte Carlo models. Hence the reconstruction parameters can be assumed to be similarly accurate to Monte Carlo studies, barring any misalignments of either the trackers or the fields.

The small discrepancy between the reconstructed data and Monte Carlo in the downstream tracker, is a direct result of the plane misalignments that produced the asymmetry in the Kuno and Pull distributions, described previously.

5.4.2 Helical Tracks

A similar procedure was conducted for the helical track reconstruction. The availability of data however was more difficult. Only one run, Dataset C, was available at the time of writing. For this run, only the upstream spectrometer solenoid was powered, and only the End2-Centre-End1 coil pack was energised. Without the additional match-coils (M1 and M2) to further stabilise the fringe field, the uniformity of the field was not ideal. An Opera [54] field model was calculated and used within the Monte Carlo to better approximate the field conditions.

The events surrounding that period of data taking, and the specific complexities associated with Dataset C are detailed in chapter 7. For this analysis, it is assumed that the field map is sufficiently similar to the real field and that the Monte Carlo geometry correctly places all detectors, such that a comparison between Monte Carlo and data may be performed for the purpose of validation.

Despite these issues, the dataset does provide helical tracks that originated from the MICE nominal 200 MeV/c muon beam, reconstructed in an approximately 4T field. As only the upstream solenoid was powered, the transmission through to TOF2 was poor, hence it was not logical to apply a similar set of cuts as used previously, as many viable tracks would not be included. Rather a TOF0 to TOF1 time cut (figure 5.11) was enforced to ensure that the majority of the beam was composed of muons, and that precisely one helical track was reconstructed in the upstream tracker. Without improving the transmission, the downstream detectors could not be utilised without a significant reduction in statistics.

In order to ensure the effect of inefficiencies is reduced in the absence of ideal PID or transmission requirements, each track that was analysed was required to have precisely 15 trackpoints. This ensured that all the tracks were entirely composed of triplet-spacepoints and therefore have very low amounts of noise.




Figure 5.11: The time-of-flight distribution between TOF0 and TOF1. The events within the TOF cut are highlighted. The secondary peak was believed to be due to a combination of poor calibration and noise, and was removed as a precaution only.

Figure 5.12 shows the distribution of the reconstructed and simulated weighted pulls at the reference plane of the upstream tracker. The width of the distribution shows good agreement between Monte Carlo and data, while the structure of the distribution is not perfectly modelled. This is likely due to the plane alignment issues discussed within the straight track validation. As only a simple correction was used to model this effect, the relatively higher track gradients of the helical tracks could further distort the distribution. Further study is required to analyse the precise description of the misalignment. Additionally, discrepancies between the field model and the actual field may provide some additional bias to the pulls. However the overall agreement in mean and width of the distribution implies that little effect should be seen in the final reconstruction parameters.

Figure 5.13 shows a comparison of the simulated and reconstructed p-values in the upstream tracker. The low statistics of Dataset C increases the difficulty in making the comparison between reconstructed data and Monte Carlo. However the close agreement that is seen is encouraging. The lower values, 0.1-0.4 appear to be overrepresented by the Monte Carlo, however the shape of the distribution is in good agreement.

The disagreement in the pull distributions would cause the p-values to ex-




Figure 5.12: The distribution of weighted trackpoint pulls at the reference frame of the upstream tracker.




Figure 5.13: A comparison of the reconstructed p-values, for the helical track reconstruction, from data (black) and the Monte Carlo simulated p-values (red) in the upstream tracker.

perience some related effect, as was found. Additionally, the effect of using an approximate field map will also alter the p-value distribution. Furthermore, the known plane alignment discrepancy would likely also subtly detriment the p-value distribution. The good agreement despite these uncertainties associated with this dataset, demonstrates that the fit is working well.

5.5 Conclusions

Overall, the low-level reconstruction was found to behave as well as expected, given the complications of variable calibration potency, noise induced digits and decay/knock-on electrons. The predominant form of noise rejection works through the track selection in that only those spacepoints that are with some criteria for forming a track are accepted. This was shown to be very effective at removing the 1 NPE noise peak from the data set.

The effects of the poor station efficiency in station 5 of the downstream tracker have not been fully analysed. However, the first indication is that only a slightly reduced efficiency is expected. Distributions of pulls and p-values are still in agreement with Monte Carlo Studies. The option of replacing components of the tracker in order to improve the efficiency is under consideration but unlikely due to the amount of man-hours required to disassemble downstream spectrometer solenoid and perform the repairs. The Monte Carlo simulation however, can natively model dead channels and poor efficiencies, hence the effect of these dead channels will be included in all future simulations.

As the efficiency of pattern recognition was shown to be very high (section 4.3) in Monte Carlo studies, and the spacepoint forming algorithms are shown to be working at a high efficiency, the behaviour of pattern recognition on data can be deduced to be within expectations, assuming that the track model is correct. A validation of this assertion is currently under study [50] whereby the efficiency of the track finding is to be tested by estimating the number of tracks expected using up- and downstream PID detectors. There will still be a small component of noise within this study, however it is expected to agree with Monte Carlo to a high precision.

The high-level reconstruction has been shown to produce global measures that are in good agreement with the Monte Carlo models. The p-value comparisons demonstrate that all the combined effects, that the track fit is sensitive to, are well modelled in the Monte Carlo geometries and configurations. This does not guarantee the performance of the fit parameters, however it does demonstrate that the entire system model is accurate. Section 4.2 can be considered

to demonstrate the sensitivity of the p-value distribution to small changes in the model.

Validation of the parameters of the track fit was conducted as required in chapters 6 & 7. The straight track reconstruction has very few systematics that can affect the fit, other than the alignment of the trackers, which is directly analysed in chapter 6. The helical fit is more complex and very sensitive to both the magnitude and uniformity of the solenoid field. This is discussed in chapter 7, where the sensitivity of different field configurations is considered in addition to the effects of a misalignment.

Due to the relative simplicity of the straight line fit, the results shown strongly suggest that it is performing as expected by the Monte Carlo model. The helical fit has been shown to have a good agreement with the Monte Carlo model, however the absolute scale of the fit parameters will be discussed in subsequent chapters.

Chapter 6

Alignment of the MICE Trackers Using Data

6.1 Introduction

The installation of the trackers was a complex procedure that involved careful manipulation of the trackers themselves within a low light environment. The available space within the Spectrometer Solenoids did not allow for the trackers to be rigidly fixed in place, rather they were installed and braced with the use of a rubber “foot”. Hence, the final location of the trackers was difficult to precisely survey, and although unlikely, they were susceptible to movement during transit. As such, a measure of the location of the two trackers, in addition to the usual survey measurements, was required to allow the detectors to be modelled to a high precision.

Additionally, performing an alignment procedure between the two trackers provides a powerful consistency check on the reconstruction parameters of the track fit. An unbiased and accurate reconstruction of both the positions and gradients is required in order to successfully align the two trackers.

No magnetic fields were present within the cooling channel hence the straight track reconstruction was used. This required elements of the track selection and fitting routines to perform as expected, in addition to the use of global track matching and calibration handling. As such, this provided the first end-to-end analysis task for the track reconstruction.

The reference planes of the trackers, both upstream and downstream, mark the point of the detectors at which the alignment was compared. A convention of aligning all components to the centre of the upstream tracker reference plane




Figure 6.1: Schematic representation of a reconstructed straight track and how a misalignment affects the global placement of the reconstruction. Note that the misalignment has been greatly exaggerated and is not to scale.

was chosen for simplicity. The alignment of the downstream tracker describes the translation and rotation of the downstream tracker reference plane from the expected position, as assumed in the reconstruction geometry. Initially the trackers were expected to be perfectly aligned in the transverse dimensions and have a known longitudinal separation, L .

The data inherently describes the transverse components of the alignment, however it would also be possible to estimate the longitudinal displacement from the expected tracker separation in future analyses.

Figure 6.1 shows a cartoon layout of the trackers, as described by the reconstruction geometry and as they would be in reality, with an example track reconstructed (note that these drawings are not to scale, and are greatly exaggerated to provide a visual representation only). They demonstrate the differences in the reconstruction of a single track using both the upstream and downstream trackers. By comparing the downstream position and gradient to the expected position and gradient, the full transverse alignment can be calculated.

6.2 Alignment Method

6.2.1 Outline

The upstream tracker was used to define the origin of the coordinate system, hence the upstream tracks were defined to be correctly reconstructed, therefore

only one set of corrections, for the downstream tracker, needed to be calculated. Each upstream reconstructed track provides position (X, Y, Z) and gradient (M_x, M_y) information at the reference plane.

The upstream tracks can then be individually propagated to the downstream tracker using a linear extrapolation to determine the expected positions (x, y) and gradients (m_x, m_y) . The downstream longitudinal position, z was used as expected from the geometry, hence only the parameter $L = z - Z$ is used. This propagation was performed using the equations:

$$\begin{aligned}
 x &= X + M_x L, \\
 y &= Y + M_y L, \\
 m_x &= M_x, \\
 m_y &= M_y.
 \end{aligned}
 \tag{6.1}$$

In practice it was more sensible to use the arc-tangent of the gradients - the track angle - hence, this is the nomenclature adopted in the following discussion. The x and y track angles (η_x and η_y) are calculated from the gradient as, $\eta = \arctan(m)$.

The downstream tracker also provides measured estimates to the track parameters, these are distinguished with a superscript m : $(x^m, y^m, \eta_x^m, \eta_y^m)$.

The comparison between the estimated track parameters and the measured track parameters permitted the downstream tracker to be aligned with respect to the upstream tracker. If there were no stochastic effects, a single track could be used to precisely determine the alignment of the trackers through the residuals between the reconstructed track parameters and the predicted track parameters:

$$\begin{aligned}
 \delta_x &\equiv x^m - x, \\
 \delta_y &\equiv y^m - y, \\
 \phi_x &\equiv \eta_x^m - \eta_x, \\
 \phi_y &\equiv \eta_y^m - \eta_y.
 \end{aligned}
 \tag{6.2}$$

Multiple Coulomb scattering causes the δ and ϕ distributions to be heavily smeared, depending on the particle species and velocity, hence the parameters δ and ϕ are only estimates for the actual tracker misalignments, Δ and Φ . Figure 6.2 provides a graphical representation for how the parameters relate in a single transverse dimension.

Repeated analysis of straight tracks would produce some distributions in each of $\delta_x, \delta_y, \phi_x, \phi_y$, which can then be used to estimate the relative alignment

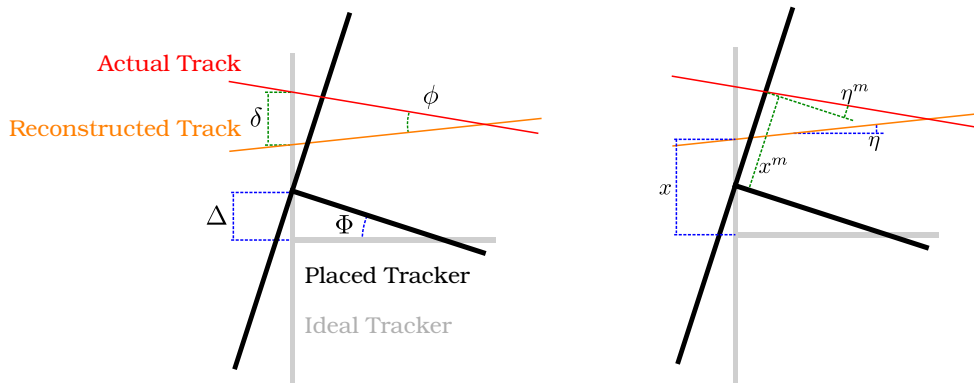


Figure 6.2: Geometrical layout of the alignment variables, and their relation to each other. This diagram describes some arbitrary axis, the x and y parametrisations are identical.

between the two trackers.

6.2.2 Assumptions

The Paraxial Approximation

From close inspection of figure 6.2 it can be seen that the equations 6.2 are only correct to first order. In fact, the value of δ experiences a parallax-style correction due to the rotation, Φ , which has been ignored under the paraxial approximation. For this reason, the alignment procedure must be iteratively applied to the geometry and re-processed, such that the residuals may be minimised in order to make the second-order corrections negligible.

No Axial Rotation

It was assumed that the trackers do not rotate about their central axis, or if they do, the angle between the two is known to a high precision and is correctly modelled in the reconstruction geometry. This assumption permits the alignment procedure to be treated as two separate problems, one in the x - z plane and one in the y - z plane¹. For this reason, x and y subscripts will be used to denote the separate systems, however the equations are identical for both systems.

6.2.3 Procedure

The runs that were used in this analysis correspond to Datasets A and B in table 5.1. Both datasets were used as the only difference between them is the

¹Due to how the 3D rotations are modelled, the two reconstruction planes are coupled by the order in which the rotations are applied, they may be treated as uncoupled under the paraxial approximation only.




Figure 6.3: Comparison of the two Time of Flight windows used in this analysis. Dataset A, corresponding to the 200 MeV/c muon beam are in the left plot and Dataset B corresponding to the 240 MeV/c muon beam are in the right plot.

mean muon momentum. Hence the only comparable effect between the two datasets will be the effect of multiple scattering. As this does not statistically bias the fit, only affecting the distribution widths, both datasets can be used without issue.

The deviation due to gravity will affect the two datasets differently as they will have different mean velocities, however the magnitude of the deviation is estimated to be on the order of 10^{-15} m across the entire length of the cooling channel, many orders of magnitude lower than the resolution of all the detectors, hence it was neglected within this analysis.

For each run, reconstructed straight tracks were processed individually at the reference plane of each tracker. Only events that produced precisely one track in each of the upstream and downstream trackers, in coincidence with passing the TOF cut, were considered. For Dataset A, the TOF cut set to (28.5, 34.0), as was used for the validation procedure in chapter 5. For Dataset B, a similar cut was imposed, (28.5, 32.0). Figure 6.3 shows the comparison between the two TOF distributions for the two datasets.

To ensure that the dataset contained only high-quality reconstructed tracks, thereby reducing the effects of noise and hard scatters, a cut on the reconstructed p-value was imposed. Only tracks with a reconstruction p-value > 0.01 were included within the analysis. For an ideally flat distribution, this would ensure that the 1% worst tracks would be removed. However due to the effects described in section 4.2, fewer than 1% of tracks are actually removed. However those that are removed, still correspond to the poorest quality fits.

An additional requirement on the radius of the projected track was also im-

posed. Each tracker only has a fiducial radius of 15 cm. Any spacepoints that could be reconstructed outside of that radius are excluded from the analysis. This ensures that the curvature of the fibres near the edge of a plane does not affect the measurements. However, the fiducial radius does affect the transmission into the downstream tracker. It manifests as a convolution of the 15 cm radial window with the approximately gaussian beam distribution, hence the actual distribution will experience some bias as a function of the tracker misalignment. By ensuring that only tracks that, once projected, fall within a selection radius are included in the analysis, the effect of the fiducial radius can be negated.

The typical amount of scattering between the two trackers corresponds approximately to a 20 mm RMS displacement. Therefore if the projection cut was chosen such that it was larger than the nominal scattering angle, but smaller than the fiducial radius, the central region of the distribution would be representative of the true distribution. The projection cut was placed at 100 mm, 2.5 standard deviations from the limit fiducial cut (containing approximately 90% tracks that pass the cut). This maintained a reasonable transmission, whilst a gaussian fit was then applied to the central peak in order to calculate the mean and standard deviation.

Each track in the upstream tracker was then propagated to the downstream tracker and compared to the downstream reconstructed track parameters. The residuals in each of the four reconstruction parameters, $(x, y, \eta_x$ and $\eta_y)$ were then histogrammed and fitted. Figure 6.4 shows these distributions as generated using the default MICE geometry.

A Monte Carlo study that made use of a precise geometry and a similar beam was performed in parallel in order to estimate the sensitivity of the method and verify that it was accurate.

The downstream tracker alignment can then be estimated from the distributions using a gaussian fit in each of the four dimensions, where,

$$\begin{aligned}\Phi_x &= \langle \phi_x \rangle, & \Delta_x &= \langle \delta_x \rangle, \\ \Phi_y &= \langle \phi_y \rangle, & \Delta_y &= \langle \delta_y \rangle.\end{aligned}\tag{6.3}$$

The total error on the fitted gaussian provides an estimate for the error in the alignment; however, the tracker resolution adds an additional smearing effect to the reconstruction. The result is that the distributions obtained were in fact a convolution of the multiple Coulomb scattering distribution with the intrinsic tracker resolution, hence the error on the alignment was over estimated. This can be approximately corrected to determine the true reconstruction precision.




Figure 6.4: Position (left) and angular (right) residuals calculated by propagating the track parameters from the upstream tracker to the downstream tracker. The position residuals have x and y standard deviations of approximately 20 mm. The angular residuals have x and y standard deviations of approximately 8.3 mrad.

Section 4.3, details the resolution of straight track reconstruction within a detailed Monte Carlo study. If the tracker resolution and the multiple Coulomb scattering distributions are assumed to be gaussian distributed, the convolution between the two distributions corresponds to the addition in quadrature of the standard deviations. Using this approximation, the resolution-independent standard deviation can be estimated.

The final results obtained from this procedure were then applied to the reconstruction geometry, and the process was repeated. Subsequent iterations provided smaller corrections to the initial estimate, and eventually converged within the expected errors. In practice, approximately 3 iterations were necessary to converge sufficiently.

6.3 Monte Carlo Validation

In order to validate the procedure outlined in section 6.2, two Monte Carlo simulations were conducted. The baseline, MICE Step IV geometry was altered such that both trackers were perfectly placed in the transverse plane and perfectly aligned with each other for the first study. Both trackers were then misaligned by a known measure for the second study. A simple gaussian muon beam with a small spread in momentum as modelled, starting immediately before the upstream tracker and allowed to propagate through the downstream tracker. No other detectors were required for the study.

Figure 6.5 shows both the position and angular residuals for the aligned geometry. The alignment residuals and their associated errors were calculated




Figure 6.5: Position (left) and track angle (right) residuals for the ideal Monte Carlo Simulation. The x and y standard deviations of the position distribution is approximately 19 mm and the x and y standard deviations of the angular distribution is approximately 8 mrad.

using a gaussian fit, applied to the histogrammed data. As indicated, alignment residuals were found to be:

$$\begin{aligned}\Delta_x &= -0.208 \pm 0.215 \text{ mm}, \\ \Delta_y &= -0.093 \pm 0.208 \text{ mm}, \\ \Phi_x &= -0.010 \pm 0.081 \text{ mrad}, \\ \Phi_y &= 0.086 \pm 0.081 \text{ mrad}.\end{aligned}$$

All components are comfortably in agreement with no misalignment. No cuts on time of flight or p-value were imposed for this Monte Carlo study. The effect of hard scatters in the trackers or other materials were therefore not correctly accounted for, but did not result in an significant biases or degradation of the results.

The misalignment introduced into the second Monte Carlo study corresponded to 5 mm translation in both x and y for the upstream tracker and a -10 mrad rotation in the y - z plane for the downstream tracker. Figure 6.6 shows the position and angular resolutions from the study. The measured alignment residuals were found to be:

$$\begin{aligned}\Delta_x &= -5.097 \pm 0.207 \text{ mm}, \\ \Delta_y &= -5.395 \pm 0.216 \text{ mm}, \\ \Phi_x &= -0.118 \pm 0.084 \text{ mrad}, \\ \Phi_y &= 10.063 \pm 0.081 \text{ mrad}.\end{aligned}$$

All components were again found to be in agreement with the known misalign-




Figure 6.6: Position (left) and track angle (right) residuals for the ideal Monte Carlo Simulation. The standard deviation of the position distribution is approximately 19 mm and the standard deviation of the angular distribution is approximately 8 mrad.

ment in the geometry. Note that the plots demonstrate the residual between the projected and reconstructed trackpoints, hence there is a sign inversion in order to calculate the actual geometric misalignment. This study validates that the measured parameters were correctly calculated.

6.4 TOF Extrapolation

A final verification of the straight track reconstruction, and a secondary alignment measurement, was the extrapolation to the TOF1 and TOF2 detectors from the upstream and downstream trackers respectively. Both offer a coarse position estimate for the track and may be used to estimate the alignment of the tracker-tracker system to the other PID detectors.

For each event that passed the cuts outlined above, the upstream and downstream tracks were treated similarly, in that they were sequentially propagated to the nearest TOF detector, and the residual between the predicted and reconstructed x and y components was histogrammed. The results of this are shown in figure 6.7.

The relatively large pixels in the TOF detectors caused the distinctly square residual distribution, in addition to having a smearing effect. Note also that the positions of the TOF detectors with respect to the trackers is not perfect, rather there is some misalignment. Due to the construction and installation of the detectors, this was expected.

The trackers were assumed to be correctly aligned with respect to the rotation about the central axis; however, no assumptions were made with respect to the TOF detectors. Figure 6.8 shows the correlation between the predicted




Figure 6.7: Upstream (left) and downstream (right) residuals calculated by propagating the track parameters to the nearest TOF detector. The upstream tracker was propagated to TOF1 and the downstream tracker was propagated to TOF2.




Figure 6.8: The correlations between the predicted and reconstructed transverse coordinates at TOF1, as compared to the upstream tracker. x is shown in the left histogram and y is shown in the right histogram. The black points mark describe the profile of the residuals in each TOF measurement bin. The gradient of the correlation was calculated from the linear fit drawn in red.

and reconstructed values of x and y for the upstream tracker-TOF1 system. Figure 6.9 shows the similar distributions for the downstream tracker-TOF2 system. These figures demonstrate that any misalignment about the central axis is minimal, as all components are correlated with a gradient of 1.0 within errors.

6.5 Axial Rotation

The trackers were assumed to be aligned in terms of the axial rotation (rotation around the beam axis). This assumption was believed to be true due to the installation tolerances, however a validation of this assertion was required.




Figure 6.9: The correlations between the predicted and reconstructed transverse coordinates at TOF2, as compared to the downstream tracker. x is shown in the left histogram and y is shown in the right histogram. The black points mark describe the profile of the residuals in each TOF measurement bin. The gradient of the correlation was calculated from the linear fit drawn in red.

When projecting a track from the upstream tracker to the downstream tracker, the η_x and η_y track angles may be used to determine the track description in cylindrical polar coordinates, $\eta = \sqrt{\eta_x^2 + \eta_y^2}$ (the polar angle with respect to the z axis) and $\theta = \arctan(\eta_y/\eta_x)$ (the azimuthal angle around z in the x - y plane).

Any rotation of the downstream tracker with respect to the upstream tracker would have been reconstructed as a residual, ω , between the predicted value, θ , and the measured value, θ^m , Figure 6.10 shows this distribution.




Figure 6.10: Residual between the projected track azimuthal angle and the measured azimuthal angle. A gaussian fit (shown in red) was applied to the peak in order to determine an approximate parametrisation for the distribution.

Run	Muons	Δ_x/mm	Δ_y/mm	Φ_x/mrad	Φ_y/mrad
07672	515	-12.25 ± 0.94	14.11 ± 1.00	1.93 ± 0.42	1.18 ± 0.40
07673	4614	-10.23 ± 0.34	14.11 ± 0.34	1.69 ± 0.14	2.51 ± 0.14
07674	3821	-9.92 ± 0.32	14.86 ± 0.32	1.83 ± 0.13	2.71 ± 0.13
07681	4298	-10.16 ± 0.34	14.29 ± 0.36	1.41 ± 0.14	2.48 ± 0.14
07682	3978	-10.33 ± 0.31	14.46 ± 0.32	1.51 ± 0.13	2.53 ± 0.13
07685	3949	-10.43 ± 0.32	15.44 ± 0.32	1.70 ± 0.13	2.81 ± 0.13
07686	2889	-9.95 ± 0.37	14.83 ± 0.38	1.61 ± 0.15	2.93 ± 0.15
07691	3906	-9.92 ± 0.32	14.77 ± 0.32	1.49 ± 0.13	2.49 ± 0.13
07693	638	-10.53 ± 0.82	15.49 ± 0.86	1.73 ± 0.32	2.87 ± 0.31
07694	2768	-9.56 ± 0.37	14.82 ± 0.38	1.85 ± 0.16	2.59 ± 0.15
07695	3328	-9.32 ± 0.41	14.73 ± 0.42	1.86 ± 0.17	2.83 ± 0.17
Combined	34740	-10.06 ± 0.11	14.71 ± 0.11	1.65 ± 0.05	2.62 ± 0.05

Table 6.1: A summary of the transverse alignment measurements for each run in the two Datasets that were analysed. The combined alignment is calculated from the weighted mean of all runs.

The results show that the trackers were aligned to a high accuracy, however a statistically significant misalignment does exist, reconstructed to be $\Omega = (-0.027 \pm 0.004)$ rad. As a result, there is a fractional, systematic error within the transverse alignment on the order of $\sin(0.027) \approx 0.027$. As the fractional errors on the reconstruction is still an order of magnitude larger than this correction, the no-axial rotation assumption is valid to first order. However with higher statistics, thereby permitting a more precise measurement, a full 3-dimensional fit would be required.

6.6 Results

The complete set of transverse alignment measurements from both Datasets A and B is presented in table 6.1. Each run was processed individually such that any deviations between consecutive measurements could be noted in the event of some systematic change. The overall alignment was calculated as the weighted mean of each individual measurement.

Section 6.5 described the process of analysing the relative axial rotation between the two trackers and section 6.4 described the comparisons between the trackers and the TOF detectors. With all measurements in hand, the combined TOF1-tracker-tracker-TOF2 system can be globally positioned with respect to

Detector	Δ_x/mm	Δ_y/mm	Φ_x/mrad	Φ_y/mrad
TOF 1	-1.57 ± 0.07	-4.77 ± 0.07	-	-
Upstream Tracker	0.0	0.0	0.0	0.0
Downstream Tracker	-10.06 ± 0.11	14.71 ± 0.11	1.65 ± 0.05	2.62 ± 0.05
TOF 2	6.96 ± 0.10	6.18 ± 0.10	-	-

Table 6.2: A summary of the relative alignment measurements for each of the trackers and TOF detectors. The upstream tracker defines the origin of the coordinate system, and hence has no relative alignment.

each other. Table 6.2 describes the relative positions and rotations between the detectors. The effect of the axial rotation was applied as a first order correction to the downstream tracker and TOF parameters.

The errors for each parameter were calculated using the residual distributions of the final iteration of the alignment procedure. This was to correct any additional smearing of the error due to the misalignments. Each parameter was assumed to be gaussian distributed and the error on the mean was calculated as appropriate given the number of entries.

6.7 Conclusions

The straight track alignment provided a powerful test of the straight track reconstruction. The alignment procedure required the reconstruction parameters to be accurate and precise in order to make a sufficiently precise measurement, such that the reconstruction geometry may be modified for future analyses. A global validation in the form of the TOF extrapolation demonstrated that the reconstruction is not experiencing any systematic errors within the tracker-tracker system, and the good agreement between all the detectors has shown that the track reconstruction is performing within expectations.

Additionally, figures 6.5 & 6.6, the Monte Carlo transverse residuals, show an excellent agreement with the results in figure 6.4, the reconstruction residuals from data. The close agreement between the width of these two distributions demonstrates that the Monte Carlo model is closely modelling the multiple Coulomb scattering due the materials within the MICE cooling channel. The small discrepancy that does exist may be explained as being due to GEANT4's model of MCS, in particular with the treatment of tail-events outside the region that is approximately Gaussian. The ongoing and detailed multiple Coulomb scattering study will provide a more thorough investigation into these results.

The transverse alignment was performed without complications, and demon-

strated that the trackers are rotationally aligned to a high accuracy, however the 10-15 mm misalignments are of concern. The construction tolerances were typically on the order of 1-5 mm, hence this demonstrates an unexpectedly poor placement of the downstream tracker. However, laser-survey measurements of the downstream solenoid are in agreement with this measurement, indicating that the cause of the misalignment was the initial placement of the downstream spectrometer solenoid. Currently a mechanical correction is being planned to improve the relative placement of the Spectrometer Solenoids.

Additional studies have/are being performed in order to provide further alignment measurements. A global detector alignment, has been performed using all downstream PID detectors in addition to an attempted extrapolation to TOF0 that required the quadrupole magnets Q4-Q9 to be unpowered. Due to the relatively large distance and very low rate without the quadrupoles, the extrapolation to TOF0 was unsuccessful. However a comparable set of results that are in agreement with those presented above were obtained [55]. Additionally, as part of the ongoing multiple Coulomb scattering study a more precise, fit-based approach using the existing library Millipede [56] is under development. This will involve a full 3D fit to spacepoint level-data at all tracker stations. The results are expected to be sensitive to the axial rotations and provide a more precise measure of the relative tracker rotations as a result.

The extrapolation to the TOF detectors demonstrated that the reconstruction within the tracker software is globally consistent with other detectors. The correlation measurements demonstrated that the axial rotation between the trackers and the TOFs is consistent with zero, as expected, and the transverse residuals highlighted some expected displacements. The laser-survey measurements are in agreement with the reconstructed residuals, which further demonstrates that the straight track reconstruction is performing as expected.

The relative axial rotation between the two trackers, Ω , was estimated using the residuals in the reconstructed azimuthal track angles, ω . The resulting distribution demonstrated a small, non-zero relative rotation of -0.027 rad. As this is a relatively small rotation, the reconstructed transverse alignment values were only approximately correct within the current errors. However with higher statistics, providing a more precise measurement, the axial rotation will become a prominent systematic error. A first order correction, based on Ω , was used in table 6.2 which represents the most accurate values obtained within this study.

Chapter 7

Measurement of the MICE Muon Beam Emittance

During the commissioning of the MICE Cooling Channel, many issues were experienced relating to the design and construction of Spectrometer Solenoids. During the 2015-2016 run period, a flaw in the design of the quench protection system meant that it was no longer practical, or safe, to run either of the two Spectrometer Solenoids until they have been retrofitted with an upgraded system.

During a non-standard commissioning exercise, an event occurred that highlighted these issues and resulted in Match Coil 1 of the downstream Spectrometer Solenoid being irreparably damaged [57]. Unfortunately only one run had so far been recorded that contained muon data with the 4T magnetic field. This chapter describes the use of these data in estimating the intrinsic beam emittance that was accepted into the MICE Cooling Channel.

7.1 Experiment Configuration

The MICE muon beamline was configured in the standard “muon beam” setting, as was used in previous analyses [33]. This provided a very pure muon beam with which this analysis was conducted. All PID detectors were present and functioning, however the trackers were being commissioned at this time and were still being calibrated. During the data taking period only the upstream tracker was functioning correctly and could provide useful data. As a result, only one run was successful and the number of muon tracks, within a solenoidal field, that could be analysed was limited.

The state of the magnet commissioning programme permitted only the End 2-Centre-End 1 coils of the upstream solenoid to be powered. This had two significant effects on the analysis:

1. The transmission into the upstream tracker was high due to presence of the focussing solenoid field. However, as the field ended after the End 1 coil, the relative transmission into the downstream PID detectors was very low. If the downstream PID was to be used for analysis, the number of muons tracks recorded in this short run would be significantly reduced, hence only upstream PID and the upstream tracker could be used without the adversely affecting the statistics of the study.
2. The default operating currents for the Spectrometer Solenoids assumed a magnetic field in both Match Coil 1 and Match Coil 2, in addition to the ECE pack. As these coils were not powered, the actual field within the tracking volume was not as uniform as expected by the design parameters. As the track fit algorithms all make extensive use of the uniform field assumption, the performance of the reconstruction during the run was subject to some non-linear systematic effect. This required detailed analysis in order to precisely model the behaviour of the reconstruction in such an environment.

An official Monte Carlo simulation was also performed in order to model the full experimental configuration, from the MICE target through to the downstream PID detectors. This allowed the analysis techniques and algorithms to be applied to a situation where the true values are known, such that they can be tested before being applied to real data. Additional simulations also permitted the systematic effects to be precisely modelled and analysed.

7.2 Procedure

In order to accurately determine the emittance of the MICE muon beam, it was imperative that all sources of systematic error be correctly modelled and accounted for. Monte Carlo studies were extensively used in order to quantify the predominant effects: field misalignment, field uniformity, and reconstruction bias.

The track fit assumes a uniform solenoidal field throughout the tracking region, perfectly aligned to the trackers. If there existed a significant misalignment, the reconstruction model would be unable to perform any reliable

reconstruction. Hence initial studies were conducted in order to establish that any effect was small.

The effect of the field uniformity was then examined. Although the reconstruction model does not anticipate a non-uniform field, it would appear as a less significant effect in the reconstruction. Even with a moderately non-uniform field, the track path would remain approximately helical. The effects would therefore be seen in the magnitude of the reconstructed momentum. As momentum corrections have already been discussed and applied successfully (section 4.3.2), this was an anticipated systematic correction for the analysis procedure.

The final systematic source of error is the reconstruction bias. The emittance calculation relies on the correlations between all phase-space variables. Hence it is possible for the reconstruction to produce statistically accurately x , y , p_x , and p_y estimates that don't reproduce an accurate emittance estimate due to: (a) off-diagonal correlation terms in the covariance matrix; and (b) a smearing effect due to the detector resolution. These effects are accounted for through the use of covariance matrix corrections. By analysing the differences between the simulated track information and reconstructed track information it is possible to estimate the parameters in this correction term.

The covariance matrix corrections could theoretically account for any field concerns, reconstruction biases, or momentum residuals. The applied reconstruction procedure was to apply the momentum corrections to the individually reconstructed tracks, before applying the calculated covariance matrix corrections. The alternative would be to neglect the momentum corrections, and rely on a more detailed covariance matrix correction that correctly accounted for the momentum residuals.

On the assumption that the Monte Carlo model is a perfect representation of the experiment, this would be a feasible approach; however, there are two concerns. Doing so requires a Monte Carlo simulation that is a very precise model of the experiment, including all field descriptions and beam parametrisations. It was not possible to generate such a model to sufficient precision and accuracy. Additionally, this would also mean that the track reconstruction is prepared without the known momentum corrections being applied, hence any analysis that was not based around the reconstruction of the covariance matrix would need to apply a momentum correction regardless.

The primary source of error was found to be the momentum residual, hence the decision was therefore taken to calculate and apply the total momentum corrections to the reconstructed tracks, and use covariance matrix corrections

to finally correct the covariance matrix before analysis.

The correction parameters were determined through a set of simulations using MAUS. A realistic model of the tracker was placed within a configurable field map. The uniformity and alignment of the field map could be varied independently. Additionally the beam being simulated was initiated at the entrance to the upstream tracker, immediately before station 5, and could be configured to model any simple distribution.

The alignment of the field was estimated using a combination of survey data and tracker alignment data, and any variations between the model and real field will correspond to a source of error. However, the effect on the emittance reconstruction can be shown to be small (see below) at the measured value (≈ 3.3 mrad). Additionally, the statistical properties of the data reconstruction have been shown to be in good agreement with the Monte Carlo model (see chapter 5), hence the tracker geometry and reconstruction biases should be accurately reproduced by the simulation. Due to time constraints there still exists some difficulty in the modelling of the canonical angular momentum, hence beams with a large transverse coupling, would not be well reproduced by the systematic studies. This effect is small however. The primary source of error in estimating the correction parameters is therefore the knowledge of the field map within the tracking region. Studies are ongoing [58] to precisely determine the agreement, however a conservative estimate is presented below.

Following the determination of the correction parameters, the analysis procedure was performed on an official Batch Monte Carlo dataset in parallel with Dataset C (table 5.1). The Batch Monte Carlo is maintained by the MICE collaboration and is undergoing constant improvements to better model the behaviour of the MICE experiment at all positions within both the beamline and the cooling channel. As such it will introduce additional effects and correlations to the transverse phase-space that are not modelled by simple gaussian approximations, but are more similar to the real experiment.

For both the Batch Monte Carlo dataset and Dataset C, an initial analysis of the uncorrected reconstruction was performed in order to determine the distribution of tracks in phase-space, and select the most appropriate correction parameters. Each track was then subjected to a momentum correction, before being added to the covariance matrix calculation. Following the determination of the covariance matrix, the emittance and other optical parameters could be calculated.

The Batch Monte Carlo analysis permitted a comparison between recon-

structed optical functions and the true beam information, such that the reconstruction and error analysis may be validated within the study. The reconstruction of Dataset C was then assumed to be accurate with the estimated errors, with an additional small uncertainty due to the field map being used within the Monte Carlo model.

7.3 Error Analysis

In order to precisely determine the true beam emittance, both the statistical and systematic errors must be determined. The statistical error is due to the sampling of the beam, as it is only possible to record a finite number of tracks, thus sampling from the true distribution

Conversely, systematic effects, derived from the tracker reconstruction and field configuration, result in the beam being incorrectly reconstructed and form a significant source of error in the final emittance measurements. Both the statistical and systematic sources of error are discussed below, with particular attention to the treatment of the systematic errors.

7.3.1 Statistical Errors

The transverse emittance of a beam is calculated by determining the 4×4 covariance matrix, Σ , and calculating the determinant such that the normalised emittance is given by,

$$\epsilon_{\perp} = \frac{|\Sigma|^{\frac{1}{4}}}{\beta\gamma}. \quad (7.1)$$

The determinant calculation can be expanded to reveal a sum of products composed of the variances and covariances of the distribution, hence if the error on a variance estimate can be determined, the error on the emittance can also be determined. It is assumed for simplicity that the error on a covariance ($\text{Cov}(x, y) = \sigma_x\sigma_y\rho$) is entirely determined by the errors of the associated standard deviations (σ_x and σ_y) and that any error associated with the correlation coefficient (ρ) can be neglected.

It can be shown (Cochran's Theorem [59]), that for n samples of a gaussian distribution with variance, σ^2 , that the sample variance, s^2 , is related to the true variance by a chi-squared distribution,

$$\frac{(n-1)\sigma^2}{s^2} \sim \chi_{n-1}^2. \quad (7.2)$$

It can therefore be shown that the variance of the variance-estimate is ap-

proximately given by,

$$\begin{aligned}\text{Var}(\sigma^2) &= \text{Var}\left(\frac{s^2}{n-1}\chi_{n-1}^2\right), \\ &= \frac{2s^4}{n-1}.\end{aligned}\quad (7.3)$$

Hence, the error on the estimate of a sample variance, s^2 , is given by $\sqrt{2s^4/(n-1)}$. It can additionally be shown that, using this estimate, the approximate error on a sample standard deviation, s , is given by $s/\sqrt{2(n-1)}$.

For a 2×2 covariance matrix Σ_{2D} , the determinant may be written in terms of the sample variances, s_x^2 , and sample covariances, $\text{Cov}(x, y) = s_x s_y \rho$, where ρ represents the correlation coefficient,

$$|\Sigma_{2D}| = \begin{vmatrix} s_x^2 & s_x s_y \rho \\ s_x s_y \rho & s_y^2 \end{vmatrix} = s_x^2 s_y^2 - s_x^2 s_y^2 \rho^2. \quad (7.4)$$

The error on the determinant can be expressed in terms of the errors on the standard deviation estimates,

$$\sigma_{|\Sigma_{2D}|}^2 = 4(1-\rho^2)^2 s_x^2 s_y^2 (s_y^2 \sigma_{s_x}^2 + s_x^2 \sigma_{s_y}^2). \quad (7.5)$$

Using the results from above,

$$\sigma_{|\Sigma_{2D}|}^2 = \frac{4(1-\rho^2)^2 s_x^4 s_y^4}{n-1} = \frac{4}{n-1} |\Sigma_{2D}|^2. \quad (7.6)$$

It therefore follows that the error on a 2×2 determinant is given by,

$$\sigma_{|\Sigma_{2D}|} = \frac{2}{\sqrt{n-1}} |\Sigma_{2D}|. \quad (7.7)$$

Extending this calculation to higher dimensions results in the formula for an $N \times N$ covariance matrix,

$$\sigma_{|\Sigma_{ND}|} = \sqrt{\frac{2N}{n-1}} |\Sigma_{ND}|. \quad (7.8)$$

It therefore follows that the error on an 4-dimensional normalised emittance measurement is given by,

$$\sigma_{\epsilon_{\perp}} = \frac{\epsilon_{\perp}}{\sqrt{2(n-1)}}. \quad (7.9)$$

7.3.2 Covariance Matrix Corrections

The emittance is a statistical quantity based on the width of the phase-space distribution and all correlations within. It is therefore sensitive to any statistical effects within the system. Although the detector performance is consistent with zero systematic fluctuations, both the resolution of the track parameters and any correlations between the coordinates of individual measurements have an effect on the reconstructed emittance. This can be seen by considering the effect of measuring many tracks of some true phase-space distribution, with some error applied to the track parameters.

Consider a vector measurement of some arbitrary phase-space, where the true phase-space vector is denoted by \mathbf{u} and the measured or reconstructed phase-space vector is given by,

$$\mathbf{m} = \mathbf{u} + \boldsymbol{\delta}, \quad (7.10)$$

where $\boldsymbol{\delta}$ is some statistical error on the measurement. The covariance matrix for the true phase-space distribution can then be written down in component form as a function of the individual covariances,

$$\Sigma_{ij} = \text{Cov}(m_i, m_j) = \langle m_i m_j \rangle - \langle m_i \rangle \langle m_j \rangle. \quad (7.11)$$

Using the definition for the measurement error and expanding,

$$\begin{aligned} \text{Cov}(m_i, m_j) &= \text{Cov}((u_i + \delta_i), (u_j + \delta_j)), \\ &= \text{Cov}(u_i, u_j) + \text{Cov}(u_i, \delta_j) + \text{Cov}(\delta_i, u_j) + \text{Cov}(\delta_i, \delta_j). \end{aligned} \quad (7.12)$$

This can be used to identify two covariance matrix corrections, \mathbf{R} and \mathbf{C} , such that the reconstructed covariance matrix may be related to the true covariance matrix by,

$$\boldsymbol{\Sigma}^{\text{meas}} = \boldsymbol{\Sigma}^{\text{true}} + \mathbf{R} + \mathbf{R}^T + \mathbf{C}, \quad (7.13)$$

where \mathbf{R} corresponds the correlations between the measurement error and the true state vector parameters and \mathbf{C} corresponds to the covariances of the measurement errors.

A detailed study based on the analysis described in section 4.3, may be used to estimate the individual elements of \mathbf{R} and \mathbf{C} , using the widths of the resolutions. In practice however, this is a poor assumption. The momentum reconstruction is derived from position information, hence there are several inherent correlations in both the \mathbf{R} and \mathbf{C} matrices.

Equations 7.14 and 7.15 outline the most significant components of the \mathbf{R} and \mathbf{C} matrices, as calculated using the Monte Carlo data in section 4.3.

$$\mathbf{R}_{\text{MC}} = \begin{pmatrix} 0.075 & 1.11 & -0.161 & -3.18 \\ -0.501 & -5.75 & -0.200 & -1.56 \\ 0.282 & 3.46 & 0.100 & 1.08 \\ 0.181 & 1.49 & -0.362 & -5.38 \end{pmatrix}, \quad (7.14)$$

$$\mathbf{C}_{\text{MC}} = \begin{pmatrix} 0.043 & 0.062 & 0 & 0.027 \\ 0.062 & 2.06 & -0.035 & 0.420 \\ 0 & -0.035 & 0.042 & 0.057 \\ 0.027 & 0.420 & 0.057 & 2.27 \end{pmatrix}. \quad (7.15)$$

Within the \mathbf{C} matrix, the individual x and y variances are consistent with the position resolutions of the detectors. The remaining components however, demonstrate the significance of the measurement correlations on the covariance matrix reconstruction.

These corrections were calculated for a uniform beam, without any momentum corrections applied to the reconstruction. Different beam distributions and different magnetic field configurations would require different correction matrices. Section 4.3 has shown that the resolution of the trackers is dependent on the momentum of the beam, hence a beam with a small momentum spread would require a different correction matrix to a beam with a large momentum spread.

In order to simplify the treatment of this source of systematic error for the analyses detailed throughout the rest of this chapter, the covariance matrix corrections were calculated and applied as single matrix \mathbf{M} . The matrix \mathbf{M} is given by,

$$\mathbf{M} = -\mathbf{R} - \mathbf{R}^T - \mathbf{C}, \quad (7.16)$$

$$M_{ij} = -\text{Cov}(m_i, \delta_j) - \text{Cov}(\delta_i, m_j) - \text{Cov}(\delta_i, \delta_j).$$

7.3.3 Field Alignment Correction

The installation of the trackers into the solenoid bore was a difficult operation due to the tight restrictions on mounting the trackers within a highly confined space. In order to reduce the amount of material within the cooling channel, the bore was constructed such that the tracker did not require any additional fixings, barring a single rubber foot. While this simplified the construction of

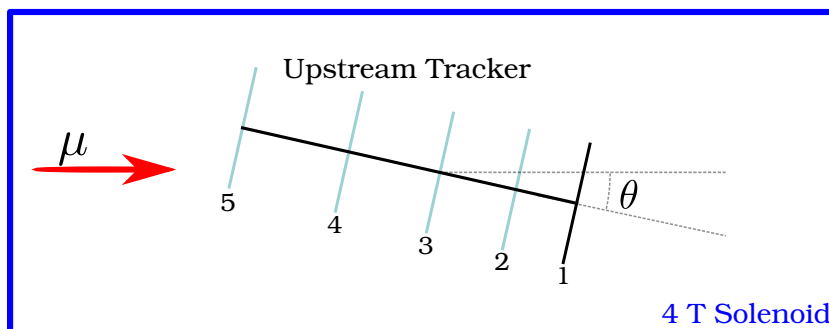


Figure 7.1: A schematic layout of the simulation geometry used to probe the effects of misaligning the tracker with respect to the solenoid field.

the solenoid and reduces the material budget within the cooling channel, the trackers were susceptible to small deviations in position. The magnet coils were also difficult to construct and only rotationally aligned to within a few mrad. Hence there was a concern that the alignment of the trackers to the magnetic field could affect the reconstruction of the beam emittance.

A series of Monte Carlo studies were undertaken in order to investigate the sensitivity of the emittance reconstruction with respect to the field alignment. A standard tracker was simulated within a uniform solenoidal field and exposed to a selection of matched muon beams with normalised transverse emittance values of 3, 6, 8 and 10 mm. The tracker was independently rotated (by angle θ around the vertical y axis) in steps of 2 mrad with respect to the field. As the tracker uses only the mean z -component of the magnetic field to perform the reconstruction, the track reconstruction was performed with no knowledge of the misalignment. Figure 7.1 describes the simulation geometry used in the Monte Carlo Studies.

At each rotation, θ , the standard tracker reconstruction was used to estimate the beam emittance at the tracker reference plane, with no corrections applied, for each of the 4 values of emittance. This was performed using a sample of 200,000 muons at each point, which were required to pass a p -value cut of 0.01 and resulted in between 10 and 100, 2000-muon ensembles for each beam. The mean and standard deviations of the calculated emittances for each ensemble were used to estimate the effect on the reconstruction of the beam emittance. The mean difference in reconstructed emittance, compared to the true beam emittance at the reference plane is shown in figure 7.2.

At large rotations, the number of particles that pass the pattern recognition

selection criteria, and the p-value cut, starts to significantly decrease, hence the statistical errors are seen to increase. Additionally the reconstructed track parameters are statistically much less accurate, further increasing the statistical error.




Figure 7.2: The mean residual between the reconstructed beam emittance and the true beam emittance, evaluated at the reference plane, as a function of the misalignment of the tracker to the field. Four difference emittances were simulated: 3, 6, 8 and 10 mm.

From the results shown, the statistical error on the emittance reconstruction is unbiased by misalignments up to 4 mrad, while a small systematic bias in the absolute value is introduced. At rotations greater than 4 mrad, deviations from the expected resolution become increasingly significant. Additionally, the number of tracks failing to be reconstructed (not shown) greatly increases above 10 mrad, which in turn increases the statistical error.

For small rotations, a small correction term may be calculated as necessary, without any additional modelling. The application of a momentum correction derived for the specific field rotation may prove sufficient to fully account for the field misalignment. For larger rotations a more detailed study would be required in order to determine the ideal progression. As the reconstruction model assumes a co-axial field, it may need to be augmented such that this requirement is more flexible, otherwise the calculated errors and correlations would significantly diverge from the correct values.

The current best estimate for the rotation of the upstream tracker with re-

spect to the field, from a combination of laser survey measurements and the tracker alignment procedures, is approximately 3.3 mrad. A more detailed study, designed to make a direct measurement of the alignment is underway [58]. However, at the time of writing, the study was not mature enough to provide a precise estimate.

As the estimated effect of the field misalignment is small, the decision was taken to estimate the systematic corrections using a model that includes the misalignment that is seen, without the addition of a specific alignment correction. This was achieved through the analysis of Monte Carlo Sample No. 4 in the following section.

7.3.4 Field Uniformity Correction

A second source of systematic error due to the field is the effect of the non-uniformity due to only having three of five magnetic coils powered. In order to precisely estimate the effect of this non-ideal field, four Monte Carlo samples were compared: (1) a perfect 4 Tesla solenoid field with a tracker simulated within; (2) a simple field approximation calculated from the coil dimensions - the native MAUS implementation; (3) an Opera field map [54] which includes the effects of the magnetic materials used in the construction of the cooling channel; (4) the same Opera field map, only rotated with respect to the upstream tracker, in accordance with the best measurements of alignment. In theory the difference in performance between the simple model and the Opera model will provide an estimate for the upper limit on the fluctuations in the reconstruction parameters. Similarly, the ideal field provides a comparison with the expected performance.

Linear momentum corrections (calculated as in section 4.3.2) could then be estimated such that the total momentum of the upstream track reconstruction is globally improved for the four different field models under study. Each track from the 4 datasets was compared against the true track parameters, such that the momentum residual as a function of total momentum could be calculated (figure 7.3). A linear fit to each dataset was performed, the results of which are presented in table 7.1. The application of the correction parameters to the analysis thus improved the reconstruction, as seen in figure 7.3.

Each field configuration generated a correction of a similar order of magnitude; however, due to the complex and subtle differences between each model, it is difficult to observe a distinct trend across them. It can be noted however, that each study was performed with the same number of particles, hence the differences in the errors on each correction parameter are due to the statistical

Field Model	Momentum Intercept [MeV/c]	Momentum Gradient [N/A]
Perfect	-1.075 ± 0.052	0.00320 ± 0.00025
MAUS Model	-1.255 ± 0.063	0.00326 ± 0.00032
Opera Model	-1.99 ± 0.11	0.00723 ± 0.00058
Misaligned Opera Model	-1.47 ± 0.12	0.00487 ± 0.00064

Table 7.1: The results of the linear fits applied to the total momentum residuals for the four different field models that were analysed.




Figure 7.3: Comparison of the total momentum residual before (left) and after (right) a linear correction factor has been applied.

fluctuations in the calculated residuals. Therefore the corrections calculated for the Opera field maps are less precise than those for the MAUS-based or perfect field maps, implying that this may not be the most precise parameterisation.

Overall, the corrections were effective and the total deviation due to the field uniformity is on the order of 0.4 MeV/c total momentum, once the correction has been applied. The corrections were calculated using a uniform beam distribution to (a) demonstrate the principle and (b) remove any biasing from the beam distribution. An issue with this approach is that the correction is sensitive to the transverse distribution as beams with a larger transverse momentum, will be better reconstructed using the Kalman filter than those with lower transverse momentum.

Figure 7.3 also demonstrates the issues with a linear correction parameterisation. The Opera field maps show a distinctly sinusoidal behaviour in the momentum distribution. As the shape of these residuals relies on several non-linear effects, it is not trivial to determine the optimal parameterisation. Hence a simple linear option was used. An alternative parameterisation where each bin in momentum receives a unique correction based on the analysis of a Monte Carlo data set could also be employed. A more detailed study however, would be required in order to correctly estimate the precision of such corrections.

7.3.5 Combined Correction Terms

The four field configurations discussed above were used to analyse the simulated emittance residuals for a range of emittances. Ten values of emittance from 1 to 10 mm, were sequentially used to generate ten beams, each initialised at the entrance to the upstream tracker, with no spread in longitudinal momentum. The nominal 200 MeV/c central momentum was assumed.

For each value of emittance, the reconstructed and true parameters were used in identical emittance calculations, allowing the resulting residuals to be analysed. This was performed multiple times, allowing different correction terms to be applied and examined in turn. Figure 7.4 shows the emittance reconstruction residuals for the uncorrected reconstruction (left) and the reconstruction using the momentum corrections determined previously (right).

The momentum correction has a significant effect on the reconstruction residuals, which was expected from the results of section 4.3.2. The emittance reconstruction residuals with the Opera field maps were reduced to less than 0.01 mm which corresponds to 0.17% for the nominal 6 mm beam. These results suggest that the Opera field map facilitates a more accurate emittance measurement within the trackers. However, this is primarily due to the fluctuations in the

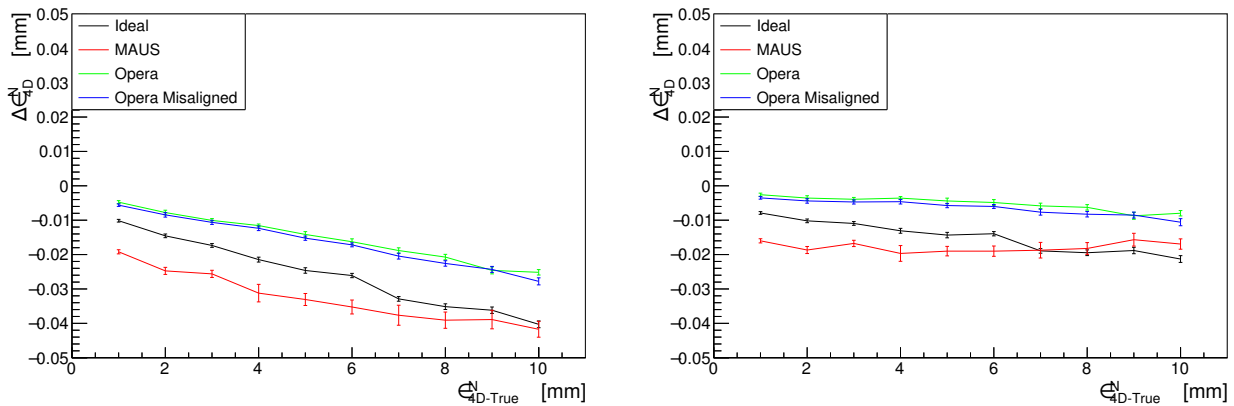


Figure 7.4: The emittance reconstruction residuals of the upstream tracker (left) and the reconstruction residual including the linear momentum correction (right), as predicted by Monte Carlo.

covariance matrix corrections discussed in section 7.3.2, in conjunction with the non-linear nature of the momentum corrections.

Using the simulated data for the 5 mm beam in each field configuration, a field-specific correction matrix was calculated using the procedure described in section 7.3.2. Figure 7.5 shows the final emittance residual when this correction matrix is applied to the reconstruction. By design, the residuals of the nominal 5 mm beam are zero within errors.

These results demonstrate the sensitivity of these correction factors for different emittance beams. Even though the correction was calculated for a 5 mm beam, the residuals are within 0.005 mm of zero for emittances in the range of 3 mm to 10 mm. It is only at very low emittances where the correction matrices must be more carefully estimated.

Early tests of this correction mechanism revealed difficulty in correctly modelling the beam. In these examples, used to demonstrate the principle and test the implementation, the beams were not generated with any correlations between the x and y planes and with zero canonical angular momentum. This is not always a good approximation for a real beamline. As a result, the actual reconstructed emittance is susceptible to further variations in both the momentum corrections and the covariance matrix corrections.

Brief comparisons between the covariance matrix corrections and momentum corrections generally revealed that the momentum corrections would vary most significantly for different beam distributions and field configurations. For this reason the error on the total systematic bias was found to be well represented by considering the 1σ variations in the momentum corrections. For a recon-




Figure 7.5: The emittance reconstruction residuals of the upstream tracker including the momentum corrections and a covariance matrix correction, calculated using the simulated 5 mm emittance beam.

structured emittance, with covariance matrix corrections applied, the momentum corrections were varied to $\pm 1\sigma$ in both the gradient and intercept, and applied to the track reconstruction. The resulting variation in the reconstructed optical parameters provides an appropriate estimate to the potential fluctuations in the systematic corrections. Due to time constraints this was found to be the simplest and most reliable estimate for the systematic error.

7.4 Monte Carlo Analysis

A final test was conducted using an official simulation of the entire MICE experiment, from the MICE target to the EMR. This Monte Carlo simulation is maintained by the MICE collaboration and is undergoing constant improvements to better model the behaviour of the MICE experiment at all positions within both the beamline and the cooling channel. As such it will introduce additional effects and correlations to the transverse phase-space that are not modelled by simple gaussian approximations, but are more similar to the real experiment.

The data were initially analysed without the corrections in order to determine the approximate magnitude of the emittance of the beam and the momentum spread. This allowed the appropriate correction matrix to be selected from the analysis performed previously (section 7.3.5). The uncorrected emittance of the beam was found to be approximately 6.6 mm. The matrix correction for a




Figure 7.6: Distribution of total momentum at the upstream tracker reference plane. The events that were included in this Monte Carlo analysis are highlighted in red.

7 mm beam was therefore used. The same dataset that was used to determine the covariance matrix correction, was also used to determine the momentum correction parameters.

The nominal beam for the MICE experiment is a 200 MeV/c muon beam. However the momentum spread of the beam is sufficiently large to cause considerable chromatic effects within analyses and significantly disrupt the optics. The simplest solution to these issues is to select a momentum window for analysis. Continuing with the nominal optics for the MICE experiment, a selection window of 10 MeV/c was placed on the central momentum of 200 MeV/c. There were 4669 simulated muons present within this momentum window, as illustrated in figure 7.6.

The momentum correction parameters were applied to each track and the appropriate covariance matrix correction was subtracted from the reconstructed covariance matrix. The transverse normalised emittance was calculated to be $6.69 \pm 0.07_{\text{Stat}} \pm 0.08_{\text{Sys}}$ mm. The statistical error was calculated using equation 7.9 and the systematic error is representative of the total variation in emittance, calculated at the limits of the momentum corrections. The Monte Carlo true emittance was calculated using an identical calculation, without any correction factors, to be 6.65 mm. Additional parameters as reconstructed from the simulated data are described in table 7.2.

Parameter	Recon.	Stat. Err	Sys. Err	MC	Units
ϵ_{\perp}	6.69	± 0.07	± 0.08	6.652	mm
ϵ_x	7.05	± 0.10	± 0.03	7.012	mm
ϵ_y	7.55	± 0.11	± 0.04	7.558	mm
β_{\perp}	211.0	± 2.0	± 2.0	213.9	mm
β_x	201.0	± 3.0	± 1.0	203.9	mm
β_y	187.0	± 3.0	± 2.0	187.4	mm
α_{\perp}	0.545	± 0.006	± 0.030	0.555	-
α_x	0.680	± 0.010	± 0.030	0.686	-
α_y	0.341	± 0.005	± 0.020	0.340	-

Table 7.2: Reconstruction results of the key optical parameters as calculated using the official Monte Carlo simulation. 4669 muons were included in the calculation following the beam selection. A statistical sampling error due to the number of muons selected and systematic error due to the momentum bias are included. The true Monte Carlo values are included for comparison.

7.5 Reconstruction Results

The data obtained from Dataset C (table 5.1) was subjected to the standard reconstruction algorithms present in MAUS, as described in chapter 3. Each event was then required to pass the following cuts in order to be included within the analysis:

1. Precisely one spacepoint must be present in TOF1. As the transmission to the downstream TOF2 detector was poor, only TOF1 could be used to ensure high transmission through SSU. Additionally, multiple spacepoints would imply either a secondary particle or a decay event, in both cases the event would be vetoed;
2. Precisely one track must be present in SSU. Additional tracks would indicate that an event beyond what was expected occurred, i.e. a decaying muon or pion, a high fluctuation in knock-on electrons or noise, or a secondary muon/pion track;
3. P-value > 0.01 . This ensures that the worst of the fitted tracks are removed from the analysis;
4. Total momentum must be within a 10MeV/c momentum window, centred on 200MeV/c. This ensured that the chromatic effects are reduced.

Following the application of these cuts, 5049 muons remained for analysis.

An initial analysis was performed in order to determine the approximate




Figure 7.7: Distribution of total momentum at the upstream tracker reference plane for Run 07469. The events that were included in this analysis are highlighted in red.

beam emittance and momentum spread. This initial measurement produced a value of 5.03 mm. The same set of momentum correction parameters that were applied to the Monte Carlo model were used during the subsequent analysis stages and the corresponding 5 mm matrix correction was used to determine the final covariance matrix.

Figure 7.7 demonstrates the distribution of total momentum at the upstream reference plane, following the cuts outlined above. The highlighted area corresponds to the selected events. Figure 7.8 shows the x and y transverse phase-space distributions at the upstream reference plane following all the cuts.

The transverse normalised emittance of the muon beam, calculated from the 5049 tracks within the 10 MeV/ c momentum window, at 200 MeV/ c , as measured at the reference plane of the upstream tracker, was calculated to be $5.08 \pm 0.05_{\text{Stat}} \pm 0.06_{\text{Sys}} \pm 0.02_{\text{Field}}$ mm. Where the statistical error was calculated using equation 7.9, the systematic error was estimated by the limits of the momentum corrections, and the field precision error is representative of the effect of different field models on the emittance residuals. The complete set of optical parameters are outlined in table 7.3, including the statistical and systematic errors.

A distinct disagreement between the actual field and the model would result in an unknown bias in both the momentum corrections and the covariance matrix corrections. Hence the field precision error is currently only an estimate, based on the typical effect of changing the field map within the Monte Carlo as seen in figure 7.4.




Figure 7.8: The x (left) and y (right) phase-space distributions for the reconstructed beam at the upstream reference plane for Run 07469.

Parameter	Recon.	Stat. Err	Sys. Err	Units
ϵ_{\perp}	5.08	± 0.05	± 0.06	mm
ϵ_x	5.31	± 0.08	± 0.03	mm
ϵ_y	6.38	± 0.09	± 0.03	mm
β_{\perp}	232.0	± 2.0	± 3.0	mm
β_x	219.0	± 3.0	± 0.2	mm
β_y	188.0	± 3.0	± 2.0	mm
α_{\perp}	0.618	± 0.006	± 0.032	-
α_x	0.775	± 0.010	± 0.027	-
α_y	0.340	± 0.005	± 0.016	-

Table 7.3: Reconstruction results of the key optical parameters as calculated from Dataset C. 5049 muons were included in the calculation following a momentum selection and p-value cut. A statistical sampling error due to the number of muons selected and systematic error due to the momentum bias are included. An additional systematic error for the uncertainty in the field map is not included.

7.6 Conclusion

A high precision emittance measurement has been calculated using the data from the MICE Step IV commissioning period. The overall accuracy has been shown to be good within the assumption of a correct field model. However, the measurement could further be improved through the reduction of both statistical and systematic errors. Due to the relatively small data set, the statistical uncertainty will not be improved until commissioning of the spectrometer solenoid magnets is resumed. Due to the $\approx 1/\sqrt{n}$ scaling, greater statistics will provide increasingly precise measurements.

The systematic uncertainty due to the knowledge of the magnetic field requires further development, in order to reduce the effect on the final measurements. If it is assumed (or at some point demonstrated) that the field is correctly mapped to some level of precision, a selection of field maps could be generated that are perturbed within the limits of that precision. By analysing the simulated emittance residual for each of the field maps, an estimate for the overall uncertainty on the emittance, due to the field map, could be derived from the spread of residuals. This would also permit the variation in covariance matrix corrections to be determined within the limits of the known field precision. It is likely however, that with optimal estimates for all systematic corrections, the precision of the field map will become the limiting systematic error for an absolute emittance measurement.

The overall accuracy of the covariance matrix corrections could also be improved by using a model that better represents the actual muon beam than the simple gaussian approach used above. Ideally the corrections would be calculated using a simulation of the beam phase-space that was directly derived from a measurement of the beam. Either another detector (TOF1), or the tracker itself, could be used to provide an approximate parameterisation of the beam, including the canonical angular momentum. Covariance matrix corrections that were derived from this simulated beam should offer a more precise correction to the reconstructed phase-space than those used above.

To further this approach, an iterative process could be used in order to converge on an more accurate set of corrections. Firstly an uncorrected measurement of the beam phase-space would be made. The parameterisation of the phase-space, using the measurement, could then be used in a simulation designed to generate the covariance matrix corrections for that specific beam. The calculated covariance matrix corrections could then be applied to the data reconstruction, to improve upon the initial measurement. The process could then be repeated whereby the corrected reconstruction could be used to provide

an increasingly precise beam model for the simulation stage. In theory, this should quickly converge on the ideal correction matrix. At present a study of the practicality of this method, or a demonstration of the performance has not been performed due to time constraints.

A more radical improvement, following the detailed systematics studies, would be to improve the track fitting process. The possibility of using an adaptive filter has been discussed previously, allowing for more precise modelling of energy loss and/or multiple Coulomb scattering processes. Furthermore, the Kalman filter could be further improved by using a Runge-Kutta propagation algorithm, rather than the coarse Euler-steps in the current implementation. A higher granularity propagation algorithm would permit a field map to be used more effectively.

Both the alignment and uniformity of the field would be correctly modelled at each step in the propagation process, which would result in a more precise determination of the track parameters. It is believed that this would improve the accuracy of the momentum reconstruction and reduce the scale of the calculated covariance matrix corrections. Detailed testing of the algorithm would be imperative to ensure that the statistical properties of the fit are either maintained or improved. At the time of writing this concept is under development.

Although the measurement is dominated by complex systematic effects, the optical properties of the muon beam, within the selected momentum window, have been accurately estimated within the assumptions made. Although there is a great scope for improvement, given more time to invest, the procedure has been demonstrated to be accurate with Monte Carlo studies and was successfully applied to data with a conservative error analysis. This analysis corresponds to one of the few high-precision emittance measurements made on a muon beam and will serve to inform future studies.

Chapter 8

Conclusions

Muon accelerators represent an exciting future for particle physics: Muon colliders offer the measurement precision of a lepton collider, but without the difficulties in overcoming synchrotron radiation. This allows the properties of the Higgs Boson to be precisely measured in an environment with much lower levels of noise and fewer backgrounds - a scenario unparalleled by hadron colliders. Additionally, muon beams can offer very pure, high luminosity neutrino beams, capable of providing a much greater sensitivity to potential CP violation and the mass hierarchy, than current experiments.

Despite the substantial efforts in research, there still exists some technological difficulties for the construction of these future accelerators. The last remaining component, required for a high luminosity muon beamline to be produced, is the process of ionisation cooling. The Muon Ionisation Cooling Experiment is charged with demonstrating and measuring the physics processes required for ionisation cooling to occur sustainably, in addition to making the first measurement of normalised transverse emittance reduction of a muon beam using ionisation cooling. Monte Carlo simulations can predict the precision and accuracy to which this is believed to be possible, and the hardware is currently being constructed and commissioned, with a complete demonstration of ionisation cooling expected by 2019.

In order to make such precise measurements, Scintillating Fibre (SciFi) Trackers were designed and constructed for use within the MICE Cooling Channel. Monte Carlo studies have been used to demonstrate the predicted resolution of these detectors, capable of reconstructing both straight and helical muon tracks in solenoidal fields. The position resolution has been shown to be approximately 0.2 mm, with an angular resolution (for straight tracks) of approximately 4 mrad, and a transverse momentum resolution (for helical tracks)

of approximately 0.88 MeV/c.

A detailed assessment of the implemented reconstruction framework has been used to validate the assumptions made during its design, in addition to a qualitative analysis of the weaknesses of the algorithm. The Kalman Filter based track fit has been shown to reliably and precisely reconstruct muon tracks, requiring only a small correction to the total momentum. The source of this correction stems from the use of a simplified reconstruction geometry, designed to streamline the online reconstruction of data. A more precise model of the materials that compose the trackers would improve the energy loss estimates for reconstructed tracks, which is expected to reduce the magnitude of the momentum corrections.

There is an additional issue with the multiple Coulomb scattering model. At present, the model overestimates the amount of scattering between measurement planes, which results in an overestimate of the statistical errors on the track parameters. The statistical measures (p-values and pulls) of the reconstructed parameters therefore differ from an ideal implementation. The accuracy of the reconstruction is not detrimentally affected. However, it does create difficulties in the justification of statistical cuts. For example, a p-value $> 5\%$ is a common requirement to ensure a high-quality reconstructed track. However, that implicitly expects a uniform p-value distribution. Statistical cuts may still be applied to good effect, however the justification and statistical understanding is more complicated.

Comparisons between Monte Carlo simulations and real data reconstructions have shown excellent agreement, the effects seen in the Monte Carlo are almost precisely reproduced in real data, requiring only small corrections to the precise positioning of measurement planes to improve the agreement. This demonstrates that the Monte Carlo model of the trackers closely agrees with the true detectors in both the physical description and the performance of the track fit. This strongly suggests that the performance seen in the Monte Carlo accurately represents the performance of the real detector.

Following the range of validation analyses, an initial study of the alignment of the two trackers was conducted. Due to the difficulties in installing and surveying each detector, this was both a necessary step, in order to precisely determine their place within the experiment hall; and a practical first test of the reconstructed data. Data from 11 different runs was split into two datasets and used to determine the transverse position and rotation of the downstream tracker reference plane, with respect to the upstream tracker. The transverse position resolution was 0.11 mm and the transverse rotation resolution was

0.05 mrad. A significant misalignment was found, and discussions are currently underway to address these issues.

A final study was conducted in order to measure the natural emittance of the MICE muon beam within the upstream tracker. Several systematic studies were conducted in order to determine the sensitivity of the track reconstruction to (a) misalignments between the tracker and the solenoid field in which it sits, and (b) the uniformity of the solenoidal field across the length of the tracker. No formal validation of the precision of the field was available, therefore four different models were investigated in order to estimate the approximate sensitivity. Momentum corrections for the track reconstruction were calculated for each field and shown to accurately correct the total momentum.

The calculation of covariance matrix corrections was also used in order to determine how correlations within the beam, and within the reconstruction of the beam, affect the final emittance measurement. A set of correction matrices, corresponding to beams with normalised transverse emittances within the range of 1 to 10 mm, were calculated and applied to a Monte Carlo beam. This resulted in the reconstruction of the emittance within the fluctuation of systematic errors.

A momentum window of 195 to 205 MeV/c was imposed upon the reconstructed tracks in addition to a p-value cut of 0.01. An ensemble of 5049 reconstructed tracks were then used to reconstruct the covariance matrix that described the muon beam at the reference plane of the upstream tracker. A Monte Carlo comparison using a dataset that was designed to replicate the results of these data was conducted in parallel and shown to accurately reproduce the true beam emittance. The statistical error on the emittance reconstruction was calculated and the final value of the normalised transverse emittance, at the reference plane of the upstream tracker, was measured to be $5.08 \pm 0.05_{\text{Stat}} \pm 0.06_{\text{Sys}} \pm 0.02_{\text{Field}}$.

Several improvements could still be applied to this process. The aforementioned additions to the reconstruction would improve the reconstruction of each track, thereby reducing the need for momentum corrections. However, there will still be some effect from the field uniformity. The primary area for improvement however, is the determination of the covariance matrix corrections. The beams that were used to generate the correction matrices were not modelled with different values of canonical angular momentum - the correlations between x and y positions and momentum. Hence the corrections did not correctly model the systematic biases that would result from these correlations. The analysis of the beam from other detectors, or from the trackers themselves, could be used to better determine the ideal correction matrix.

Bibliography

- [1] M. Palmer, “An overview of the US Muon Accelerator Program,” in *COOL13*, 2013.
- [2] J.-P. Delahaye *et al.*, “Enabling Intensity and Energy Frontier Science with a Muon Accelerator Facility in the U.S.: A White Paper Submitted to the 2013 U.S. Community Summer Study of the Division of Particles and Fields of the American Physical Society,” in *Community Summer Study 2013: Snowmass on the Mississippi (CSS2013) Minneapolis, MN, USA, July 29-August 6, 2013*, 2013.
- [3] P. Kyberd *et al.*, “Neutrinos from stored muons letter if intent to the fermilab physics advisory committee,” *arXiv:1206.0294*, 2012.
- [4] D. Adey *et al.*, “Neutrinos from STOREd Muons. proposal to Fermilab PAC,” *arXiv:1308.6822v1*, 2013.
- [5] A. Letchford *et al.*, “Status of the RAL Front End Test Stand,” in *IPAC Conf. Proc.*, 2012.
- [6] K. McDonald *et al.*, “The MERIT high-power target experiment at the CERN PS,” in *IPAC10*, 2010.
- [7] S. Bogacz, “Beam dynamics of low energy muon acceleration,” *Nuclear Physics B - Proceedings Supplements*, vol. 155, no. 1, pp. 334 – 335, 2006. Proceedings to the 7th International Workshop on Neutrino Factories and Superbeams Proceedings to the 7th International Workshop on Neutrino Factories and Superbeams.
- [8] D. J. Summers, L. M. Cremaldi, R. Godang, B. R. Kipapa, H. E. Rice, and R. B. Palmer, “Muon acceleration to 750 GeV in the Tevatron tunnel for a 1.5 TeV $\mu^+\mu^-$ collider,” *Conf. Proc.*, vol. C070625, p. 3178, 2007. [IEEE Nucl. Sci. Symp. Conf. Rec.,3178(2007)].

- [9] A. N. Skrinsky *et al.* *Sov. J. Part. Nucl.*, vol. 12, no. 223, 1981.
- [10] Muon Ionisation Cooling Experiment. <http://www.mice.iit.edu/>.
- [11] M. Bogomilov *et al.*, “The MICE Muon Beam on ISIS and the beam-line instrumentation of the Muon Ionization Cooling Experiment,” *JINST*, vol. 7, p. 5009, 2012.
- [12] M. Reiser, *Theory and Design of Charged Particle Beams*. Wiley-VCH, 2008.
- [13] G. Penn and J. S. Wurtele, “Beam envelope equations for cooling of muons in solenoid fields,” *Phys. Rev. Lett.*, vol. 85, pp. 764–767, Jul 2000.
- [14] S. van der Meer, “Stochastic cooling and the accumulation of antiprotons,” *Rev. Mod. Phys.*, vol. 57, pp. 689–697, Jul 1985.
- [15] G. I. Budker, “An effective method of damping particle oscillations in proton and antiproton storage rings,” *Soviet Atomic Energy*, vol. 22, no. 5, pp. 438–440, 1967.
- [16] G. I. Budker *et al.* *Particle Accelerators*, vol. 7, no. 197.
- [17] E. G. Bessonov and K.-J. Kim, “Radiative cooling of ion beams in storage rings by broad-band lasers,” *Phys. Rev. Lett.*, vol. 76, pp. 431–434, Jan 1996.
- [18] J. Beringer *et al.*, “Review of particle physics,” *Phys. Rev. D*, vol. 86, 2012. <http://pdg.lbl.gov/>.
- [19] R. C. Fernow and J. C. Gallardo, “Muon transverse ionization cooling: Stochastic approach,” *Phys. Rev. E*, vol. 52, pp. 1039–1042, Jul 1995.
- [20] B. Rossi, *High Energy Particles*. Prentice Hall, Inc., Englewood Cliffs, 1952.
- [21] T. Carlisle and J. Cobb, “Ionisation cooling in MICE Step IV,” in *IPAC11*, 2011.
- [22] R. Palmer, “Emittance and cooling Lecture II: Ionization cooling,” in *MICE Collaboration Meeting 38.*, 2014. <https://indico.cern.ch/event/290685/>.
- [23] The MICE Collaboration, “Design and expected performance of the MICE demonstration of ionisation cooling,” *Forthcoming*, 2016.

- [24] D. J. S. Findlay, “Isis - pulsed neutron and muon source,” in *2007 IEEE Particle Accelerator Conference (PAC)*, pp. 695–699, June 2007.
- [25] M. Bonesini, “Progress of the MICE experiment,” *PoS*, vol. EPS-HEP2015, p. 521, 2015.
- [26] C. N. Booth *et al.*, “The design, construction and performance of the MICE target,” *To be published in JINST*, 2013. <http://arxiv.org/abs/1211.6343>.
- [27] D. Orestano, “The detector system of the MICE experiment,” *Nucl. Inst. Phys. A*, vol. 617, pp. 45 – 47, 2010.
- [28] D. M. Kaplan, M. Drews, D. Rajaram, M. Winter, L. Cremaldi, D. Sanders, and D. Summers, “Progress on Cherenkov Reconstruction in MICE,” in *17th International Workshop on Neutrino Factories and Future Neutrino Facilities Search (NuFact15) Rio de Janeiro, Brazil, August 10-15, 2015*, 2016.
- [29] R. Bertoni, A. Blondel, M. Bonesini, G. Cecchet, A. de Bari, J. Graulich, Y. Karadzhov, M. Rayner, I. Rusinov, R. Tsenov, S. Terzo, and V. Verguilov, “The design and commissioning of the MICE upstream time-of-flight system,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 615, no. 1, pp. 14 – 26, 2010.
- [30] U. Bravar *et al.*, “MICE: the Muon Ionization Cooling Experiment. Step I: First Measurement of Emittance with Particle Physics Detectors,” in *Particles and fields. Proceedings, Meeting of the Division of the American Physical Society, DPF 2011, Providence, USA, August 9-13, 2011*, 2011.
- [31] F. Ambrosino *et al.*, “Calibration and performances of the KLOE calorimeter,” *Nucl. Instrum. Meth.*, vol. A598, pp. 239–243, 2009.
- [32] R. Asfandiyarov, “A totally active scintillator calorimeter for the Muon Ionization Cooling Experiment (MICE). Design and construction,” *Nucl. Instrum. Meth.*, vol. A732, pp. 451–456, 2013.
- [33] M. Bogomilov *et al.*, “Pion contamination in the MICE muon beam,” 2015.
- [34] S. Ozaki, R. Palmer, M. Zisman, and J. Gallardo, “Feasibility study-II of a muon based neutrino source,” tech. rep., BNL-52623, 2001.

- [35] D. Rajaram and P. Snopok, “Muon ionisation cooling experiment Step VI,” in *IPAC13, Shanghai*, 2013.
- [36] M. R. Wayne, “The d0 tracker upgrade and projections for b physics,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 351, no. 1, pp. 77 – 83, 1994.
- [37] D. Adey, *Beam instrumentation and investigations into muon cooling at MICE*. PhD thesis, Warwick University, 2012.
- [38] C. D. Tunnel and C. T. Rogers, “MAUS: Mice User Analysis Software,” in *IPAC Conf. Proc.*, 2011.
- [39] M. Buehler *et al.*, “Magnetic measurements for a MICE spectrometer solenoid,” *MICE Internal Note - No. 397*, 2012.
- [40] A. Dobbs. Private Communication, 2016.
- [41] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35 – 45, 1960.
- [42] G. C. Goodwin and K. S. Sin, *Adaptive Filtering, Prediction and Control*. Prentice-Hall, Inc., 1984.
- [43] R. Fruhwirth, “Application of kalman filtering to track and vertex fitting,” *Nucl. Instrum. Methods Phys. Res., A*, vol. 262, p. 444. 19 p, Jun 1987.
- [44] F. J. P. Soler *et al.*, “Kalman filter tracking and vertexing in a silicon detector for neutrino physics,” *Nucl. Instrum. Methods Phys. Res., A*, vol. 486, pp. 639–662, 2002.
- [45] H. Bichsel, “Straggling in thin silicon detectors,” *Rev. Mod. Phys.*, vol. 60, pp. 663–699, Jul 1988.
- [46] S. Agostinelli *et al.*, “Geant4 - a simulation toolkit,” *Nucl. Instr. Meth. A*, vol. 506, pp. 250 – 303, 2003.
- [47] H. W. Lewis, “Multiple scattering in an infinite medium,” *Phys. Rev.*, vol. 78, pp. 526–529, Jun 1950.
- [48] V. L. Highland, “Some practical remarks on multiple scattering,” *Nuclear Instruments and Methods*, vol. 129, no. 2, pp. 497 – 499, 1975.

- [49] The MICE Collaboration, “The reconstruction software for the mice scintillating fibre trackers,” *Forthcoming*, 2016.
- [50] M. Uchinda, A. Dobbs, and P. Kyberd. Private Communication, 2016.
- [51] A. Dobbs, C. Hunt, K. Long, E. Santos, and C. Heidt, “The reconstruction software for the mice scintillating fibre trackers,” *MICE Note No. 451*, 2014.
- [52] M. Ellis, P. Hobson, P. Kyberd, J. Nebrensky, A. Bross, J. Fagan, T. Fitzpatrick, R. Flores, R. Kubinski, J. Krider, R. Rucinski, P. Rubinov, C. Tolian, T. Hart, D. Kaplan, W. Luebke, B. Freemire, M. Wojcik, G. Barber, D. Clark, I. Clark, P. Dornan, A. Fish, S. Greenwood, R. Hare, A. Jamdagni, V. Kasey, M. Khaleeq, J. Leaver, K. Long, E. McKigney, T. Matsushita, C. Rogers, T. Sashalmi, P. Savage, M. Takahashi, A. Tapper, K. Yoshimura, P. Cooke, R. Gamet, H. Sakamoto, Y. Kuno, A. Sato, T. Yano, M. Yoshida, C. MacWaters, L. Coney, G. Hanson, A. Klier, D. Cline, X. Yang, and D. Adey, “The design, construction and performance of the {MICE} scintillating fibre trackers,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 659, no. 1, pp. 136 – 153, 2011.
- [53] E. Overton, “Data quality/calibration,” in *MICE Video Conference No. 183*, 2016. <https://indico.cern.ch/event/502816/>.
- [54] Opera FEA - Simulation Software, 2015. See also <http://operafea.com/>.
- [55] F. Drielsma. Private Communication - Publication forthcoming, 2016.
- [56] V. Blobel and C. Kleinwort, “A New method for the high precision alignment of track detectors,” in *Advanced Statistical Techniques in Particle Physics. Proceedings, Conference, Durham, UK, March 18-22, 2002*, pp. URL-STR(9), 2002.
- [57] A. Bross, “Spectrometer solenoid plans for step IV,” in *MICE Collaboration Meeting 44*, 2016. <https://indico.cern.ch/event/485764/>.
- [58] C. Rogers, “Tracker to solenoid alignment,” Collaboration Meeting 43, MICE, 2015. <https://indico.cern.ch/event/436141/>.
- [59] M. Kendall, *Kendall’s Advanced Theory of Statistics*. London: Griffin, 5th ed., 1987.

- [60] R. Brun and F. Rademakers, “Root - an object oriented data analysis framework,” *Nucl. Inst. & Meth. in Phys. Res. A*, vol. 389, pp. 81–86, September 1997. See also <http://root.cern.ch/>.

Appendix A

Pattern Recognition in MAUS

A.1 The Method of Least Squares

The method of least squares is a standard tool used for data analysis. It provides a powerful and simple method to estimate the optimal solution to a set of overdetermined equations, that is there are more constraints than unknowns. This can be interpreted as fitting a model parameterised by a vector of parameters, $\boldsymbol{\beta}$, of length N , to a set of M data points, $\mathbf{y} = \{y_1, y_2, y_3 \dots y_M\}$.

In order to apply the method, the problem must be correctly constructed. In the most general form, each measurement is allowed to depend on multiple independent variables, x_i , and can be written as some unknown function, parametrised by the set of parameters, β ,

$$y_i = \boldsymbol{\beta}^T \mathbf{x} = \beta_1 + \beta_2 x_2 + \beta_3 x_3 + \dots + \beta_N x_N, \quad (\text{A.1})$$

where x_1 was assumed to be 1, allowing for a constant term. This can be generalised to include the whole dataset \mathbf{y} , by constructing a matrix \mathbf{X} that contains all the independent variables, hence,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} = \begin{pmatrix} 1 & x_{12} & x_{13} & \cdots & x_{1N} \\ 1 & x_{22} & x_{23} & \cdots & x_{2N} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{M2} & x_{M3} & \cdots & x_{MN} \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_N \end{pmatrix}. \quad (\text{A.2})$$

A scalar function, $S(\boldsymbol{\beta})$ can then be defined such that the optimal linear solution represents a minimum of S . As S is positive-definite by definition, a value of zero represents a perfect minimum in that all constraints are precisely met by the dataset provided. Although theoretical systems can be designed to emulate this effect, it is not typically seen in data analysis. Rather some minimisation of S provides an optimal solution to the problem:

$$S(\boldsymbol{\beta}) = |\mathbf{y} - \mathbf{X}\boldsymbol{\beta}|^2 = (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})^\top (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}). \quad (\text{A.3})$$

In a typical case, however, the data points are accompanied by an associated error, or more generally a covariance matrix, $\mathbf{V} = \mathbf{W}^{-1}$, which effectively acts as to weight each measurement according to it's accuracy - the measurements with the smallest error should provide the most accurate information. Hence, equation A.3 may be updated to reflect this:

$$S(\boldsymbol{\beta}) = (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})^\top \mathbf{W}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}). \quad (\text{A.4})$$

Note that assuming the measurement errors are gaussian distributed and the parameters $\boldsymbol{\beta}$ are known, repeated calculations of S with different datasets would actually form a χ^2 distribution, hence this is commonly referred to as Chi-Squared Minimisation.

This problem can be analytically solved in matrix form, by calculating the values of $\boldsymbol{\beta}$ which allows the derivative of S to equal zero,

$$\frac{dS}{d\boldsymbol{\beta}} = \frac{d}{d\boldsymbol{\beta}} \left[\mathbf{y}^\top \mathbf{W} \mathbf{y} - \mathbf{y}^\top \mathbf{W} \mathbf{X} \boldsymbol{\beta} - \boldsymbol{\beta}^\top \mathbf{X}^\top \mathbf{W} \mathbf{y} + \boldsymbol{\beta}^\top \mathbf{X}^\top \mathbf{W} \mathbf{X} \boldsymbol{\beta} \right] = 0. \quad (\text{A.5})$$

Making use of (a) each term is a scalar and hence is transpose invariant, and (b) $\mathbf{W} = \mathbf{W}^\top$,

$$\frac{dS}{d\boldsymbol{\beta}} = \frac{d}{d\boldsymbol{\beta}} \left[\boldsymbol{\beta}^\top \mathbf{X}^\top \mathbf{W} \mathbf{X} \boldsymbol{\beta} - 2\mathbf{y}^\top \mathbf{W} \mathbf{X} \boldsymbol{\beta} \right], \quad (\text{A.6})$$

hence the solution can then be derived:

$$\boldsymbol{\beta} = (\mathbf{X}^\top \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{W} \mathbf{y}. \quad (\text{A.7})$$

The routines as written in the MAUS framework implement equation A.7 to perform the least squares fits to the track information.

A.2 Straight Track Fitting

For the straight track reconstruction, the x and y transverse coordinates are independent, and can be treated as such. For an arbitrary transverse plane, e.g. the x - z plane, the track model is the typical straight line parameterisation,

$$\begin{aligned} x &= X_0 + zt_x, \\ y &= Y_0 + zt_y, \end{aligned} \tag{A.8}$$

where X_0 is the intercept of the track with the reference plane and t_x is the gradient of the track in the x - z plane. Note that this model does not allow for any stochastic effects. The parameters in the β vector correspond to the intercept and gradient, with one independent variable, z_i corresponding to the z -position of the tracker stations. The dependent variables, i.e. the spacepoint measurements m_i , correspond to the y -vector. For the x - z plane the system is described by,

$$\mathbf{y} = \mathbf{X}\beta = \begin{pmatrix} 1 & z_0 \\ 1 & z_1 \\ 1 & z_2 \\ 1 & z_3 \\ 1 & z_4 \end{pmatrix} \begin{pmatrix} X_0 \\ t_x \end{pmatrix} = \begin{pmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}. \tag{A.9}$$

The covariance matrix is determined from the error on each measurement. The measurement error is assumed to be the variance of top-hat function with width equal to the channel width, a . Hence the covariance matrix \mathbf{V} for the x - z system is given by,

$$\mathbf{V} = \begin{pmatrix} a/\sqrt{12} & 0 & 0 & 0 & 0 \\ 0 & a/\sqrt{12} & 0 & 0 & 0 \\ 0 & 0 & a/\sqrt{12} & 0 & 0 \\ 0 & 0 & 0 & a/\sqrt{12} & 0 \\ 0 & 0 & 0 & 0 & a/\sqrt{12} \end{pmatrix}. \tag{A.10}$$

A.3 Helical Track Fitting

For the helical track fit, the model is inherently non linear. An analytical treatment without approximations is therefore not possible, however using a

specific formulation of the equation for a helix, a fit is still possible. The goal of pattern recognition is not to provide a precise set of reconstruction parameters for each track, rather it requires enough precision to determine whether a viable track has been found.

Firstly, the fit is decomposed into transverse and longitudinal coordinates. In the transverse plane, a circle fit is required, while in the longitudinal plane a non-linear sinusoidal fit is required. This can be avoided by using the results of the circle fit to estimate the rotation around the circumference for each spacepoint, hence a fit to the path length of the track with respect to the z position is performed instead. This corresponds to a straight line fit.

For the circle fit, the model is described in the x - y plane by the radius ρ , and the circle centre (X_0, Y_0) ,

$$(x - X_0)^2 + (y - Y_0)^2 = \rho^2, \quad (\text{A.11})$$

which can be more practically parameterised as,

$$\alpha(x^2 + y^2) + \beta x + \gamma y = \kappa; \quad (\text{A.12})$$

$$\begin{aligned} \alpha &= \frac{1}{\rho^2 - (X_0^2 + Y_0^2)}, \\ \beta &= -2X_0\alpha, \\ \gamma &= -2Y_0\alpha, \\ \kappa &= 1. \end{aligned}$$

In equation A.12, if the definition of dependent and independent variables is suspended, the model parameters form a linear polynomial. This is not a bad assumption as it implies that the parameter κ is the dependent variable and the x and y values are independent variables. Formulating the fit for this system will therefore attempt to generate a set of parameters for which κ is closest to 1,

$$\mathbf{y} = \mathbf{X}\beta = \begin{pmatrix} (x_0^2 + y_0^2) & x_0 & y_0 \\ (x_1^2 + y_1^2) & x_1 & y_1 \\ (x_2^2 + y_2^2) & x_2 & y_2 \\ (x_3^2 + y_3^2) & x_3 & y_3 \\ (x_4^2 + y_4^2) & x_4 & y_4 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix} = \begin{pmatrix} \kappa \\ \kappa \\ \kappa \\ \kappa \\ \kappa \end{pmatrix}. \quad (\text{A.13})$$

As all measurements have an equivalent error, the corresponding covariance

matrix would be diagonal, hence it can be factored out of the calculation. Although conceptually this is not correct, in practice it provides a simple method of applying a circle fit to the spacepoint data. The resulting covariances and correlations would likely be incorrectly modelled, however for the coarse requirements of pattern recognition this was found to be more than sufficient.

The results of the circle fit are then used to assign each spacepoint a value of s , the path length around the corresponding circle circumference. The second stage - a fit in the s - z plane - is then conducted similarly for the straight line fit where t_s is the gradient in s - z space and S_0 is the intersection of the path length at the reference plane:

$$s = S_0 + zt_s, \quad (\text{A.14})$$

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} = \begin{pmatrix} 1 & z_0 \\ 1 & z_1 \\ 1 & z_2 \\ 1 & z_3 \\ 1 & z_4 \end{pmatrix} \begin{pmatrix} S_0 \\ t_s \end{pmatrix} = \begin{pmatrix} s_0 \\ s_1 \\ s_2 \\ s_3 \\ s_4 \end{pmatrix}. \quad (\text{A.15})$$

In all cases the matrix calculation is performed using the algorithms provided by the ROOT framework [60].

Appendix B

Reconstruction Validation Data

A complete set of plots, as produced from the reconstruction of Datasets A & B (table 5.1) is detailed below. Three different plots are included:

1. The raw distributions of reconstructed digits,
2. The reconstructed trackpoint pulls for each plane,
3. The distribution of Kuno-sums for spacepoints that are included in the track fit.

The distribution of reconstructed digits clearly reveals the channels that produce excessive numbers of digits due to noise, in addition to dead channels.

The trackpoint pulls show an excellent agreement between Monte Carlo and data, which results in the close agreement between simulated and reconstructed p-value distributions (section 5.4). Additionally, the difference in widths between the first and last planes in each tracker is attributed to the poor modelling of multiple Coulomb scattering, which in turn incorrectly weights the error associated with each measurement.

The distributions of Kuno-sums demonstrate the spacepoint production and channel cabling is correct. Together with the pull distributions, the misalignments between neighbouring planes within a station can be observed. The asymmetry in each distribution can be reproduced by Monte Carlo by modelling a small displacement of the centre of each plane. This was attempted for the upstream tracker, but only approximately. The improvement in agreement between Monte Carlo and data can be clearly seen when contrasted with the downstream tracker.

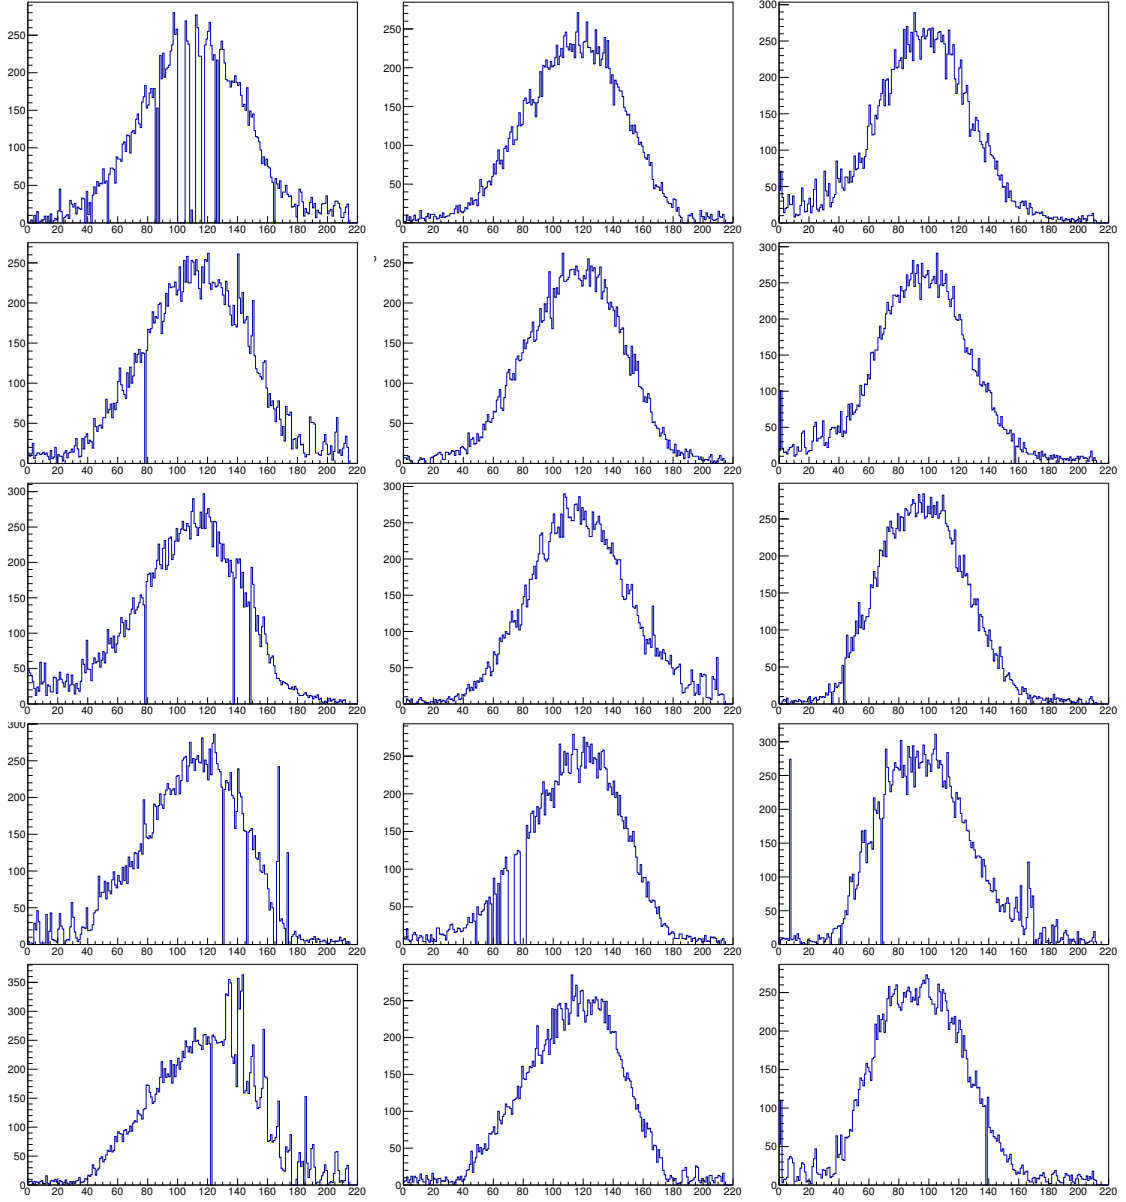


Figure B.1: Distributions of reconstructed digits from Dataset A for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the upstream tracker.




Figure B.2: Distributions of reconstructed digits from Dataset A for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the downstream tracker.

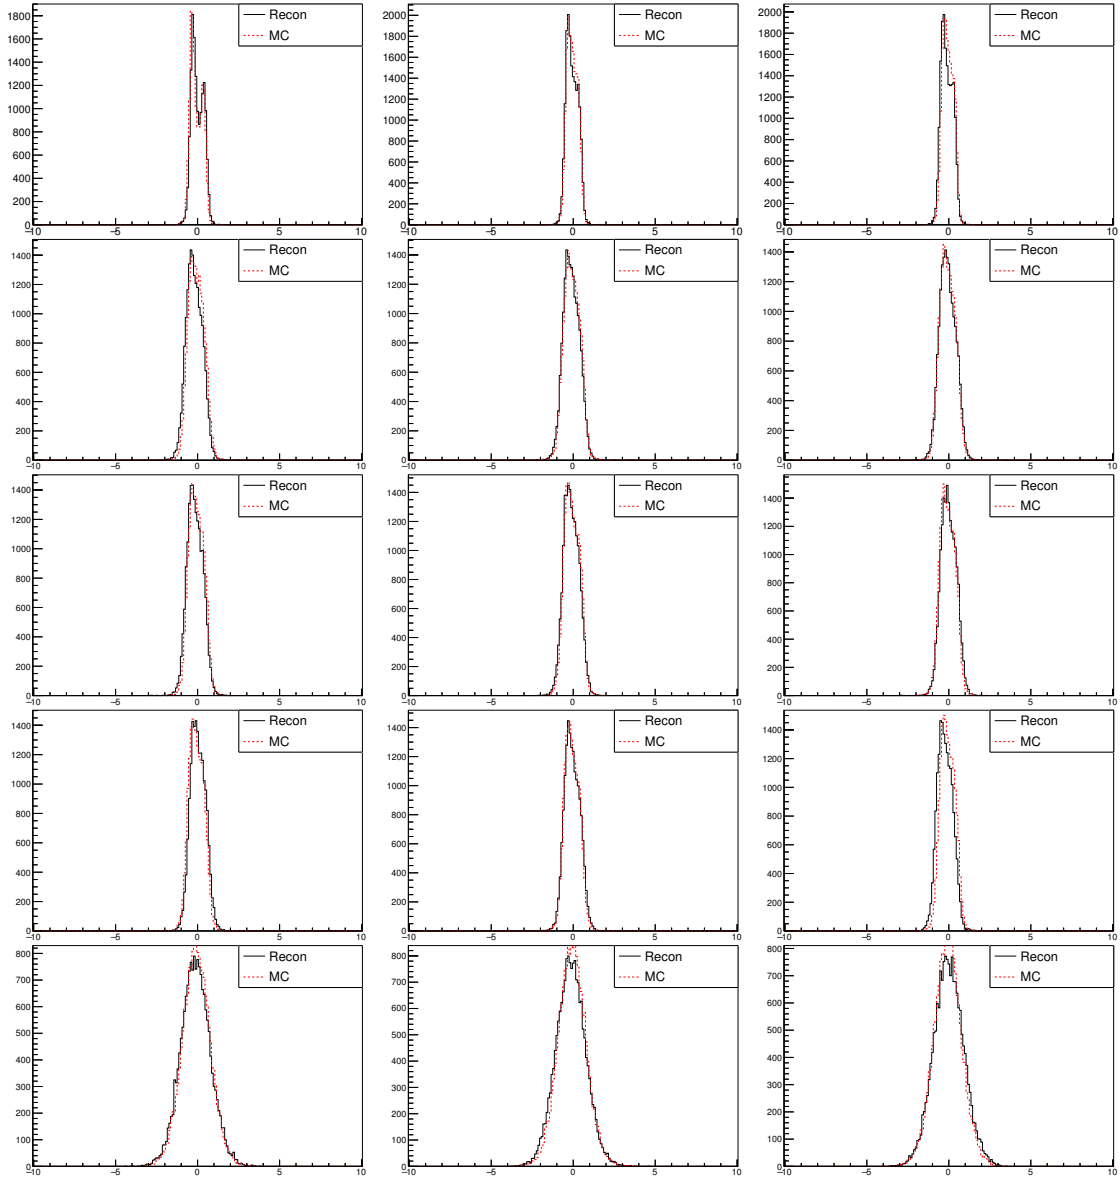


Figure B.3: Distributions of reconstructed pulls from Dataset A for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the upstream tracker.

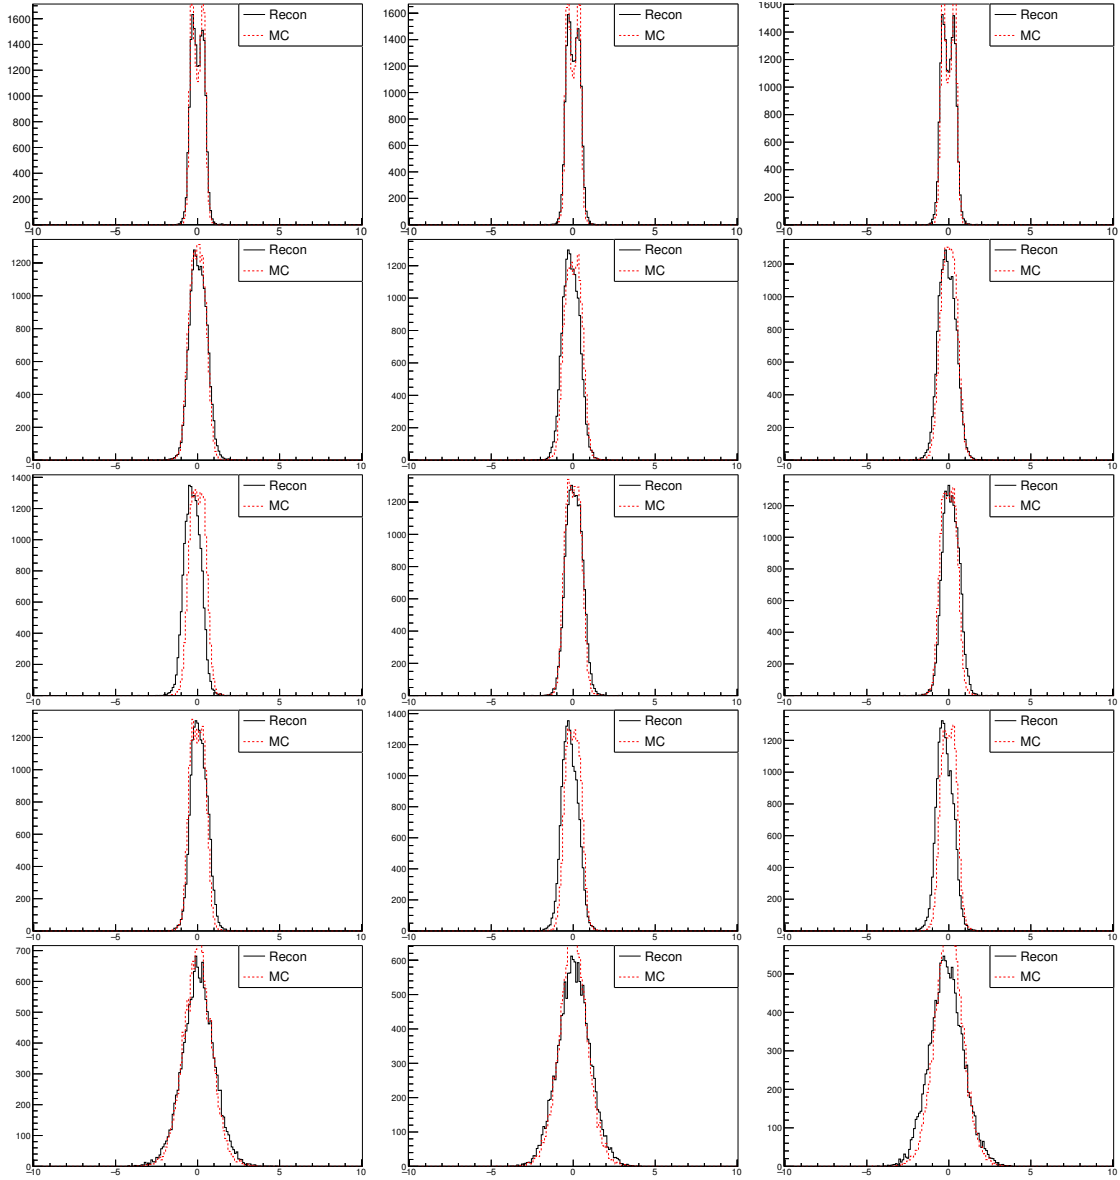


Figure B.4: Distributions of reconstructed pulls from Dataset A for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the downstream tracker.

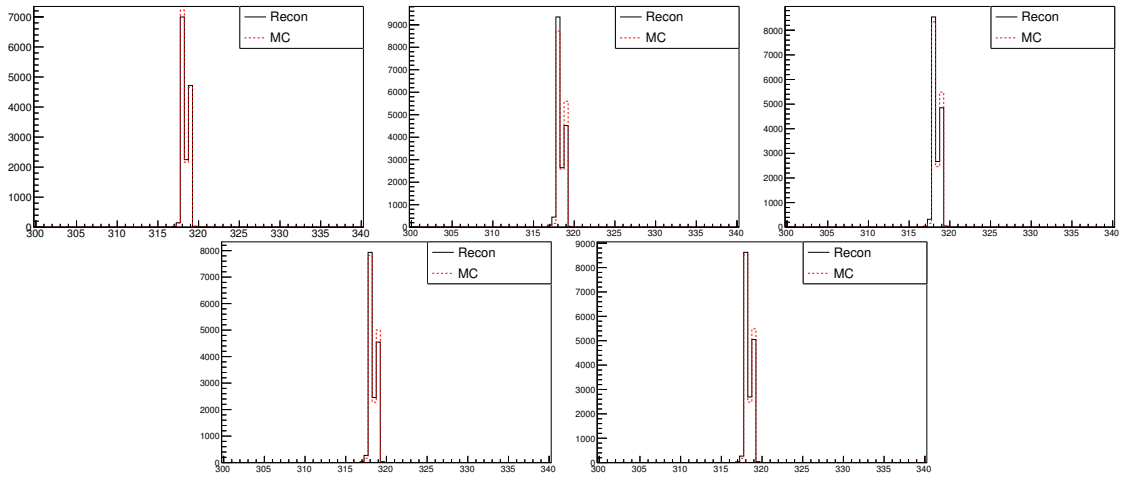


Figure B.5: Distributions of the Kuno-Sums for spacepoints used to form a track from Dataset A in stations 1 to 5 in the upstream tracker.

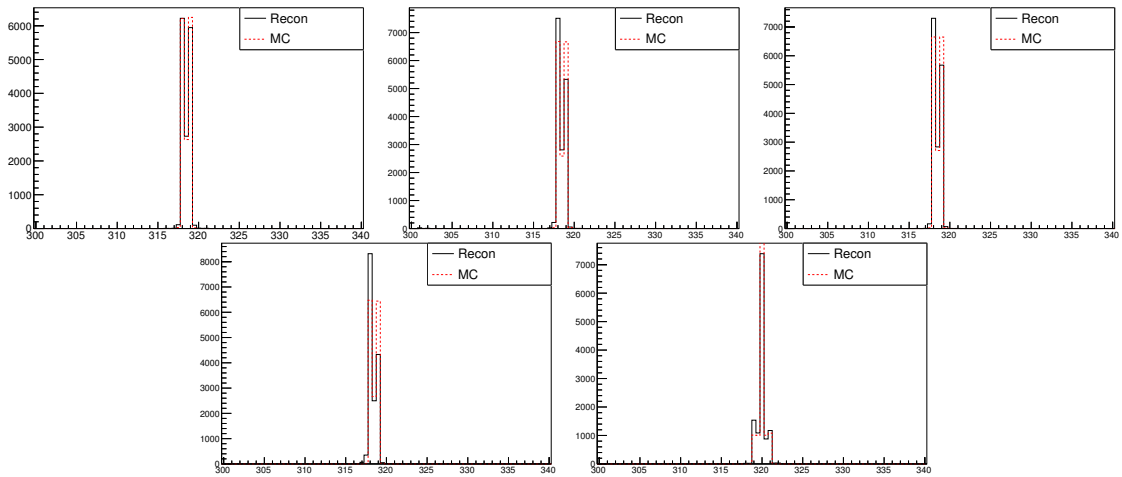


Figure B.6: Distributions of the Kuno-Sums for spacepoints used to form a track from Dataset A in stations 1 to 5 in the downstream tracker.

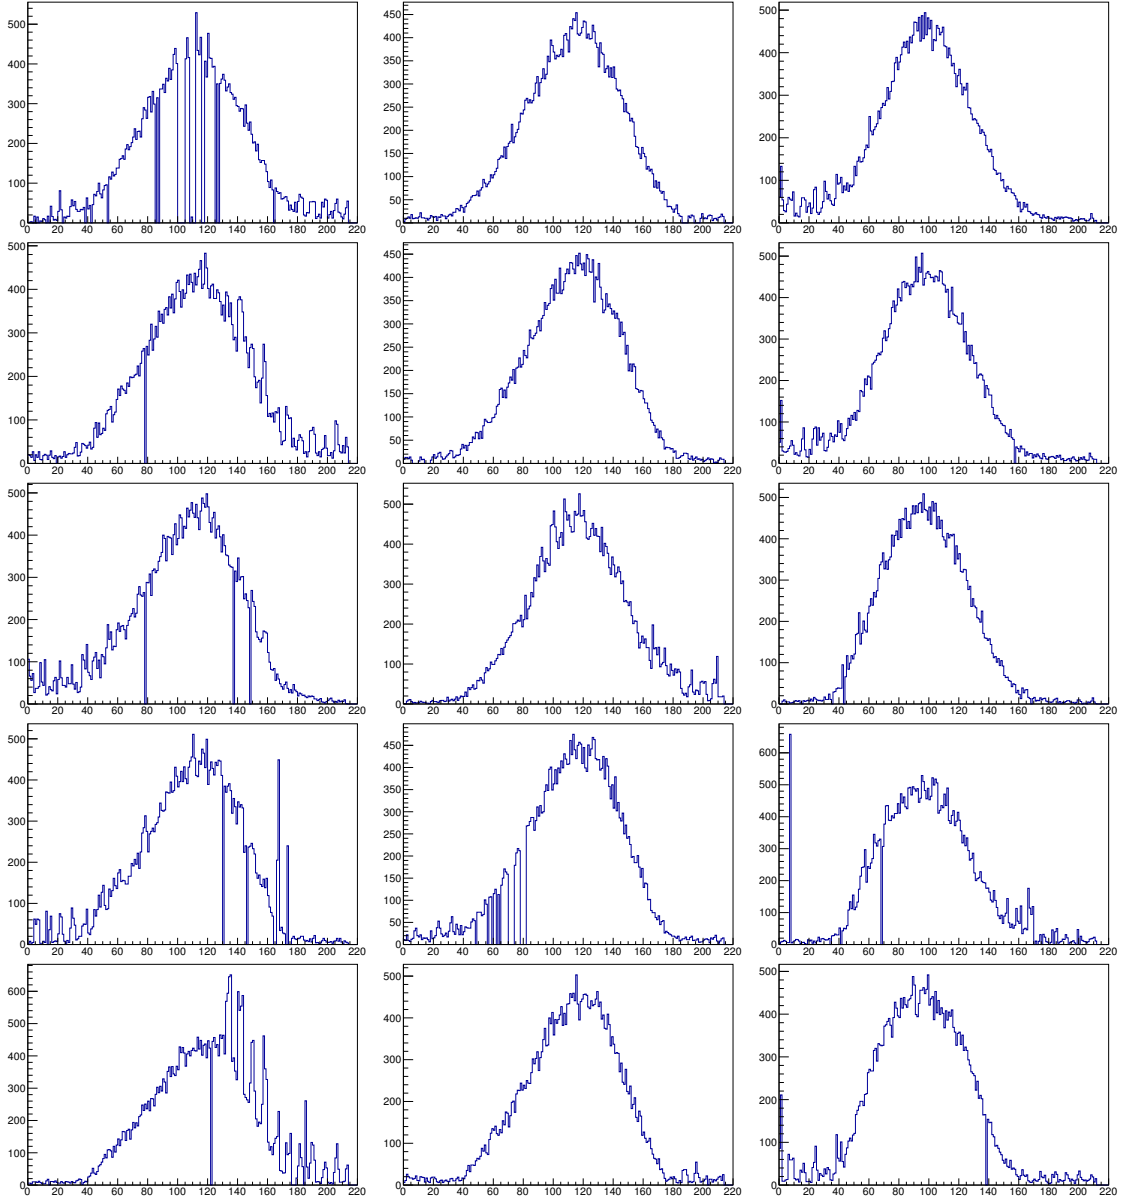


Figure B.7: Distributions of reconstructed digits from Dataset B for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the upstream tracker.




Figure B.8: Distributions of reconstructed digits from Dataset B for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the downstream tracker.

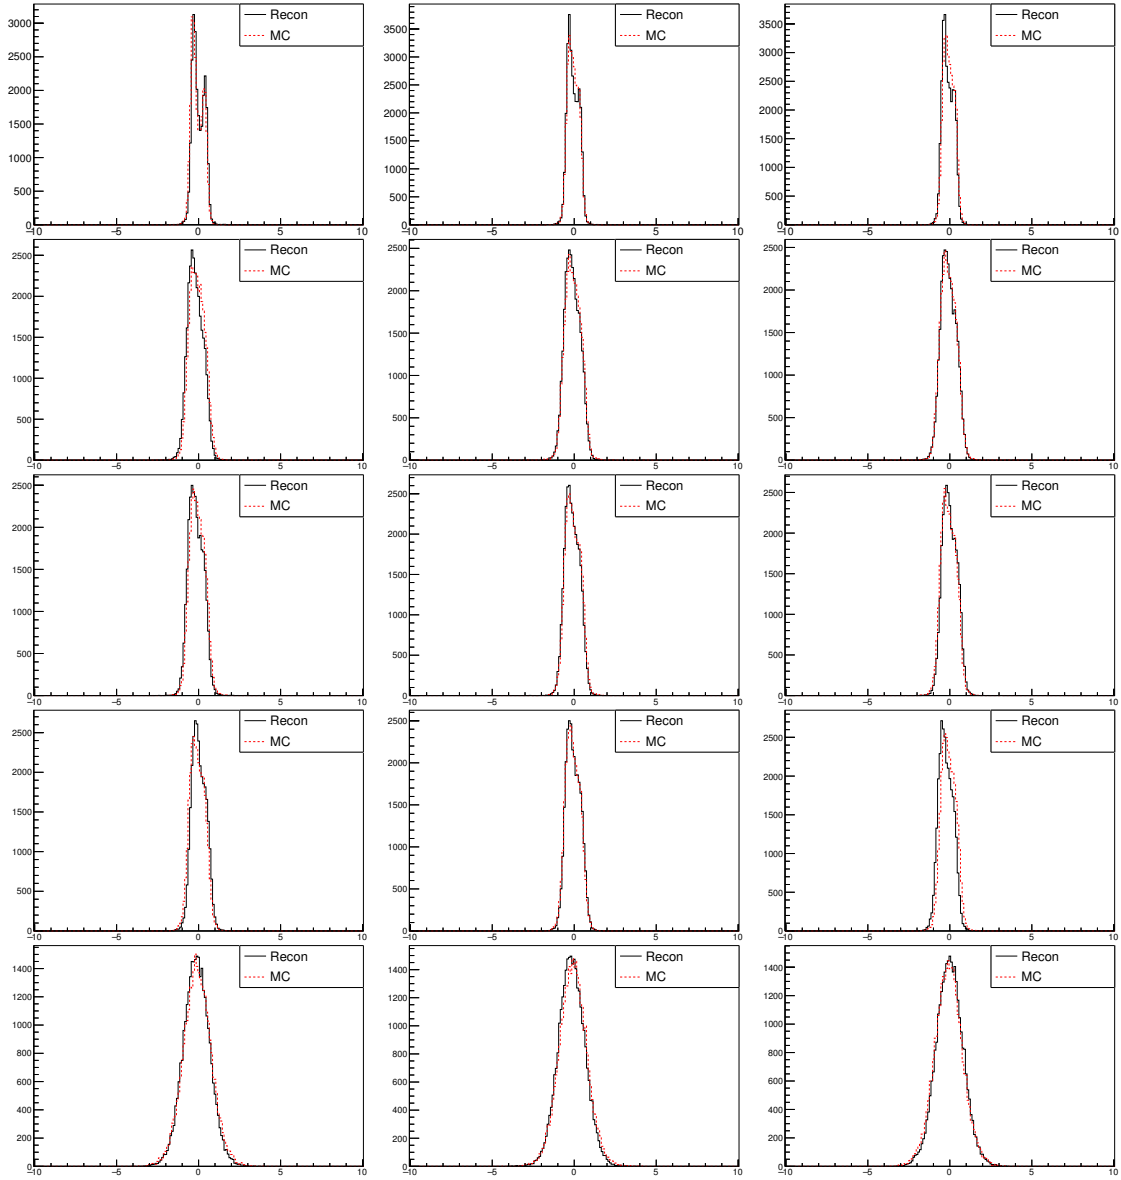


Figure B.9: Distributions of reconstructed pulls from Dataset B for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the upstream tracker.

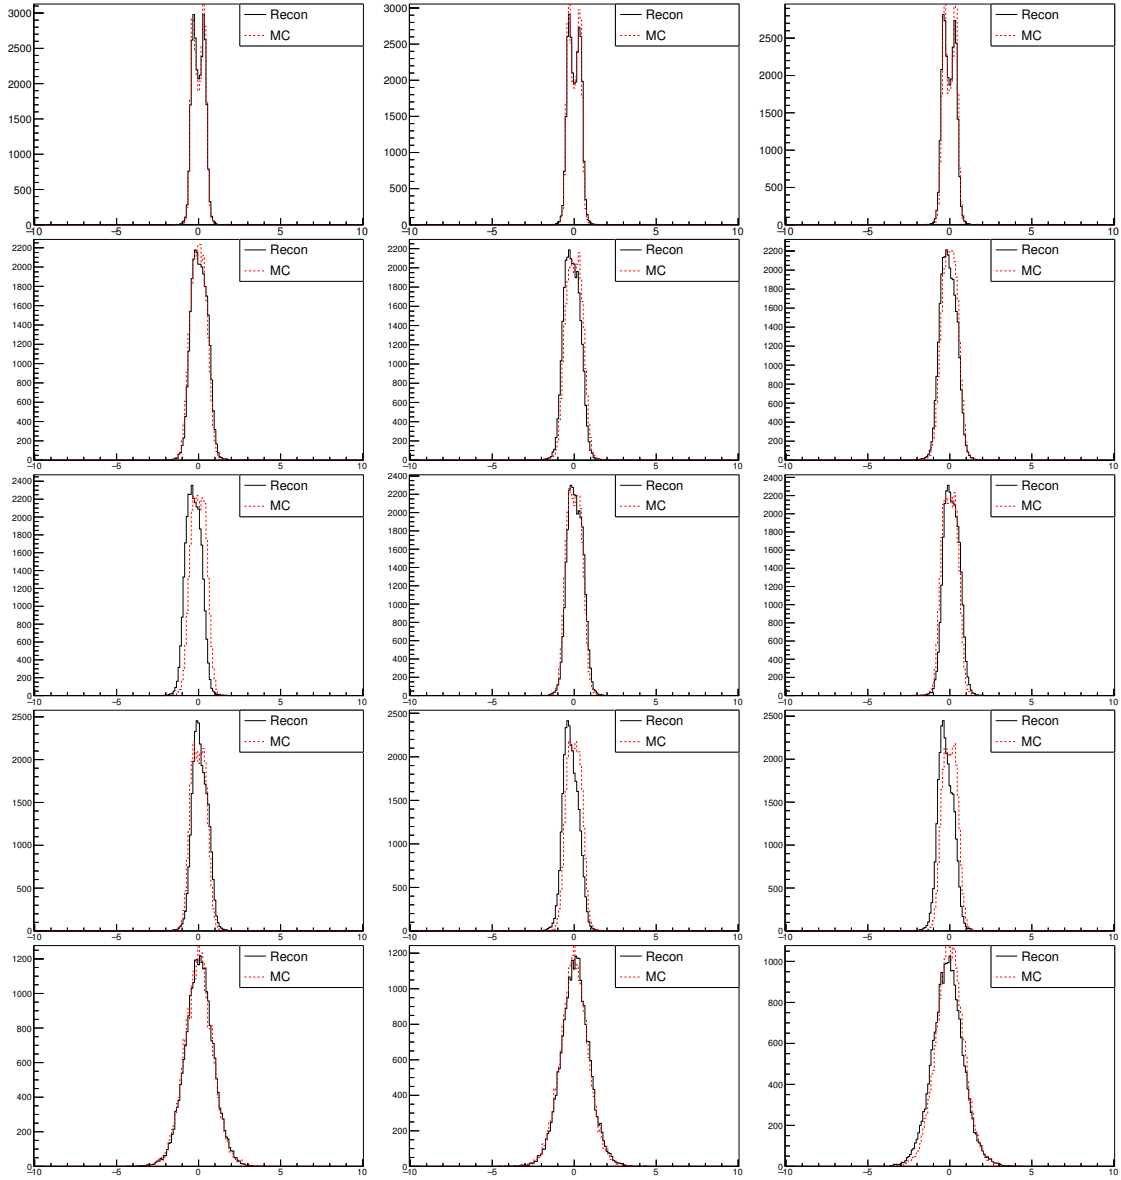


Figure B.10: Distributions of reconstructed pulls from Dataset B for all planes, 0 to 2, left to right, and all stations 1 to 5, top to bottom in the downstream tracker.

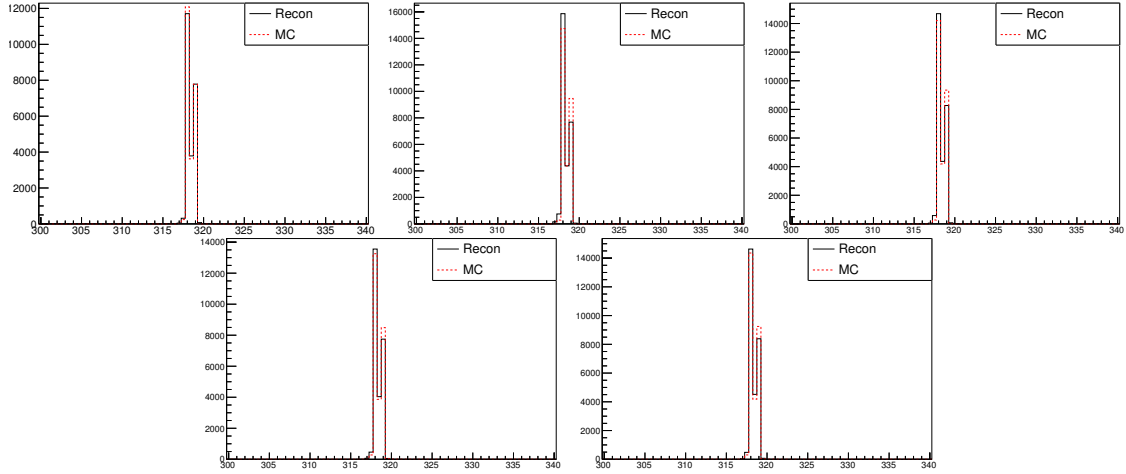


Figure B.11: Distributions of the Kuno-Sums for spacepoints used to form a track from Dataset B in stations 1 to 5 in the upstream tracker.

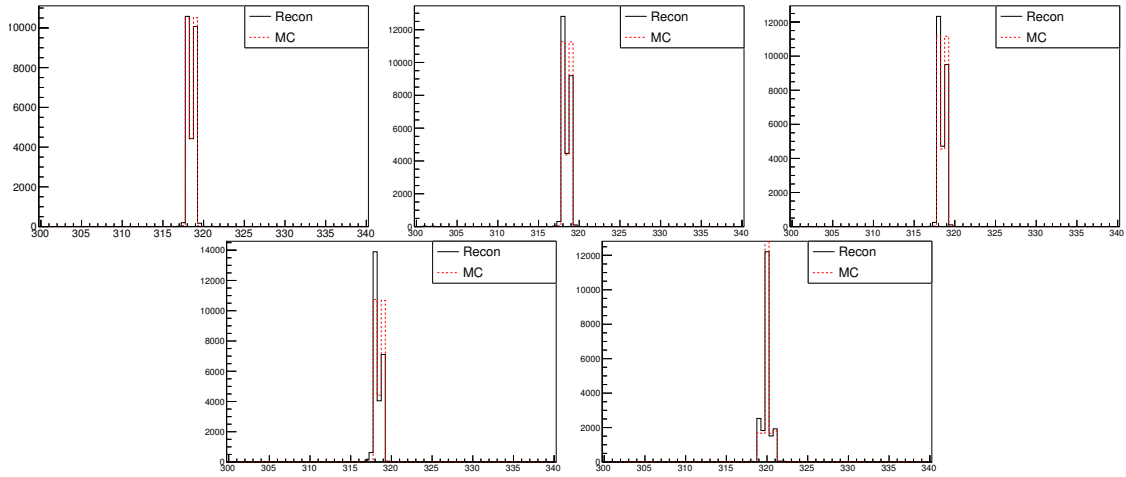


Figure B.12: Distributions of the Kuno-Sums for spacepoints used to form a track from Dataset B in stations 1 to 5 in the downstream tracker. Note that station 5 was constructed with an additional channel in plane 2, hence the peak is appropriately displaced.