

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/97733>

Copyright and reuse:

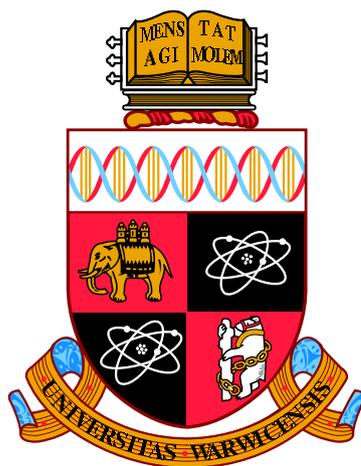
This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



Multiple Coulomb Scattering of Muons in MICE

by

Jan R. Greis

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

September 2017

THE UNIVERSITY OF
WARWICK

Contents

List of Tables	v
List of Figures	vi
Acknowledgments	xiii
Declarations	xiv
Abstract	xv
Chapter 1 Introduction	1
1.1 History of the Neutrino	1
1.1.1 Prediction & Discovery	1
1.1.2 The Solar Neutrino Problem & Neutrino Oscillation	2
1.2 Neutrino Oscillations & Mass	3
1.3 Current Experiments	6
1.3.1 Neutrino Oscillation Parameters	6
1.4 Next Generation Oscillation Experiments	8
1.4.1 Deep Underground Neutrino Experiment (DUNE)	8
1.4.2 Hyper-Kamiokande	8
1.4.3 Jiangmen Underground Neutrino Observatory (JUNO)	9
1.4.4 Precision IceCube Next Generation Upgrade (PINGU)	10
1.5 The Neutrino Factory	10
1.5.1 Proton Driver	12
1.5.2 Target	12
1.5.3 Muon Front-End	13
1.5.4 Acceleration	13
1.5.5 Storage / Decay Ring	14

1.6	Muon Colliders	15
Chapter 2 MICE		17
2.1	Beam Emittance	17
2.1.1	RMS Emittance	19
2.2	Ionisation Cooling	20
2.3	MICE Target & Beamline	22
2.3.1	The ISIS Proton Synchrotron	22
2.3.2	The Target	22
2.3.3	Beam Magnets	23
2.3.4	Proton Absorber	25
2.3.5	Diffuser	26
2.3.6	Partial Return Yoke	26
2.4	Cooling Channel	27
2.4.1	Step IV Configuration	27
2.4.2	Demonstration of Ionisation Cooling	28
2.5	Detectors	30
2.5.1	Luminosity Monitor	30
2.5.2	Time-of-Flight Detectors	30
2.5.3	Cherenkov Detectors	32
2.5.4	Scintillating Fibre Trackers	32
2.5.5	KLOE-Light Calorimeter	34
2.5.6	Electron Muon Ranger	35
2.6	Software Framework	36
2.6.1	Software Elements	36
2.6.2	Framework Design	37
2.6.3	Detector Reconstruction	38
2.7	Outlook	40
Chapter 3 Global Track Reconstruction		41
3.1	Global Data Structure	41
3.2	Track Matching	42
3.2.1	Propagation Method	43
3.2.2	TOF1 & TOF2	45
3.2.3	TOF0	45
3.2.4	Cherenkov Detectors	46

3.2.5	KL	46
3.2.6	EMR	46
3.2.7	Through the Absorber	46
3.2.8	Performance	47
3.3	Particle Identification (PID)	50
3.3.1	PID Variables	54
3.4	Track Fitting	54
3.5	Conclusion	55
Chapter 4 Multiple Scattering in the MICE Absorbers		56
4.1	Physics of Multiple Scattering	56
4.1.1	Single Scattering	56
4.1.2	Plural and Multiple Scattering	57
4.1.3	The Molière Theory	58
4.1.4	Other Scattering Theories	59
4.1.5	Algorithms	59
4.2	Previous Scattering Measurements	61
4.3	Scattering Measurements in MICE	63
4.4	Measurement Strategy	64
4.4.1	The Convolution Approach	66
4.4.2	Verification with Toy Monte Carlo	68
4.5	Selection Criteria	70
4.5.1	Number of Tracker Tracks and TOF0/1 Space Points	71
4.5.2	Tracker Track Fit χ^2/NDF	71
4.5.3	Number of Tracker Clusters	74
4.5.4	Minimum Transverse Momentum in the Upstream Tracker	74
4.5.5	No Energy Gain in Absorber	75
4.5.6	Fiducial Volume Cut	75
4.5.7	Momentum Selection	76
4.5.8	PID Selection	77
4.6	Systematic Errors	79
4.6.1	Tracker Track Fit χ^2/NDF Cut	80
4.6.2	Number of Tracker Clusters Cut	80
4.6.3	Minimum Transverse Momentum in the Upstream Tracker Cut	81
4.6.4	No Energy Gain in Absorber Cut	81
4.6.5	Fiducial Volume Cut	81

4.6.6	Momentum Selection	81
4.6.7	PID Selection	81
4.6.8	Amount of Absorber Material	82
4.6.9	Beamline Alignment	82
4.6.10	Alignment Correction	83
4.6.11	Magnetic Field Strength	86
4.7	Results	86
Chapter 5 Summary		96
Appendix A Covariance Matrices		98
A.1	Covariance Matrices for Individual Systematic Errors	99
A.2	Combined Covariance Matrices	147

List of Tables

1.1	Global best fit neutrino oscillation parameters	7
2.1	Summary of the MICE beamline magnets	25
3.1	Default track matching tolerances	47
4.1	Remaining events after cuts for runs with 140 MeV/c nominal beam momentum	72
4.2	Remaining events after cuts for runs with 170 MeV/c nominal beam momentum	72
4.3	Remaining events after cuts for runs with 200 MeV/c nominal beam momentum	73
4.4	Remaining events after cuts for runs with 240 MeV/c nominal beam momentum	73
4.5	χ^2/NDF values for all momentum selections	88

List of Figures

1.1	Illustration of normal and inverted neutrino mass hierarchy and flavour content of mass states.	6
1.2	DUNE beamline	8
1.3	Hyper-Kamiokande detector	9
1.4	Proposed detector design for Juno	10
1.5	Layout for a 25 GeV neutrino factory	11
1.6	Concept design for the liquid mercury jet target system for a neutrino factory	12
1.7	Racetrack design for the neutrino factory storage ring	14
1.8	Conceptual block diagrams for a neutrino factory and a muon collider . .	16
2.1	A phase-space ellipse described by the parameters α_t , β_t , and γ_t	19
2.2	Illustration of the mechanism of ionisation cooling	21
2.3	The MICE Step IV Beamline	22
2.4	Schematic of the MICE target assembly	23
2.5	Model of the MICE Diffuser	26
2.6	Design of the MICE partial return yoke	27
2.7	The MICE absorbers	28
2.8	MICE steps IV, V, and VI as originally planned	28
2.9	“Demonstration of Ionisation Cooling” configuration of the MICE beamline	29
2.10	Descoped “Demonstration of Ionisation Cooling” configuration of the MICE beamline	29
2.11	MICE Luminosity Monitor	31
2.12	Schematic of the time-of-flight detectors	31
2.13	Cherenkov threshold curves for electrons, muons, and pions in the MICE Cherenkov detectors	32

2.14	Schematics of a MICE tracker frame and the arrangement of tracker planes and coils	33
2.15	Arrangement of fibres and layers in a tracker plane	34
2.16	Schematic layout of scintillating fibres and extruded lead in the KL	34
2.17	Diagrams of the active volume and arrangement of triangular bars in the EMR	35
3.1	Illustration of the difference between true path length and distance between start end end point of a step for curved tracks	45
3.2	Illustration of the track matching process	47
3.3	Position residuals for MC truth propagation in the downstream tracker	49
3.4	Transverse momentum residuals for MC truth propagation in the downstream tracker	49
3.5	Track matching residuals for TOF0	51
3.6	Track matching residuals for TOF1	51
3.7	Track matching residuals for TOF2	52
3.8	Track matching residuals for the KL	52
3.9	Track matching residuals for the EMR	53
3.10	Track matching residuals through the absorber	53
4.1	Illustration of the bulk effect of multiple single scatters in a material	59
4.2	The MuScat detector	61
4.3	MuScat scattering distribution for 0.24 mm of iron	62
4.4	MuScat scattering distribution for 159 mm of liquid hydrogen	63
4.5	Normalised distributions of the maximum radial displacement from the beamline centre in the upstream tracker	64
4.6	An illustration of the propagation required to calculate scattering angles in MICE in the presence of a magnetic field.	65
4.7	Toy MC scattering distributions without absorber, absorber-only, and the convolution of the two.	69
4.8	Comparison of scattering distributions in the toy Monte Carlo	70
4.9	χ^2/NDF distribution of the tracker track fit before and after the cut.	71
4.10	Number of Tracker Clusters before and after the cut.	74
4.11	Uncertainty in p_z as a function of p_t	75
4.12	Reconstructed momentum in the downstream tracker as a function of reconstructed momentum in the upstream tracker	76

4.13	Maximum transverse position of the path inside the tracker before and after the cut	77
4.14	Momentum distribution in the upstream tracker before and after momentum selection	78
4.15	TOF0–TOF1 time-of-flight distribution before and after PID selection . .	78
4.16	Illustration of the sampling smoothing used for evaluation of systematic uncertainties	80
4.17	Sample covariance matrix for the momentum selection systematic for LiH data at 170 MeV/c	82
4.18	Sample covariance matrices for alignment and rotation systematics for LiH data at 170 MeV/c	83
4.19	Scattering distributions for 180, 200, and 220 MeV/c before shift correction	83
4.20	MC-based momentum-dependent shift in reconstructed scattering angle as result of a 5 mm spatial misalignment of the downstream tracker . . .	84
4.21	Shift correction graph for the dataset with a nominal beamline momentum of 200 MeV/c	84
4.22	Sample covariance matrix for the alignment correction systematic for LiH data at 170 MeV/c	85
4.23	Sample covariance matrices for magnetic field uncertainty systematics for LiH data at 170 MeV/c	87
4.24	RMS widths of the scattering angle distributions for LiH data and convolution between “empty” data and absorber-only MC scattering distribution	89
4.25	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 120 MeV/c	90
4.26	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 130 MeV/c	90
4.27	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 140 MeV/c	91
4.28	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 150 MeV/c	91
4.29	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 160 MeV/c	92
4.30	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 170 MeV/c	92
4.31	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 180 MeV/c	93

4.32	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 190 MeV/c	93
4.33	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 200 MeV/c	94
4.34	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 220 MeV/c	94
4.35	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 230 MeV/c	95
4.36	Comparison between LiH data and convolution between “empty” data and absorber-only MC scattering distribution at 240 MeV/c	95
A.1	Statistical and alignment covariance matrices for the 120 MeV/c LiH dataset	99
A.2	Momentum selection and field uncertainty covariance matrices for the 120 MeV/c LiH dataset	100
A.3	Statistical and alignment covariance matrices for the 120 MeV/c convoluted dataset	101
A.4	Momentum selection and field uncertainty covariance matrices for the 120 MeV/c convoluted dataset	102
A.5	Statistical and alignment covariance matrices for the 130 MeV/c LiH dataset	103
A.6	Momentum selection and field uncertainty covariance matrices for the 130 MeV/c LiH dataset	104
A.7	Statistical and alignment covariance matrices for the 130 MeV/c convoluted dataset	105
A.8	Momentum selection and field uncertainty covariance matrices for the 130 MeV/c convoluted dataset	106
A.9	Statistical and alignment covariance matrices for the 140 MeV/c LiH dataset	107
A.10	Momentum selection and field uncertainty covariance matrices for the 140 MeV/c LiH dataset	108
A.11	Statistical and alignment covariance matrices for the 140 MeV/c convoluted dataset	109
A.12	Momentum selection and field uncertainty covariance matrices for the 140 MeV/c convoluted dataset	110
A.13	Statistical and alignment covariance matrices for the 150 MeV/c LiH dataset	111
A.14	Momentum selection and field uncertainty covariance matrices for the 150 MeV/c LiH dataset	112

A.15 Statistical and alignment covariance matrices for the 150 MeV/c convoluted dataset	113
A.16 Momentum selection and field uncertainty covariance matrices for the 150 MeV/c convoluted dataset	114
A.17 Statistical and alignment covariance matrices for the 160 MeV/c LiH dataset	115
A.18 Momentum selection and field uncertainty covariance matrices for the 160 MeV/c LiH dataset	116
A.19 Statistical and alignment covariance matrices for the 160 MeV/c convoluted dataset	117
A.20 Momentum selection and field uncertainty covariance matrices for the 160 MeV/c convoluted dataset	118
A.21 Statistical and alignment covariance matrices for the 170 MeV/c LiH dataset	119
A.22 Momentum selection and field uncertainty covariance matrices for the 170 MeV/c LiH dataset	120
A.23 Statistical and alignment covariance matrices for the 170 MeV/c convoluted dataset	121
A.24 Momentum selection and field uncertainty covariance matrices for the 170 MeV/c convoluted dataset	122
A.25 Statistical and alignment covariance matrices for the 180 MeV/c LiH dataset	123
A.26 Momentum selection and field uncertainty covariance matrices for the 180 MeV/c LiH dataset	124
A.27 Statistical and alignment covariance matrices for the 180 MeV/c convoluted dataset	125
A.28 Momentum selection and field uncertainty covariance matrices for the 180 MeV/c convoluted dataset	126
A.29 Statistical and alignment covariance matrices for the 190 MeV/c LiH dataset	127
A.30 Momentum selection and field uncertainty covariance matrices for the 190 MeV/c LiH dataset	128
A.31 Statistical and alignment covariance matrices for the 190 MeV/c convoluted dataset	129
A.32 Momentum selection and field uncertainty covariance matrices for the 190 MeV/c convoluted dataset	130
A.33 Statistical and alignment covariance matrices for the 200 MeV/c LiH dataset	131
A.34 Momentum selection and field uncertainty covariance matrices for the 200 MeV/c LiH dataset	132

A.35 Statistical and alignment covariance matrices for the 200 MeV/c convoluted dataset	133
A.36 Momentum selection and field uncertainty covariance matrices for the 200 MeV/c convoluted dataset	134
A.37 Statistical and alignment covariance matrices for the 220 MeV/c LiH dataset	135
A.38 Momentum selection and field uncertainty covariance matrices for the 220 MeV/c LiH dataset	136
A.39 Statistical and alignment covariance matrices for the 220 MeV/c convoluted dataset	137
A.40 Momentum selection and field uncertainty covariance matrices for the 220 MeV/c convoluted dataset	138
A.41 Statistical and alignment covariance matrices for the 230 MeV/c LiH dataset	139
A.42 Momentum selection and field uncertainty covariance matrices for the 230 MeV/c LiH dataset	140
A.43 Statistical and alignment covariance matrices for the 230 MeV/c convoluted dataset	141
A.44 Momentum selection and field uncertainty covariance matrices for the 230 MeV/c convoluted dataset	142
A.45 Statistical and alignment covariance matrices for the 240 MeV/c LiH dataset	143
A.46 Momentum selection and field uncertainty covariance matrices for the 240 MeV/c LiH dataset	144
A.47 Statistical and alignment covariance matrices for the 240 MeV/c convoluted dataset	145
A.48 Momentum selection and field uncertainty covariance matrices for the 240 MeV/c convoluted dataset	146
A.49 Combined (incl. statistical errors) covariance matrices 120 MeV/c	147
A.50 Combined (incl. statistical errors) covariance matrices 130 MeV/c	147
A.51 Combined (incl. statistical errors) covariance matrices 140 MeV/c	148
A.52 Combined (incl. statistical errors) covariance matrices 150 MeV/c	148
A.53 Combined (incl. statistical errors) covariance matrices 160 MeV/c	149
A.54 Combined (incl. statistical errors) covariance matrices 170 MeV/c	149
A.55 Combined (incl. statistical errors) covariance matrices 180 MeV/c	150
A.56 Combined (incl. statistical errors) covariance matrices 190 MeV/c	150
A.57 Combined (incl. statistical errors) covariance matrices 200 MeV/c	151
A.58 Combined (incl. statistical errors) covariance matrices 220 MeV/c	151
A.59 Combined (incl. statistical errors) covariance matrices 230 MeV/c	152

A.60 Combined (incl. statistical errors) covariance matrices 240 MeV/c 152

Acknowledgments

Firstly, I would like to thank my supervisor, Steve Boyd. I genuinely think that having the most cynical person I know as a supervisor was quite helpful in getting through those four years. Steve, you still owe me that scythe.

Those who have come and gone in office 450, I have certainly enjoyed the 1/2/3/4 (please circle appropriate choice) years I have known you all.

There are a great many excellent people on MICE deserving of thanks, but in particular I want to mention Adam Dobbs, Durga Rajaram, Chris Rogers, and Melissa Uchida, for all their support over the years. Also, John Cobb, for his help with understanding multiple scattering theories.

At the intersection of the two groups above is Celeste Pidcott, whose whimsical sense of humour certainly brightened up many days. Also, she foolishly agreed to proof-read this thesis.

I want to thank Friedo Zölzer, whose most excellent teaching played a rather significant role in my decision to study Physics.

I would also like to thank my family, who have all supported me in many different ways. Especially my brother Kian, without whom my path in life would have been quite different.

Finally, I want to thank my wife, Steph. Going through the process of doing a PhD together and being able to support each other certainly helps with maintaining some semblance of sanity.

Declarations

Chapter 1 is drawn entirely from available literature and publications. Chapter 2 consists of information from MICE publications (published as well as in preparation), internal notes, as well as some literature. Chapter 3 is largely the work of the author with some information drawn from internal documentation, a recent MICE thesis, and available literature. Sections 4.1 and 4.2 use information from publications as well as (in the case of 4.1) the GEANT4 documentation. From section 4.3 onwards, all work presented in this thesis is the sole work of the author. This thesis has not been submitted for a degree at another university.

Abstract

The Muon Ionisation Cooling Experiment (MICE) was designed to provide the first demonstration of muon ionisation cooling by passing a muon beam through a cooling channel consisting of low- Z absorber material as well as radio-frequency (RF) cavities, with the net effect of a reduction in transverse momentum spread at constant total momentum. The performance that can be achieved in such a cooling channel depends on both energy loss and multiple coulomb scattering in the absorber, the latter of which is currently not understood to very high precision for muons. Verification of muon multiple scattering algorithms – especially in low- Z materials – is therefore crucial for accurate modelling of a muon ionisation cooling channel, which would be used in future facilities such as a Neutrino Factory or a Muon Collider.

A track matching algorithm was implemented in order to combine locally reconstructed detector output for further processing and analysis. Multiple coulomb scattering of muons in the lithium hydride (LiH) MICE absorber has been studied in the presence of magnetic fields using a convolution method, which has the advantage over typically used deconvolution methods in that it does not depend on the chosen value of the regularisation strength parameter. It does, however, only allow comparison between data and Monte Carlo, rather than direct extraction of the “true” scattering distributions. The presence of magnetic fields – while providing significantly superior momentum resolution – introduces additional systematic errors, mainly due to the uncertainties the field maps used.

The results of this analysis provide support for the use of convolution methods and suggest that scattering algorithms in GEANT4 have significantly improved since the time of MuScat (which was an experiment designed specifically to measure muon scattering in a variety of materials). There is some indication, that they may have changed from vastly overestimating to slightly underestimating large angle scatters though the uncertainties on the final results are insufficiently small to make definite statements on this point.

Chapter 1

Introduction

“Physics is really nothing more than a search for ultimate simplicity, but so far all we have is a kind of elegant messiness.”

Bill Bryson

1.1 History of the Neutrino

1.1.1 Prediction & Discovery

The neutrino was first postulated by Wolfgang Pauli in a letter to colleagues at a physics conference in Tübingen dated 4 December 1930[1]. In it, he proposed a new fundamental particle – which at the time he dubbed the neutron – in order to explain the continuous energy spectrum of β decay. He himself referred to the postulation as a desperate resort to save the principle of conservation of energy as well as what is now known as the Pauli Exclusion Principle and stated that at the time he did not dare to publish his idea before receiving feedback about the feasibility of proving the existence of such a particle.

Further study of the energy spectrum of β decay in the following years – specifically evidence for a strict cut-off of the electron energy – provided growing support for the existence of a neutrino, but it took another 26 years to achieve detection of neutrinos. In the Cowan-Reines neutrino experiment[2], antineutrinos from a nuclear reactor interacted with protons in water tanks via the interaction



with the resulting positron annihilating with an electron producing two gamma rays. Additionally, the neutron was captured by cadmium in the water, producing a third

gamma ray. By timing the coincidence of the 3 gamma rays, it was possible to conclusively determine the presence of neutrinos, a result that won Frederick Reines the Nobel Prize in physics in 1995.

In 1962, Lederman, Schwartz, and Steinberger demonstrated the existence of the muon neutrino by showing that the interaction of neutrinos from muon decays with nucleons produced muons, but never electrons. This discovery earned them the Nobel Prize in 1988.[3]

The tau neutrino finally was discovered in 2000 by the DONUT (Direct Observation of the NU Tau) experiment at Fermilab.[4] ν_τ were primarily produced from $D_S \rightarrow \tau + \bar{\nu}_\tau$ as well as the subsequent τ decays, and observed by the characteristic kink from the decay of the τ produced in the interaction of the $\nu_\tau^{(-)}$ with the detector.

This discovery completed the set of 6 Standard Model leptons, as experimental evidence strongly suggests the existence of exactly three neutrino flavours with $m_\nu < m_Z/2$ (the combined fit from the experiments ALEPH, DELPHI, L3, and OPAL lies at 2.984 ± 0.008)[5]. It is however still possible that additional types of neutrinos exist that are either sterile (i.e. don't participate in weak interactions), very heavy, or both.

1.1.2 The Solar Neutrino Problem & Neutrino Oscillation

The Homestake Experiment[6] in the late 1960 was one of the first experiments designed to count the solar neutrinos produced in fusion reactions in the Sun's interior. Located 4,400 m underground in the Homestake gold mine at Lead, South Dakota it used a detector containing 390,000 litres of liquid tetrachloroethylene (C_2Cl_4) to detect neutrinos from the capture reaction



The argon produced was extracted by bubbling the tank with helium and then separated out via gas chromatography.

The experiment found a neutrino count significantly smaller than what was expected from experimental models, a result that was confirmed by experiments such as Kamiokande-II[7], SAGE[8], GALLEX[9], Super-Kamiokande[10], and SNO[11], which consistently measured rates between half and a third of what would have been expected from the Standard Solar Model (SSM), a phenomenon that was named the "Solar Neutrino Problem".

Additionally, studies of muon neutrinos from cosmic rays also found a deficit, which the Super-Kamiokande experiment demonstrated to be a function of the zenith-angle[12], i.e. the ratio of the number of detections from cosmic ray events on the far

and the near side of the planet was significantly lower than expected.

All these results provided more and more support for neutrino oscillations – neutrinos changing their flavour as they propagate. The most crucial piece of evidence then came from SNO in 2002[13]. Using a heavy water Cherenkov detector, the experiment was able to detect three neutrino flavours via three different interactions, charged current (CC), neutral current (NC), and elastic scattering (ES):

$$\nu_e + d \rightarrow p + p + e^- \quad (1.3)$$

$$\nu_x + d \rightarrow p + n + \nu_x \quad (1.4)$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \quad (1.5)$$

where $x = e, \mu, \tau$.

While the CC interaction probed by SNO only occurs for ν_e , the NC interaction has equal sensitivity for all flavours, whereas the ES interaction has reduced sensitivity to ν_μ and ν_τ with respect to ν_e . By comparing the relative detection rates for the three channels, SNO demonstrated the existence of non- ν_e solar neutrinos at a confidence level of 5.3σ , with the total number of detected neutrinos consistent with predictions from the SSM[13]. As there is no known process by which the Sun would produce ν_μ or ν_τ , this provided proof for neutrino oscillation and solved the solar neutrino problem.

1.2 Neutrino Oscillations & Mass

The discovery of neutrino oscillations carried with it a significant implication for the nature of neutrinos – that they must have non-zero mass. This is because oscillations require that a neutrino propagates as a mass state that is not identical to its flavour state. Thereby the flavour of the neutrino detected can differ from the flavour at creation of the neutrino. In the following description, which follows [14], we will first calculate the transition amplitude for the n flavour case and then proceed to the transition probability in the case of $n = 3$.

For an arbitrary number n of orthonormal eigenstates we can relate the n flavour eigenstates $|\nu_\alpha\rangle$ to the n mass eigenstates $|\nu_i\rangle$ via a single unitary mixing matrix U :

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle \quad (1.6)$$

conversely

$$|\nu_i\rangle = \sum_{\alpha} (U^\dagger)_{i\alpha} |\nu_\alpha\rangle = \sum_{\alpha} U_{\alpha i}^* |\nu_\alpha\rangle \quad (1.7)$$

For antineutrinos:

$$|\bar{\nu}_\alpha\rangle = \sum_i U_{\alpha i}^* |\bar{\nu}_i\rangle \quad (1.8)$$

The mass eigenstates then propagate with time and position dependence as

$$|\nu_i(x, t)\rangle = e^{-ix_\mu p^\mu} |\nu_i(0, 0)\rangle \quad (1.9)$$

where

$$x_\mu = (t, \mathbf{x}), \quad p^\mu = (E, \mathbf{p}) \quad (1.10)$$

As neutrinos are created and detected as pure flavour eigenstates, we can write the state of a neutrino created as $|\nu_\alpha\rangle$ at $t = 0$ using equations 1.6, 1.7, 1.9, and 1.10 as

$$\begin{aligned} |\nu(x, t)\rangle &= \sum_i U_{\alpha i} |\nu_i(x, t)\rangle \\ &= \sum_i U_{\alpha i} e^{-iE_i t} |\nu_i(x, 0)\rangle \\ &= \sum_i U_{\alpha i} e^{-iE_i t} e^{ipx} |\nu_i(0, 0)\rangle \\ &= \sum_{i, \beta} U_{\alpha i} U_{\beta i}^* e^{i(px - E_i t)} |\nu_\beta\rangle \end{aligned} \quad (1.11)$$

As can be seen from the phase factor $e^{i(px - E_i t)}$, the flavour content of the neutrino state changes as it propagates, as long as the masses of the $|\nu_i\rangle$ states differ. Hence, as mentioned above, oscillation implies that neutrinos have non-zero mass or more specifically, all mass states must have different masses i.e. at least two of the three mass states have non-zero mass (see also equation 1.14).

We can thus write the transition amplitude between flavour states $\nu_\alpha \nu_\beta$ as

$$\begin{aligned} A(\alpha \rightarrow \beta)(t) &= \langle \nu_\beta | \nu(x, t) \rangle = \sum_i U_{\beta i}^* U_{\alpha i} e^{i(px - E_i t)} \\ &= \sum_i U_{\beta i}^* U_{\alpha i} \exp\left(-i \frac{m_i^2 L}{2E}\right) \end{aligned} \quad (1.12)$$

where for the second line we have used the fact that neutrinos are ultra-relativistic and

hence

$$E_i = \sqrt{m_i^2 + p_i^2} \simeq p_i + \frac{m_i^2}{2p_i} \simeq E + \frac{m_i^2}{2E} \quad (1.13)$$

and $L = x = ct$ is the distance between neutrino source and detector.

The transition probability can then be calculated from $|A(\alpha \rightarrow \beta)|^2$ which for the three-flavour case is given by

$$P(\nu_\alpha \rightarrow \nu_\beta) = \delta_{\alpha\beta} - 4 \sum_{i>j=1}^3 \Re(K_{\alpha\beta,ij}) \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right) + 4 \sum_{i>j=1}^3 \Im(K_{\alpha\beta,ij}) \sin\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \cos\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \quad (1.14)$$

where $K_{\alpha\beta,ij} = U_{\alpha i} U^*_{\beta i} U_{\alpha j} U^*_{\beta j}$ and $\Delta m_{ij}^2 = m_i^2 - m_j^2$.

The matrix U is the Pontecorvo–Maki–Nakagawa–Sakat matrix U_{PMNS} , with a common parameterisation (for Dirac neutrinos; the case of Majorana neutrinos requires two additional complex phases) given by:

$$U_{\text{PMNS}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13} e^{i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13} e^{i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.15)$$

$$= \begin{pmatrix} c_{12} c_{13} & s_{12} c_{13} & s_{13} e^{-i\delta_{\text{CP}}} \\ -s_{12} c_{23} - c_{12} s_{23} s_{13} e^{i\delta_{\text{CP}}} & c_{12} c_{23} - s_{12} s_{23} s_{13} e^{i\delta_{\text{CP}}} & s_{23} c_{13} \\ s_{12} s_{23} - c_{12} s_{23} s_{13} e^{i\delta_{\text{CP}}} & -c_{12} s_{23} - s_{12} c_{23} s_{13} e^{i\delta_{\text{CP}}} & c_{23} c_{13} \end{pmatrix}$$

where the shorthand

$$s_{ij} = \sin(\theta_{ij}) \quad c_{ij} = \cos(\theta_{ij}) \quad (1.16)$$

has been used.

This matrix denotes the relationship between neutrino flavour and mass states as

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{\text{PMNS}} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (1.17)$$

U_{PMNS} is defined by a total of four parameters, which in the above representation are the mixing angles θ_{12} , θ_{13} , and θ_{23} , as well as the CP-violating phase δ_{CP} which as

of yet has not conclusively been shown to be non-zero, though a best fit on current data slightly favours $\sin(\delta_{\text{CP}}) \simeq -0.9$ and disfavours $\sin(\delta_{\text{CP}}) \simeq +0.9$ at almost 3σ [15].

Due to the way the parameters appear in the oscillation equations, the sign of the mass difference between two given neutrino mass states (Δm) is lost in the measurement. As a consequence of this, the so-called mass hierarchy of neutrinos is currently unknown, it could either be *normal* with small m_1 and m_2 and large m_3 or *inverted* with small m_3 and large m_1 and m_2 (see figure 1.1). Determination of the mass hierarchy is possible due to the way it affects the relative oscillation rates of $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ while passing through matter (which contains a vast excess of electrons as opposed to positrons) and the normal hierarchy is favoured by recent results at $\sim 2.0\sigma$ [16].

1.3 Current Experiments

In the last two decades, the field of neutrino physics has grown significantly, both in number and scope of active experiments. The main focus of the field is currently in the following areas: Precise measurement of the neutrino mixing angles, mass squared differences Δm_{21}^2 and Δm_{31}^2 , and the CP violating phase δ_{CP} (supported by neutrino cross section measurements), determination of the absolute neutrino mass scale and the neutrino mass hierarchy (figure 1.1), and the question whether the neutrino is a Dirac or Majorana particle (i.e. whether it is its own antiparticle).

1.3.1 Neutrino Oscillation Parameters

Table 1.1 shows current global best fit neutrino oscillation parameters. In order to probe the full parameter space, it is necessary to combine results from various experiments, as no current experiment has sufficient sensitivity in all regions of interest. The following description of the sources for the data on which the fit is based is based on [15].

$\sin^2(\theta_{12})$ and Δm_{21}^2 are primarily constrained by results from the KamLAND

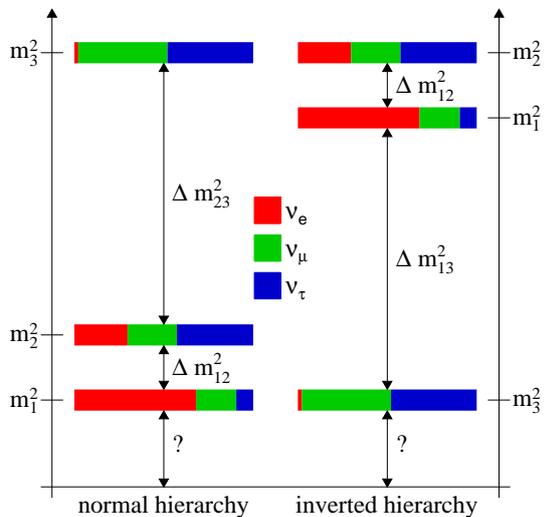


Figure 1.1: Illustration of normal and inverted neutrino mass hierarchy and flavour content of mass states. $\Delta m_{13}^2 \gg \Delta m_{12}^2$, so $\Delta m_{13}^2 \approx \Delta m_{23}^2$. Flavour composition drawn as $\sin^2(\theta_{12}) = 0.304$, $\sin^2(\theta_{13}) = 0.0219$, $\sin^2(\theta_{23}) = 0.51$.

Parameter	Best Fit (NH)	Best Fit (IH)
$\sin^2(\theta_{12})$	0.304 ± 0.014	
$\sin^2(\theta_{13})$	$(2.19 \pm 0.12) \times 10^{-2}$	
$\sin^2(\theta_{23})$	0.51 ± 0.05	0.50 ± 0.05
Δm_{21}^2	$(7.53 \pm 0.18) \times 10^{-5} \text{ eV}^2$	
Δm_{32}^2	$(2.44 \pm 0.06) \times 10^{-3} \text{ eV}^2$	$(2.51 \pm 0.06) \times 10^{-3} \text{ eV}^2$

Table 1.1: Global best fit oscillation parameters as per [5]. NH and IH stand for normal and inverted mass hierarchy respectively.

(Kamioka Liquid Scintillator Antineutrino Detector) experiment[17], an 18 meter diameter liquid scintillator detector, whose $\bar{\nu}_e$ flux primarily comes from 56 Japanese nuclear power reactors surrounding it; as well as various solar neutrino experiments and the SNO experiment.

Bounds on all three mixing angles, $\sin^2(\theta_{12})$, $\sin^2(\theta_{13})$, and $\sin^2(\theta_{23})$ as well as both Δm_{21}^2 and Δm_{32}^2 can be obtained by combining this with data from long baseline neutrino experiments, in particular from T2K[18, 19] and NO ν A[20, 21]. T2K (Tokai to Kamioka) uses a 295 km baseline between J-PARC in Tokai and Super-Kamiokande, a 50 kt water Cherenkov detector. Two near detectors, one on-axis (INGRID) and one off-axis (ND280) are used and the far detector is aligned 2.5° off-axis in order to achieve a fairly narrow 0.6 GeV peak neutrino energy [22, 19]. NO ν A uses a longer baseline (810 km) and higher energies (1-3 GeV) with a 14 kt far detector made of polyvinyl chloride (PVC) extrusions filled with liquid scintillator. Both the far and the near detector are aligned 0.84° off-axis.[20]

Significantly tighter bounds on $\sin^2(\theta_{13})$ and Δm_{32}^2 come from short baseline reactor experiments, especially Daya Bay and RENO, both of which use gadolinium doped liquid scintillator detectors to pick up the neutrino flux from nearby (less than 2 km) nuclear reactors.[23, 24] The combined dataset also allows some probing into the phase space of δ_{CP} .

A further improvement finally comes from adding in atmospheric neutrino results from Super-Kamiokande – the detector also used for the T2K experiment – and IceCube DeepCore, a densely instrumented subsection of the IceCube detector in Antarctica sensitive to lower energies than most of IceCube. These results predominantly help constrain $\sin^2(\theta_{23})$ and Δm_{32}^2 , but also provide some indication concerning the other oscillation parameters.[25, 26]

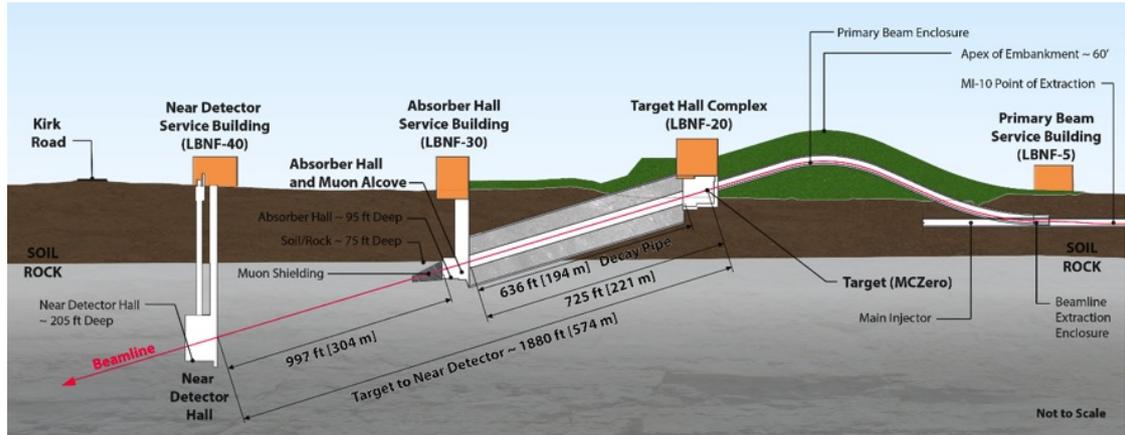


Figure 1.2: The DUNE beamline[28]

1.4 Next Generation Oscillation Experiments

Four next generation experiments are currently under development, designed to probe the neutrino oscillation parameters at unprecedented precision, two long baseline (DUNE and Hyper-Kamiokande), one medium baseline (JUNO), and one for atmospheric neutrinos (PINGU):

1.4.1 Deep Underground Neutrino Experiment (DUNE)

DUNE will send an extremely high intensity neutrino beam from the Long Baseline Neutrino Facility (LBNF) at Fermilab to a liquid argon time-projection chamber (LArTPC) detector with a total fiducial mass of 40 kt at the Sanford Underground Research Facility (SURF) – a beamline length of 1300 km. The near detector design is currently under consideration. With regard to neutrino oscillations, DUNE’s primary aims are measurement of θ_{23} and δ_{CP} , as well as the determination of the neutrino mass hierarchy. Within the first 15 years of data-taking, DUNE is expected to resolve θ_{23} to within 1° resolution for the entire possible parameter range and δ_{CP} to within 5 %. A 5σ result on the neutrino mass hierarchy should be available after at most 8.5 years of data-taking for the least favourable values of δ_{CP} . [27, 28]

1.4.2 Hyper-Kamiokande

Hyper-Kamiokande will use the same neutrino source, near detector, and 295 km beamline setup as T2K, though upgrades to the existing, or even new near detectors are under review. The main upgrade with respect to T2K will be the far detector: One or

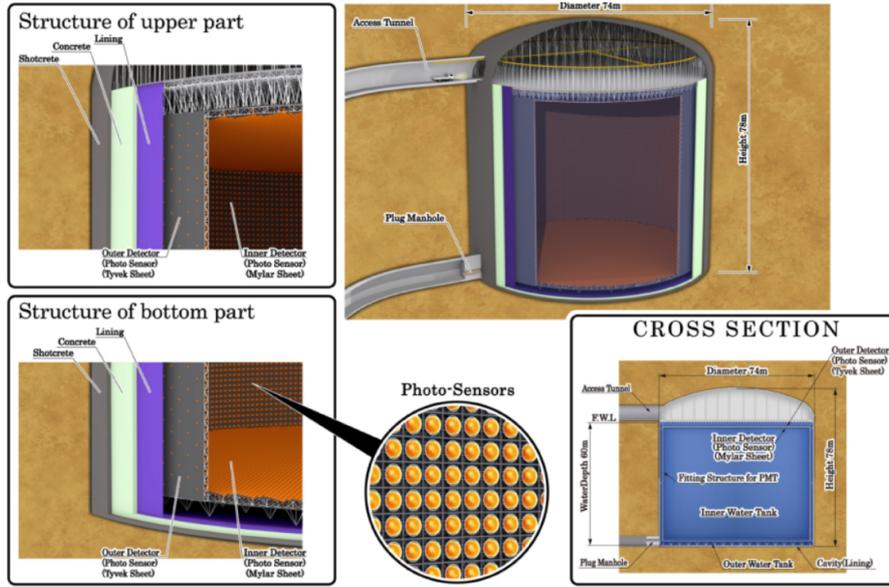


Figure 1.3: Schematic design of the first tank of the Hyper-Kamiokande detector. The second tank (if constructed) may be placed at a different site.[30]

two tanks, each holding 258,000 tons of ultra-pure water and instrumented by 40,000 20” diameter PMTs, with an additional surrounding layer more sparsely instrumented by 8” PMTs. The expected discovery potential for the scenario where a second detector is built at a separate site in Korea 1.5° off-axis is a 5σ result on CP violation for 60 % of the possible values of δ_{CP} and rejection of the wrong mass ordering at 5σ for all possible values of δ_{CP} within 20 years at 1.3 MW beam power. More details regarding the discovery potential for various possible scenarios can be found in [29].

1.4.3 Jiangmen Underground Neutrino Observatory (JUNO)

JUNO is a reactor neutrino experiment with an unusually long baseline of 53 km to be constructed in Jiangmen, China. The detector design is not finalised, though one proposed option can be seen in figure 1.4, using 20 kt of a linear alkylbenzene based liquid scintillator as the main detector volume. A significant and currently unresolved technical challenge is the required energy resolution of $3\%/\sqrt{E(\text{MeV})}$. If successfully constructed, 6 years of running will provide a $3 - 4\sigma$ result on the neutrino mass hierarchy, as well as better than 1 % resolution on $\sin^2\theta_{12}$, Δm_{21}^2 , and $|\Delta m_{31}^2|$.[31]

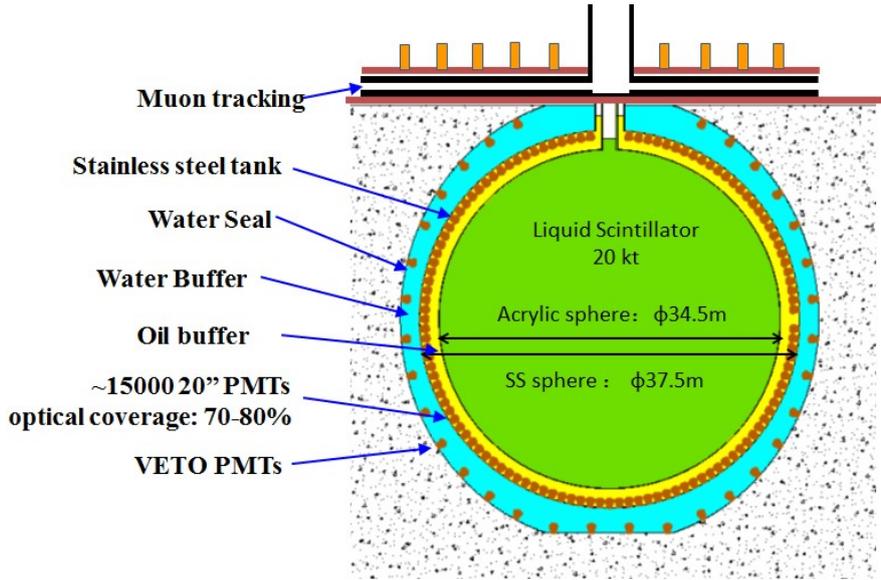


Figure 1.4: One proposed detector design for JUNO[31]

1.4.4 Precision IceCube Next Generation Upgrade (PINGU)

A relatively low-cost (compared to other next generation experiments) upgrade to IceCube, PINGU will add several thousand additional optical modules within the IceCube DeepCore region to achieve higher particle rates and sensitivity to significantly lower energy threshold (< 5 GeV). It aims to measure θ_{23} and Δm_{32}^2 , as well as determine the mass ordering to 3σ within the first four years of operation. Within that time, a 90 % C.L. for the octant of θ_{23} can be achieved for $0.46 \lesssim \sin^2 \theta_{23} \lesssim 0.55$ given a normal mass hierarchy and $0.44 \lesssim \sin^2 \theta_{23} \lesssim 0.57$ given an inverted mass hierarchy.[32]

1.5 The Neutrino Factory

Current neutrino experiments – and even the next generation accelerator-based experiments described above – are subject to limitations which severely limit performance both in terms of statistics and systematic errors. These limitations are related to the typical method of producing neutrinos: Protons are fired into a target producing a spray of pions, which is sign-selected and focussed. The pions decay to muons and neutrinos at which point the muons are filtered out of the beam by a beam dump. The neutrinos – mostly unimpeded – continue on to the detector(s). A neutrino beam produced in this way has a very large transverse momentum spread and thus – especially for a

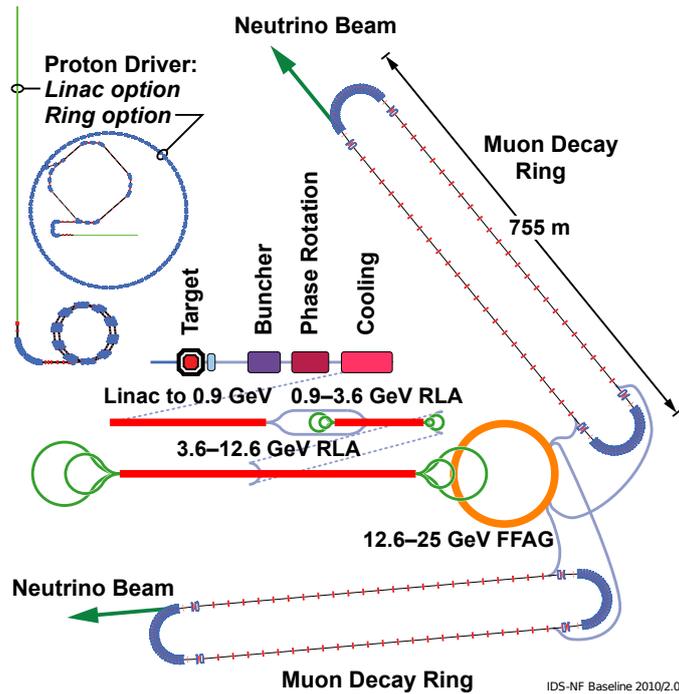


Figure 1.5: Layout for a 25 GeV neutrino factory[33]

baseline of hundreds to thousands of kilometres – the neutrino flux at the detector will be fairly low. Additionally, the exact spectrum of pions produced in proton interactions with the target is somewhat poorly understood, which makes accurate calculations of the produced neutrino spectrum difficult.

The proposed solution for these problems is the Neutrino Factory – a facility that can produce well-defined neutrino beams of unparalleled intensity. Rather than filtering out the muons produced in the pion decays and using the neutrinos, a neutrino factory collects, focusses and accelerates the muons before placing them in a storage ring. Neutrinos are then obtained from muon decay:

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu \quad \mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \quad (1.18)$$

There are a number of significant technical challenges associated with the construction of a neutrino factory – primarily the target system and cooling channel – which are and have been the subject of significant R&D efforts. The neutrino factory baseline as envisaged by the International Design Study for the Neutrino Factory (IDS-NF)[34] is outlined below. The specifications have been optimised to produce 10^{14} decays of ~ 25 GeV muons per second.

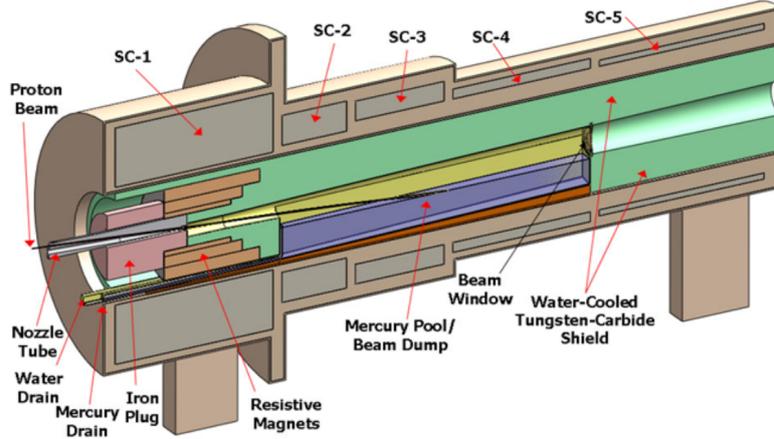


Figure 1.6: Concept design for the liquid mercury jet target system[34]

1.5.1 Proton Driver

The proton driver for a neutrino factory needs to deliver a 4 MW beam of 5–15 GeV protons to the target. Additionally, a specific timing structure is imposed, consisting of three very short (1–3 ns RMS length) bunches separated by about $120 \mu\text{s}$, which requires a dedicated bunch compression system. Three proton driver designs are proposed, for construction of a neutrino factory at CERN, FNAL, and RAL respectively. At CERN, the beam would be delivered by a 5 GeV upgrade of the Superconducting Proton Linac (SPL). FNAL would require an upgrade of the proposed Project X linac, which at base specifications will have a beam power significantly lower than required. At RAL, the neutrino factory would not be fed by a linac, but by an upgrade of the ISIS accelerator ring.

1.5.2 Target

The significant required input beam power of a neutrino factory would quickly destroy most conventional target systems. A free liquid mercury jet was proposed [35] to mitigate this issue, with the mercury cycling through and being collected in a pool which at the same time serves as a proton beam dump. In order to allow for the capture of both signs of charged pions, the target system uses solenoid magnets creating a field that tapers from 20 T to 1.5 T within a total target assembly length of 15 m. A side effect of the strong magnetic field is a mitigation of the jet disruption created by the incident proton beam, allowing for the desired 50 Hz repetition rate.

The total mercury flow speed is 20 m/s at a jet diameter of 8 mm, resulting in a

flow rate of 1.0 l/s. At this speed, the jet curvature due to gravity is negligible compared to the jet diameter over two nuclear interaction lengths (≈ 30 cm) and each new beam pulse (at a 20 ms separation) will encounter at least two nuclear interaction lengths of target material undisturbed by the previous pulse.

1.5.3 Muon Front-End

The pions produced in the target travel along a decay and drift channel of 57.7 m length where they decay as

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad \pi^- \rightarrow \mu^- + \bar{\nu}_\mu \quad (1.19)$$

with the analogous decay to electrons / positrons having a near-zero branching fraction due to helicity suppression.

Additionally, the originally very tight pion bunches disperse due to their large energy spread, resulting in a position-energy correlation of both pions and daughter muons. The expected rate of muons at the end of the decay channel is about 0.4 of each sign per incident 8 GeV proton.

The decay channel is followed by a buncher – which forms the muons into a bunch train using RF cavities – and a phase-energy rotator. The latter exploits the position-energy correlation created in the drift length to reduce the longitudinal momentum spread by decelerating earlier bunches and accelerating later ones.

The final component of the muon front-end is the cooling channel. In it, the *emittance* of the muon beam (a measure for the the phase-space volume it occupies) is reduced by passing it through alternating layers of a low- Z absorber, which reduces the beam momentum in all directions, and RF cavities, which increase the momentum in beam-line direction only. This technique, called *ionisation cooling*, is used as conventional cooling methods are too slow compared to the 2.2 μ s lifetime of the muons and is described in more detail in section 2.2. A neutrino factory requires the beam to be cooled in 4 phase-space dimensions (i.e. transverse cooling), whereas for a muon collider (see section 1.6), 6D cooling is required.

1.5.4 Acceleration

When exiting the cooling channel, the muons have a mean energy of 0.24 GeV. A linac increases this to 0.9 GeV, making the muons strongly relativistic ($\gamma \sim 8.6$), preparing them for efficient acceleration in the following stages.

Next, the muons enter a set of two recirculating linear accelerators (RLAs) via two double chicanes, increasing the energy to 3.6 and 12.6 GeV respectively. The RLAs

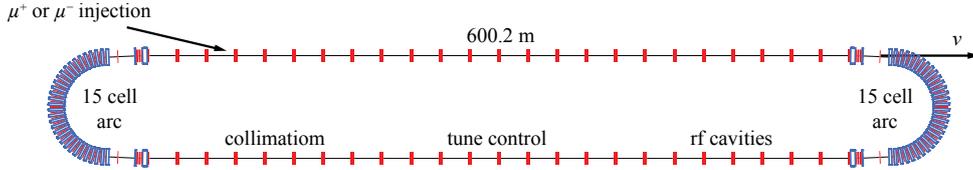


Figure 1.7: Racetrack design for the neutrino factory storage ring[34]

have a “dog-bone” shape (see figure 1.5), facilitating the simultaneous acceleration of muons of both charge signs. In each RLA, 4.5 acceleration passes are performed, with each pass in RLA 1 providing 0.6 GeV and each pass in RLA 2, 2.0 GeV.

Final acceleration to the desired 25 GeV occurs in a fixed-field alternating gradient accelerator (FFAG), where the muons perform a total of 11 passes before being extracted by kicker magnets and passed into the storage ring.

1.5.5 Storage / Decay Ring

The storage / decay ring (see figure 1.7) is the final segment of the neutrino factory. One of the straight segments is pointed at the far detector (angled into the ground at 18° (36°) for a baseline length of 4000km (7500km)), whereas along the return straight the muons undergo collimation, RF, and tune control. Alternative designs under consideration allow counter-rotating μ^+ and μ^- beams, or use a triangular shape for two neutrino beam directions. In addition to the challenging optical properties of the storage ring arising from the requirement that the angular divergence of the muon beam be much smaller (~ 0.4 mrad) than the natural divergence of the resulting neutrinos due to the kinematics of muon decay, precise instrumentation is also required to measure the muon energy and polarisation, as well as beam divergence and current. For the beam current, standard beam current monitors are sufficient, but the details regarding the other measurements are still under investigation.

Design efforts for a strongly simplified version of a Neutrino Factory, a facility called ν STORM (Neutrinos from Stored Muons), are currently ongoing. ν STORM would rely entirely on established technology (most notably omitting the cooling channel) to generate neutrino beams with similar characteristics as a Neutrino Factory (though of course significantly lower intensity).[36]

1.6 Muon Colliders

Muon cooling can not only tremendously benefit the field of neutrino physics, it can also be used to significantly push the boundaries of collider physics. For the probing of electroweak interactions, lepton colliders offer a significant advantage over hadron colliders in that they collide elementary particles with exactly defined initial states. This leads to much cleaner collisions and thus higher precision measurements, as well as having the entire particle energy – rather than only the parton energy – available for interactions. In the case of electrons / positrons however, as is the case with all current lepton colliders, this comes with a substantial drawback, as the energy achievable in a lepton collider of realistic size is severely limited by synchrotron radiation, with the classical energy loss per revolution δE [5] given by

$$\delta E = \frac{4\pi}{3} \frac{e^2}{R} \beta^3 \gamma^4, \quad (1.20)$$

where e is the particle charge, R is the radius of the accelerator ring, and β , γ are $\frac{v}{c}$ and the Lorentz factor, respectively, or

$$\delta E(\text{MeV}) \approx 0.0885 \frac{[E(\text{GeV})]^4}{R(m)} \quad (1.21)$$

for the specific case of ultrarelativistic electrons or positrons. Hence to keep the energy loss at feasible levels, one can either increase the radius of the ring (problematic due to a high dependence of cost on the radius) or the mass of the lepton (as the factor of γ^4 provides a mass dependence of m^{-4}). Hence the idea of a muon collider, which – due to the significantly higher mass of the muon – reduces synchrotron radiation by almost a factor of 2×10^9 with respect to an electron-positron collider of the same energy.

As mentioned in section 1.5.3, muon colliders and neutrino factories share the problem that due to the short muon lifetime, single-pass cooling is necessary before the muons can be accelerated to suitable energy, and in fact both facilities are essentially identical up to the initial cooling step (see figure 1.8). A muon collider requires additional 6D cooling, and the acceleration requirements differ, but the overlap is large enough that a neutrino factory can be considered the first stage of the construction of a muon collider. A muon collider can also serve as a Higgs factory and would be uniquely suitable to provide a precision measurement of the width and shape of the standard model Higgs resonance [37], with the $\mu\mu \rightarrow \text{H}$ cross section enhanced by a factor of $\left(\frac{m_\mu}{m_e}\right)^2$ with respect to $ee \rightarrow \text{H}$.

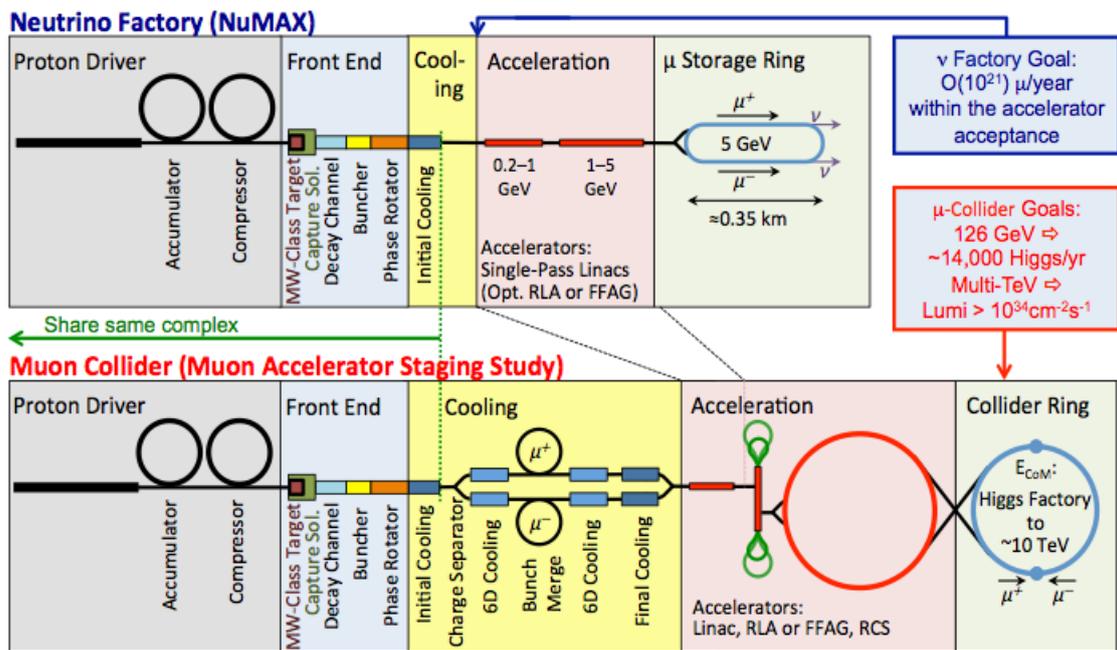


Figure 1.8: Conceptual block diagrams for a neutrino factory and a muon collider[37], demonstrating the high proportion of shared elements. Note that a lower energy neutrino factory than the International Design Study is depicted, and for an IDS-NF, there would also be significant overlap in the acceleration segment (see also figure 1.5).

Chapter 2

MICE

“We are stuck with technology when what we really want is just stuff that works.”

Douglas Adams

An integral part of both a muon collider and a neutrino factory is the cooling channel, a beamline segment that reduces the phase-space volume occupied by the muon beam (i.e. its *emittance*) before it is injected into the storage ring or accelerator. The Muon Ionisation Cooling Experiment (MICE) was constructed in order to demonstrate a viable method of muon cooling in order to lay part of the groundwork for future construction of such facilities.

2.1 Beam Emittance

The following discussion of beam emittance is based on information from [38, 39].

Particles in a beamline can be represented by their six-dimensional phase space coordinates, e.g. as (x, p_x, y, p_y, s, E) , where p_x, p_y are the transverse momenta, s is the spatial coordinate along the particle’s trajectory, and E is the particle energy. Other parameterisations are possible, and may be more convenient depending on the application (e.g. using x' and y' , the trajectory slopes, rather than p_x and p_y , and cp rather than E). In linear beam dynamics, the coupling between the horizontal, vertical, and longitudinal plane is typically sufficiently small to ignore, which allows us to treat the three corresponding 2D phase spaces as independent. We can describe such a 2D phase space as an ellipse (see figure 2.1) described by the equation

$$\gamma_t x^2 + 2\alpha_t x x' + \beta_t x'^2 = \epsilon \tag{2.1}$$

where α_t , β_t , and γ_t are the Twiss parameters (also known as Courant-Snyder functions, the subscript t is used to distinguish them from the relativistic β and γ) and ϵ (in units of $\text{mm} \cdot \text{rad}$) is the beam emittance and corresponds to the area of the ellipse divided by π (some conventions include the factor of π in ϵ).

A special case of the emittance ellipse, namely for a quadrupole magnet, where the tilt of the ellipse is 0, can be derived from Hill's equation:

$$x'' + k(s)x = 0 \quad (2.2)$$

where $k(s)$ is a coefficient describing the distribution of focusing strength and is periodic on the scale of a beamline cell. This has a solution of the form

$$x = \sqrt{\beta(s)\epsilon} \cos[\phi(s) + \phi_0] \quad (2.3)$$

where $\beta(s)$ and ϵ are equivalent to β_t and ϵ from equation 2.1, $\phi' = \beta^{-1}$, and ϕ_0 is a constant phase. Differentiating with respect to s , we obtain

$$x' = -\frac{\epsilon}{\beta(s)} \sin[\phi(s) + \phi_0] + \left[\frac{\beta'(s)}{2} \right] \sqrt{\frac{\epsilon}{\beta(s)}} \cos[\phi(s) + \phi_0] \quad (2.4)$$

Taking $\beta'(s) = 0$, this describes an ellipse with no tilt, semi-axes $\sqrt{\epsilon\beta_t}$ and $\sqrt{\epsilon/\beta_t}$ and area $\pi\epsilon$, corresponding to the shape of the general ellipse in figure 2.1, with $\alpha_t = 0$. As transverse momentum decreases, transverse position increases and vice versa, i.e. a focussing/defocussing effect.

From the geometry of ellipses it can be shown that the Twiss parameters are correlated as

$$\beta_t\gamma_t - \alpha_t^2 = 1 \quad (2.5)$$

According to Liouville's Theorem [39], a beam that is only subject to conservative forces – such as e.g. from multipoles or solenoids – has a constant phase-space (and therefore constant ϵ). The individual values of the Twiss parameters are functions of s (the position along the direction of the beamline), causing the ellipse to rotate and skew, but the area of the ellipse is unaffected. The correlation between position and momentum is set by the parameter α and is defined such that for $\alpha > 0$ the beam is converging and for $\alpha < 0$ the beam is diverging. The physical size of the beam is described by β : a larger β increases the physical extent of the beam but reduces its momentum spread. γ is determined by α and β as given by equation 2.5. In a beam transport lattice, the parameters are periodic, i.e. at equivalent points in each lattice

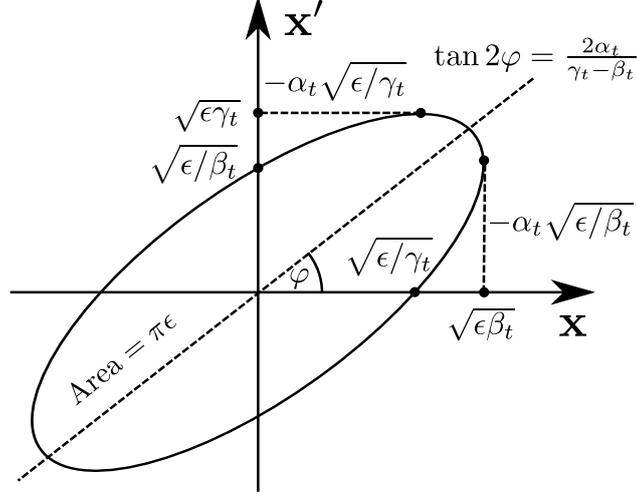


Figure 2.1: A phase-space ellipse described by the parameters α_t , β_t , and γ_t . x is the transverse spatial dimension and x' its derivative with respect to the beamline direction.

cell, the values of all parameters will be identical.

When dealing with accelerated beams, since the coordinates x and x' are not canonical (i.e. they do not fulfil the fundamental Poisson bracket relations of Hamiltonian mechanics, which Liouville's description is based on), Liouville's Theorem does not hold for the definition of emittance given in equation 2.1: the emittance appears to shrink as the beam is accelerated. For this reason, an alternate definition of the emittance, the *normalised* emittance ϵ_n is sometimes used, defined as

$$\epsilon_n = \beta\gamma\epsilon \quad (2.6)$$

where γ is the Lorentz factor, and $\beta = \frac{v}{c}$. ϵ_n will remain constant even under acceleration¹.

2.1.1 RMS Emittance

Real particle beams are generally not sharply defined but have density distributions which appear Gaussian. It is therefore useful to work with the RMS emittance ϵ_{rms} which describes a phase-space ellipse containing all particles out to a certain number of standard deviations of that Gaussian distribution. In the case of protons, ϵ_{rms} by

¹Note that in the case of electrons in a typical accelerator, ϵ will not be constant as the particles will radiate significant energy when accelerated and therefore Liouville's theorem does not apply.

convention describes a 2σ spread. Since for electron machines safe aperture sizes are much larger and more variable, 1σ was the chosen convention to avoid ambiguity. This is the convention that will be used in this document.

We define the RMS emittance as

$$\epsilon_{rms} = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2} \quad (2.7)$$

and the normalised RMS emittance as

$$\epsilon_{rms,n} = \beta\gamma\epsilon_{rms} \quad (2.8)$$

2.2 Ionisation Cooling

Beam cooling – i.e. a reduction of the beam emittance – is a requirement for many particle accelerators and depending on the exact requirements and setup, a number of different cooling methods can be used, such as stochastic cooling, electron cooling, or laser cooling:

Stochastic cooling measures the displacement of the centre of charge of a beam at one location in the beamline and then use this measurement to control a kicker magnet at another location to counteract this displacement.

Electron cooling uses an electron beam of well-defined momentum to cool a proton beam, as energy exchange between the two particle species results in the proton beam converging to the momentum of the (constantly renewed) electron beam.

Laser cooling uses a laser tuned to a Doppler-shifted difference between two ion energy levels to explicitly target particles of specific momenta.

Muon cooling presents a unique challenge, however, in that the short mean lifetime of the particles ($\sim 2.2\ \mu\text{s}$) renders most methods (that aren't already excluded for other reasons, such as laser cooling which only works with ions) unusable, as a significant fraction would decay before appreciable cooling has taken place. The proposed solution is ionisation cooling, whereby the muons are passed through a low- Z absorber, where ionisation energy loss causes a reduction of all momentum components, and re-accelerated by radio-frequency (RF) cavities, increasing the longitudinal momentum only. The net effect is then a reduced transverse momentum spread at constant total momentum (see figure 2.2). Longitudinal cooling can also be achieved, by dispersing a beam across a wedge absorber, such that higher momentum muons will pass through more absorber

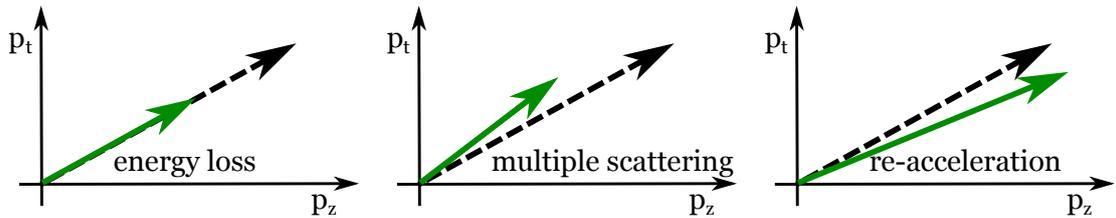


Figure 2.2: Illustration of the mechanism of ionisation cooling. The dashed black arrow represents the particle’s momentum vector before cooling, the green arrow the resulting vector due to the indicated physical processes.

material. Ionisation cooling is not feasible for most other particles, as hadronic interactions in the material (for hadrons) and bremsstrahlung (for electrons) would adversely affect cooling performance.[39]

While the absorber material is chosen in order to minimise multiple scattering, some degree of beam heating (i.e. an increase in emittance) in the absorber is unavoidable. The net cooling effect is thus determined by the interplay of cooling and heating in the absorber, and is given by the ionisation cooling equation (first put forward in [40], used in the MICE technical design report [41], and with a fit parameter of 13.6 MeV taken from [5]):

$$\frac{d\varepsilon}{ds} = \left(\frac{-\varepsilon}{\beta^2 E} \left\langle \frac{dE}{ds} \right\rangle \right)_{\text{cooling}} + \left(\frac{\beta_{\perp} (13.6 \text{ MeV})^2}{2\beta^3 E m_{\mu} X_0} \right)_{\text{heating}} \quad (2.9)$$

where E is the beam energy, $\beta = \frac{v}{c}$, β_{\perp} is the transverse beta function (i.e. one of the Twiss parameters, described above), m_{μ} is the muon mass, and X_0 is the radiation length of the material. The parameter of 13.6 MeV comes from a Gaussian approximation to a Coulomb scattering fit according to the Molière theory (discussed in more detail in section 4.1.3).

From this we can obtain the equilibrium emittance of the beam where $\frac{d\varepsilon}{ds} = 0$, i.e. the lowest possible emittance that can be achieved with the given absorber and beam parameters:

$$\varepsilon_{\text{eq}} = - \frac{\beta_{\perp} (13.6 \text{ MeV})^2}{2\beta m_{\mu} X_0 \left\langle \frac{dE}{ds} \right\rangle} \quad (2.10)$$

Hence to optimise the performance of the cooling channel (achieve maximal focussing), we want to maximise radiation length and energy loss in the absorber (which prescribes low- Z absorber materials) and minimize the transverse beta function of the beam. It should be noted that if longitudinal beam emittance is a concern, an additional heating

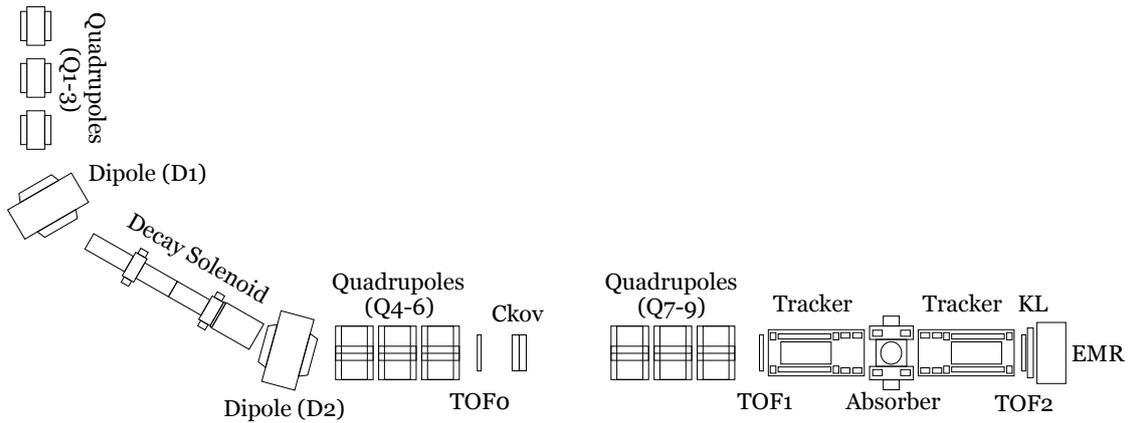


Figure 2.3: The MICE Step IV Beamline. The Quadrupoles are numbered Q1–Q9 starting from the top left, the dipoles D1 & D2, and the Time-of-Flight detectors as TOF0, TOF1, and TOF2. The Trackers are typically referred to as upstream and downstream tracker. Chkov, KL, and EMR are the Cherenkov detectors, the KLOE-Light Calorimeter, and the Electron Muon Ranger (see below).

factor due to energy straggling has to be considered.

2.3 MICE Target & Beamline

2.3.1 The ISIS Proton Synchrotron

ISIS is fed by an ion source that produces H^- ions, which are then passed into a 202.5 MHz radio-frequency quadrupole (RFQ) to shape the beam into discrete bunches with a separation of 4.94 ns. The following linac uses RF cavities to accelerate the beam to 70 MeV, providing 200 μs pulses to the synchrotron, where an aluminium oxide foil strips off all electrons. Once the ring is filled, the RF system shapes the beam into two bunches which are then accelerated to 800 MeV over about 10,000 orbits and finally sent to the two ISIS target stations by a kicker magnet. The whole process occurs at a repetition rate of 50 Hz, so the ring is filled and emptied every 20 ms.[42]

2.3.2 The Target

MICE runs parasitically off the ISIS proton beam by dipping the MICE target – a cylinder of inner radius 2.275 mm and outer radius 2.975 mm forming the end of a 528 mm long titanium rod – into the ISIS beam halo. In order to only sample protons at close to maximum energy, the target may enter the beam only for the last 1-2 ms of the 10 ms beam cycle. Due to the shrinking of the beam during the cycle as well as the

dip depth necessary to create the required muon flux, this results in a minimum travel distance of 24 mm at an acceleration of around 900 ms^{-2} . The required driving force is provided by a linear actuator consisting of sintered neodymium-iron-boron magnets on the target shaft placed inside a twenty-four coil stator.

Operation under these parameters can cause significant wear, frictional heating, and dust-production, thus minimising the coefficient of friction between shaft and bearings is crucial. As an added complication, the vacuum environment of the target prohibits the use of lubricating films. Vespel® SCP-5000 – a radiation-hardened polyimide – was chosen for the bearings, while the contact surfaces on the target shafts itself were coated in Diamond-like carbon (DLC).

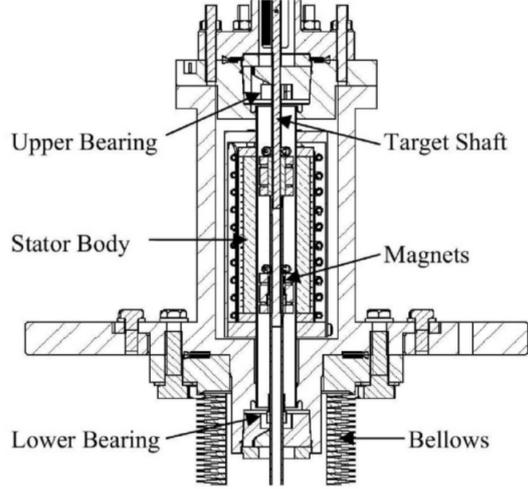


Figure 2.4: Schematic of the MICE target assembly[43]

The position of the target is measured via an optical vane read out by three fibre-coupled solid state lasers. The vane has 157 slots of width 0.3 mm providing a measurement resolution of $150 \mu\text{m}$. Using this readout, the dip depth can be set remotely via a parameter referred to as beam centre distance (BCD) – the distance of the tip of the target from the nominal beam centre. Detailed information about the MICE target can be found at [44].

A single dip of the target will typically result in tens of particles being captured in the MICE beamline, though these are then separated out by the trigger system into triggers containing usually only a single particle (see also section 3).

2.3.3 Beam Magnets

Particles produced in the target are captured by a quadrupole triplet (Q1–3) and pass through two dipole magnets (D1 & D2) for momentum selection which sandwich the so-called decay solenoid (DS). Transport to the cooling channel then occurs via two further sets of quadrupole triplets (Q4–9). A summary of the beamline positions, effective lengths, and maximum fields/field gradients for the magnets is given in table 2.1 and a schematic of the beamline in figure 2.3.

Quadrupoles

Q1-3 are Type-IV quadrupoles originally used in the NIMROD synchrotron between 1964 and 1978. They feature an effective circular aperture of 101.5 mm radius and are each powered by a 200 A, 30 V power supply. The section of the beamline passing through Q1-3 is contained within an aluminium vacuum pipe.

The other two sets of quadrupole triplets, Type QC quadrupoles, were previously used on DESY. They are powered by 400 A, 70 V power supplies and each have an effective circular aperture of 176 mm radius.

Dipoles

D1 & D2 are NIMROD Type-1 dipoles with horizontal and vertical apertures of 508 mm and 152 mm respectively. Together, they are used to select the momenta of both the particles (mostly pions) entering, and those (mostly muons) exiting the decay solenoid. D1 bends the beam by about 60° and is powered near yoke saturation by a power supply limited to 440 A, 240 V. D2 bends the beam by a much shallower angle ($\sim 30^\circ$) and is powered by a significantly weaker 200 A, 100 V supply. By changing the relative field strengths of the two dipoles, it is also possible to select between pion mode – where the momentum selection of both dipoles is the same – and muon mode – where the momentum selection of D2 is about half of that of D1 – which only select backwards-decayed muons and almost entirely eliminates pion contamination:

For a backwards-decaying muon from pion decay, the momentum in the lab frame is given by (in natural units):

$$\begin{aligned} p'_\mu &= \gamma (v_\pi E_\mu - p_\mu) \\ &= \frac{\sqrt{p_\pi^2 + m_\pi^2}}{m_\pi} \left(\frac{p_\pi}{\sqrt{p_\pi^2 + m_\pi^2}} \frac{m_\pi^2 + m_\mu^2}{2m_\pi} - \sqrt{\left(\frac{m_\pi^2 + m_\mu^2}{2m_\pi} \right)^2 - m_\mu^2} \right) \end{aligned} \quad (2.11)$$

where the order of terms in the Lorentz transform has been reversed from the common form to account for the relative directions.

MICE operates in a pion momentum range roughly around $m_\pi < p_\pi < 3.5m_\pi$. Setting $p_\pi = xm_\pi$ (where x is a scaling variable), and evaluating the constant terms, the above equation simplifies to

$$p'_\mu \approx \sqrt{1 + x^2} \left(\frac{x}{\sqrt{1 + x^2}} 109.8 \text{ MeV} - 29.8 \text{ MeV} \right) \quad (2.12)$$

Magnet	Distance from target	L_{eff}	Max field/gradient	
	(mm)	(mm)	(T)	(T/m)
Q1	3000.0	888	-	1.6
Q2	4400.0	888	-	1.6
Q3	5800.0	888	-	1.6
D1	7979.1	1038	1.6	-
DS	12210.7	5000	5.7	-
D2	15808.1	1038	0.85	-
Q4	17661.6	660	-	2.3
Q5	18821.6	660	-	2.3
Q6	19981.6	660	-	2.3
Q7	25293.7	660	-	2.3
Q8	26453.7	660	-	2.3
Q9	27613.7	660	-	2.3

Table 2.1: Summary of the MICE beamline magnets.[43] The distance from target is measured along the nominal beam axis.

For the aforementioned momentum range, this implies $0.48p_\pi \lesssim p'_\mu \lesssim 0.56p_\pi$, i.e. a momentum selection of D2 approximately half that of D1. Due to some technical issues causing a reduced particle rate such as the failure of one of the tracker matching coils (see section 2.5.4) and intermittent availability of the decay solenoid, MICE is currently operated in pion mode, which provides higher statistics while the increased pion contamination can be compensated for due to the stronger time-of-flight separation between muons and pions compared to muon mode.

Decay Solenoid

The decay solenoid (DS) is the only superconducting magnet in MICE outside of the cooling channel. Its purpose is to increase the muon flux by decreasing the amount of particles that will exit the beamline transversely. Originally used in the $\mu E4$ beam (a low energy muon beam at the Paul Scherrer Institute in Switzerland), the solenoid has an aperture radius of 57.5 mm and in MICE is typically operated at a current of 870 A, well below its rating of 1000 A, to produce a 5 T operating field.

Detailed information about the MICE beam magnets can be found at [43].

2.3.4 Proton Absorber

When running with positive beam polarity, a significant proton flux can reach as far down the beamline as the first time-of-flight detector, TOF0, resulting in a large proton

background in the detector. To reduce this background, borated polyethylene sheets can be lowered into the beamline between the decay solenoid and D2 in order to absorb protons (as well as some neutrons), while having a limited effect on pions and muons due to their differing energy loss behaviour in the material. There are four sheets of thicknesses 15 mm, 29 mm, 49 mm, and 54 mm, which can be independently lowered into the beamline thereby providing a range of total thicknesses up to 147 mm of material, depending on the expected momentum ranges due to the settings of D1 & D2. 147 mm of borated polyethylene are sufficient to stop protons with momenta of up to ~ 500 MeV/c. [43]

2.3.5 Diffuser

As the aim of MICE is to measure a reduction in beam emittance, it is important that the emittance of the incoming beam can be accurately controlled. For this purpose, a diffuser consisting of 4 irises that can be individually opened and closed is placed in the beamline just before the first scintillating fibre tracker (see section 2.5.4). The materials and thicknesses of the irises (2.97 mm brass, 5.94 mm brass, 2.80 mm tungsten, 5.60 mm tungsten) are chosen such that the total material inserted into the beamline can be set to up to 3 radiation lengths (X_0) in steps of $0.2 X_0$. The diffuser is operated by an air-driven actuator, as operation by electric motors would not be possible in the 4 T field of the upstream spectrometer solenoid (see section 2.5.4) in which it is located.[43]

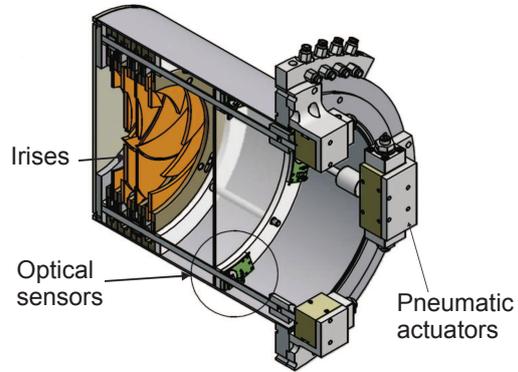


Figure 2.5: Model of the MICE Diffuser[43]

2.3.6 Partial Return Yoke

To mitigate the stray fields produced by the spectrometer solenoids housing the scintillating fibre trackers (section 2.5.4), a partial return yoke encloses the entire cooling channel, including both tracker assemblies on four sides (i.e. except for top and bottom, hence “partial” return yoke) in around 60 tonnes of low carbon steel.

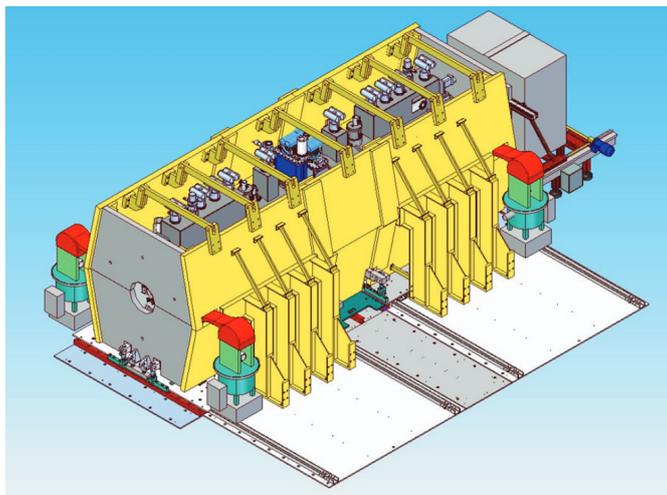


Figure 2.6: Design of the MICE partial return yoke[45]

2.4 Cooling Channel

2.4.1 Step IV Configuration

The Step IV configuration of MICE does not provide for re-acceleration of muons and only contains a single absorber station. The absorber can be switched out between runs, so that data taking can proceed with both liquid hydrogen (contained in a 21 litre, thin aluminium vessel, figure 2.7a) and lithium hydride (in the form of a 65 mm disk, figure 2.7b). The possibility of inserting a wedge-shaped absorber, which would allow 6D cooling, is under investigation. Due to the difficulty of manufacturing a lithium hydride absorber in this shape, the wedge absorber would consist of polyethylene, which offers slightly worse cooling performance.

The so-called absorber focus coil (AFC) (figure 2.7a) actually consists of two superconducting magnet coils in an assembly designed to both house the absorber itself and – in the case of the liquid hydrogen absorber – provide the connection to the liquid hydrogen system. The coils themselves can have their polarities set individually, thus allowing MICE to run both in “solenoid mode” – where the direction of the magnetic field is uniform throughout the beamline – and “flip mode” – where the direction flips in the centre of the absorber (with the downstream tracker solenoid consequently switching polarity as well). Solenoid mode requires lower currents to operate at sufficient field strength, but for an extended cooling channel consisting of many cells (as would be required in a neutrino factory or muon collider) it would result in canonical angular momentum growth and thus a mismatch of the beam to the lattice.

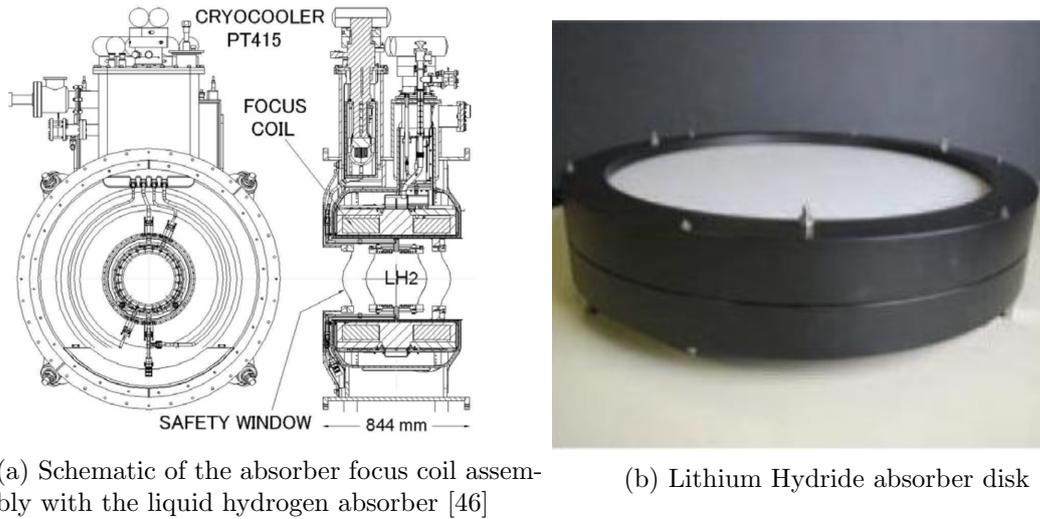


Figure 2.7

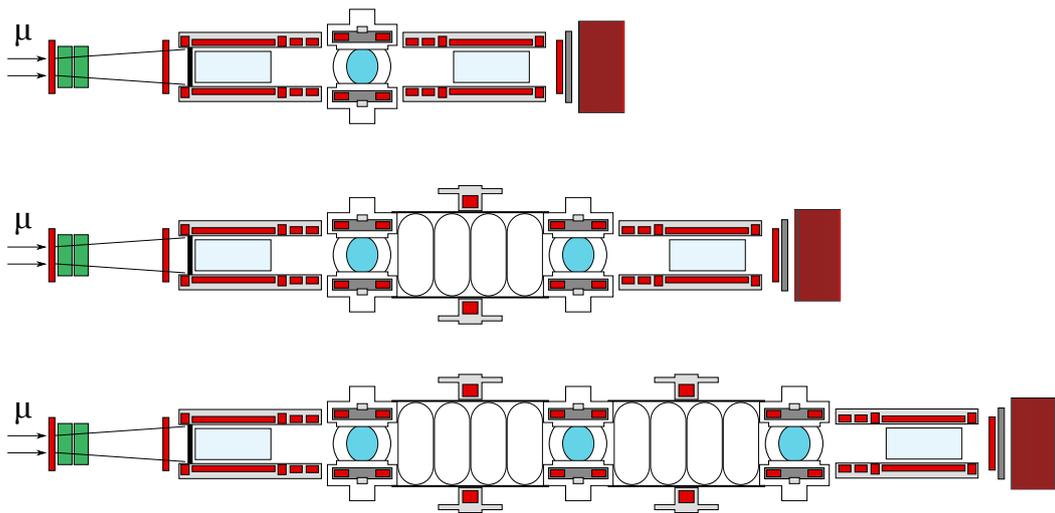


Figure 2.8: MICE steps IV, V, and VI as originally planned where each step after VI adds an additional absorber as well as an RF module consisting of four RF cavities.

2.4.2 Demonstration of Ionisation Cooling

In order to demonstrate *sustainable* muon ionisation cooling, i.e. a loss in emittance without an associated loss in total momentum, re-acceleration of the muons is required. To that end, MICE was originally envisaged to culminate in what was designated Step VI, with a cooling channel consisting of three absorbers and two RF cavity modules

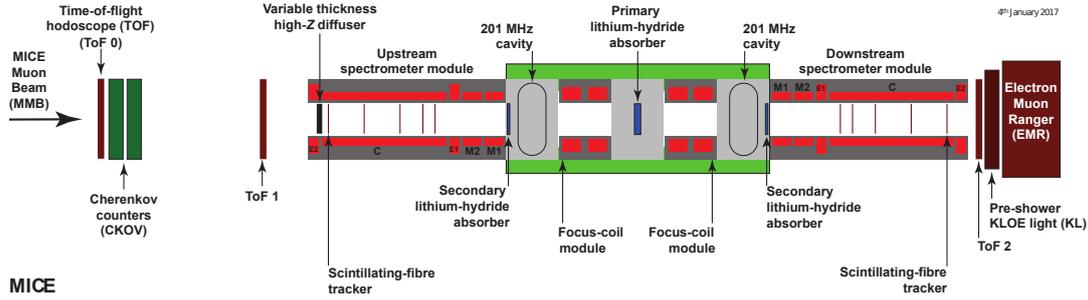


Figure 2.9: “Demonstration of Ionisation Cooling” configuration of the MICE beamline[47]

(figure 2.8). A number of delays, partly due to faulty components, as well as the US Department of Energy’s decision to cease US involvement in MICE in 2017, followed by decisions of other agencies not to extend funding of MICE, resulted in a redesign of the beamline for the planned “Demonstration of Ionisation Cooling” (figure 2.9). The new design contains two single RF cavities and three lithium hydride absorber disk, placing the focus coil modules that would have been used for two liquid hydrogen targets in between the absorber disks.

The failure of one of the matching coils in the downstream tracker led to the design of a descoped version of this configuration (figure 2.10) which entirely removes the downstream tracker coils and could provide a cooling demonstration at very low cost. This descoped beamline would contain a single RF cavity as well as one of the focus coil modules, sandwiched between two lithium hydride absorber disks. The downstream tracker would be modified to only consist of three stations without a magnetic field and the

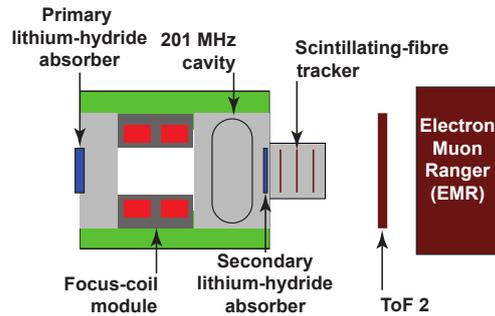


Figure 2.10: Downstream section of the descoped “Demonstration of Ionisation Cooling” configuration of the MICE beamline[47]

KL (see section 2.5.5) would be removed to improve the momentum reconstruction of the EMR, since the straight tracks in the downstream tracker would provide no momentum information. Funding for either of these configurations is currently being sought.

2.5 Detectors

2.5.1 Luminosity Monitor

The luminosity monitor is the only MICE detector not actually placed in the MICE beamline but inside the ISIS ring 10 meters from the target and at an angle to it equal and opposite to that of the first section of the MICE beamline (see figure 2.11b). It provides a measurement of the particle rate in MICE independent of the beam loss monitors installed in the ISIS ring. The monitor itself consists of two scintillator pairs read out by photomultiplier tubes, separated by a 150 mm thick sheet of polyethylene that shields out low-energy protons and pions. Coincidence counts of both pairs, as well as of the whole four scintillator system are recorded to determine the particle rate.[48]

2.5.2 Time-of-Flight Detectors

The time-of-flight detectors (TOFs) primarily provide PID discrimination by measuring the particle speed which – together with the approximate momentum as set by the dipole magnets (section 2.3.3) as well as the measured momentum in the upstream tracker (section 2.5.4) – can be used to estimate the particle mass and thus particle type. Furthermore, TOF1 is typically used as a trigger, though if needed other detectors can be used for this purpose as well.

All TOFs have a similar layout, consisting of two planes of 25 mm thick scintillator bars (40 mm wide for TOF0, 60 mm wide for TOF1 & TOF2) arranged horizontally and vertically, respectively (see figure 2.12). Flat fish-tail PMMA lightguides are then used to direct the scintillation light into R4998 Hamamatsu photomultiplier tubes (PMTs) at both sides of each bar.[49, 50].

The TOF detectors all provide timing resolutions between 50 and 60 ps.[51]

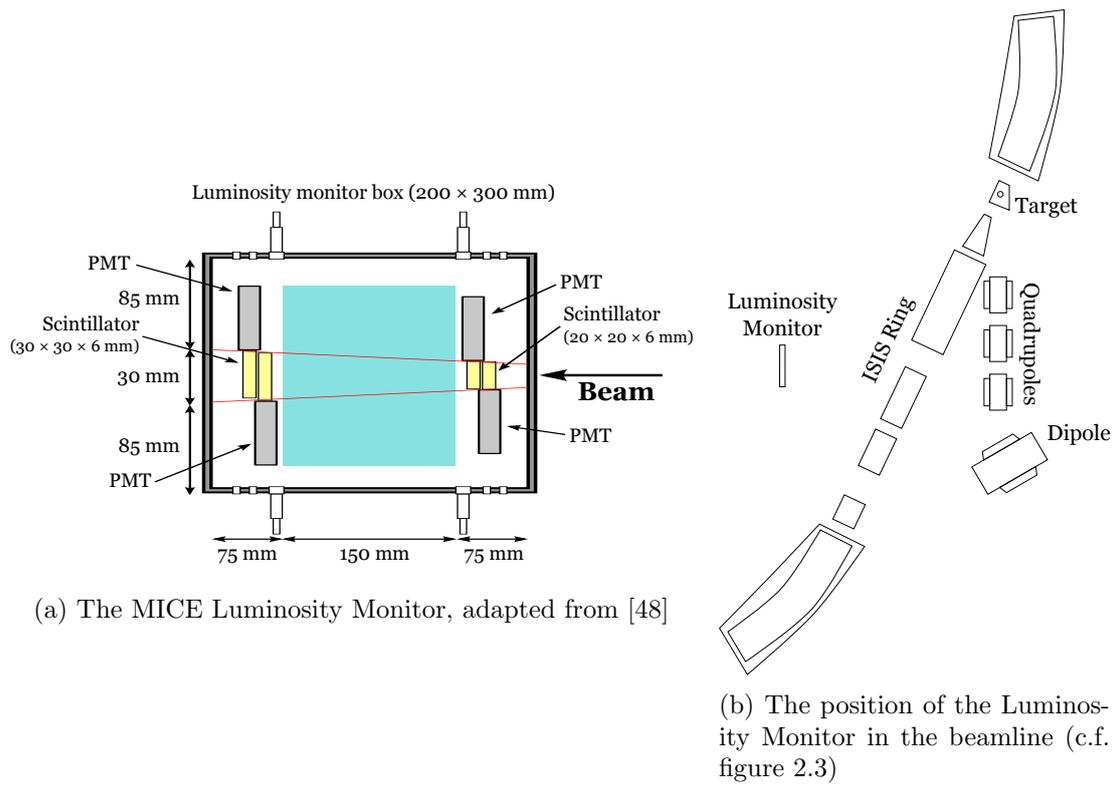


Figure 2.11

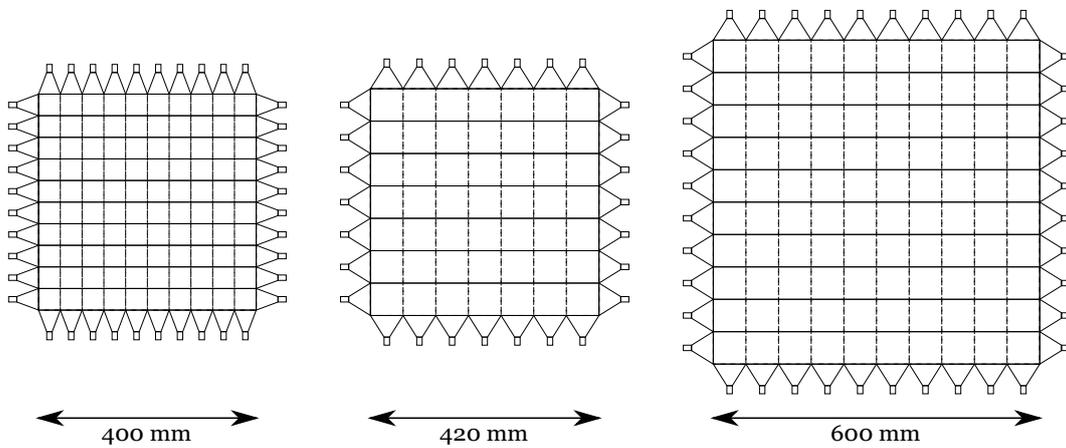


Figure 2.12: Schematic of the time-of-flight detectors. From left to right TOF0, TOF1, TOF2

2.5.3 Cherenkov Detectors

The Cherenkov detectors are primarily designed to provide μ - π separation in the higher momentum ranges, where TOF separation is not sufficient for conclusive particle identification. In order to provide separation over a large range of momenta, two aerogel Cherenkov detectors (CkovA & CkovB) with refractive indices $n = 1.07$ and $n = 1.12$ are used. They are each read out by four 200 mm EMI9356KA PMTs and placed directly one after another in the beamline. Their respective thresholds provide different responses in four distinct momentum ranges (see figure 2.13). Below the CkovB muon threshold of about 209 MeV/c, where there is no μ - π separation, the TOFs provide good separation (e.g. $\Delta t_{\text{TOF}1/0} \geq 2ns$)

whereas the momentum range above the CkovA pion threshold (≈ 306 MeV/c) is outside of the MICE running parameters.[41]

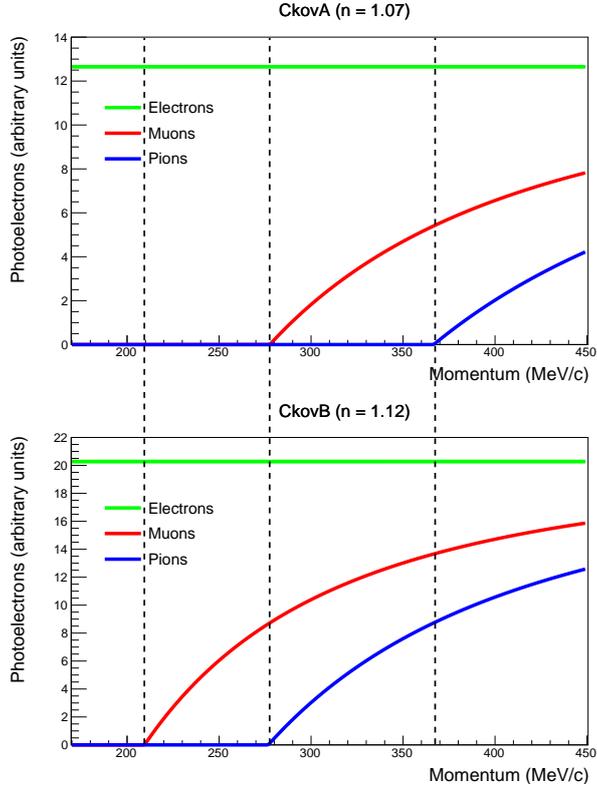


Figure 2.13: Cherenkov threshold curves for electrons, muons, and pions in both detectors. Four momentum ranges are marked which provide different PID separation. Figure adapted from [41]

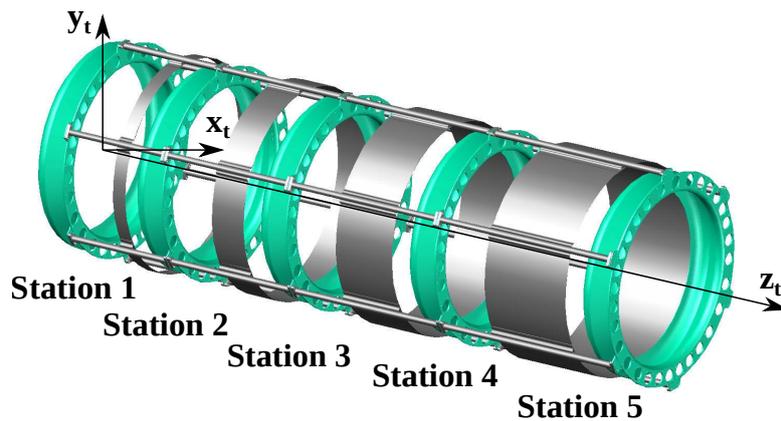
2.5.4 Scintillating Fibre Trackers

The scintillating fibre trackers on either side of the cooling channel form the most complex and important part of the detector system. They provide precise position and momentum information, thereby allowing the emittance of the beam before and after the cooling channel to be calculated.

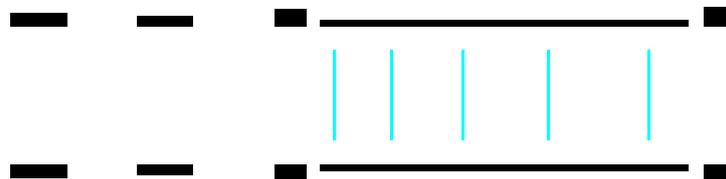
Each tracker consists of five scintillating fibre stations (see figures 2.14a and 2.14b) mounted on a carbon fibre frame at varying spacing (20, 25, 30, and 35 cm) in order to eliminate degeneracy in the path reconstruction, placed inside a 4 T magnetic

field.

The field is provided by a number of superconducting coils; a centre coil flanked on both sides by end coils, producing a field inside the detector volume uniform to within 0.3 %, and two matching coils two match the field to that of the absorber focus coils (see section 2.4.1). The magnet assembly of each detector consists of a total of 55,346 m of superconducting wire inside a 1390 kg cold mass, cooled to 4.2 K by two cryogenic coolers. Due to a fault in one of the matching coils of the downstream tracker in 2016, the downstream matching coils are not powered, resulting in lower transmission.



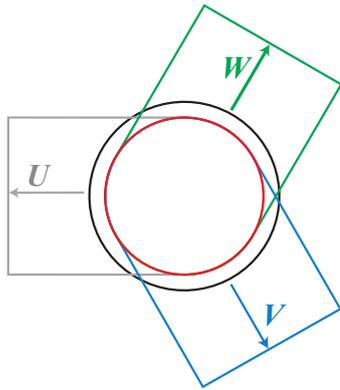
(a) Schematic of a tracker frame. The z-orientation corresponds to that of the downstream tracker. [52]



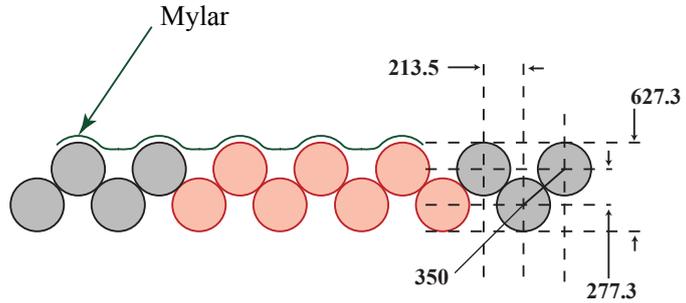
(b) Schematic of the arrangement of the tracker planes and coils. Coils are black, fibre planes light blue. The centre coil is surrounded by two end coils and there are two matching coils at the end of the tracker assembly facing the cooling channel. Total distance between the outermost fibre planes is 110 cm.

Figure 2.14

Every tracker station contains three doublet layers of 350 μm scintillating fibres, doped with (by mass) 1.25 % para-terphenyl (pT) and 0.25 % 3-hydroxyflavone (3HF). The layers are oriented at 120° towards each other, producing a partially degenerate internal coordinate system (see figure 2.15a). The arrangement of the fibres within a



(a) The arrangement of the layers in a tracker plane forming the u-v-w internal coordinate system. [53]



(b) Cross-section of part of a tracker plane. Marked are fibre-spacing and pitch as well as the the pattern of the seven-fibre grouping read out by a single light-guide. [53]

Figure 2.15

layer which can be seen in figure 2.15b ensures there is no dead inactive area within the layer that a particle could pass through. Every doublet layer consists of a total of 1491 fibres (1477 for the “u” view) connected in groups of 7 to a single 1.05 mm clear-fibre light guide, resulting in 213 (211 for “u” view) channels per layer and a $470 \mu\text{m}$ spatial resolution per layer. The instrumented radius is approximately 160 mm with a fiducial radius of 150 mm. The fibres are mirrored on the other end to maximise light yield. Layers are separated by $25 \mu\text{m}$ Mylar sheets to eliminate cross-talk between them. Read-out is done by high-gain visible light photon counters cooled to 9 K.

For details about track reconstruction in the MICE trackers, see section 2.6.3. For additional information about their design and construction, refer to [53].

2.5.5 KLOE-Light Calorimeter

The KLOE-Light Calorimeter (KL) is a sampling calorimeter based on the design of the electromagnetic calorimeter of the KLOE experiment in Italy [54]. It consists of scintillating fibres sandwiched between extruded lead foils in a volume ratio of $\sim 2 : 1$ (see figure 2.16) using significantly thinner lead foil than in the KLOE detector – hence the name KLOE-Light.

The fibres are arranged in 21 horizontal cells

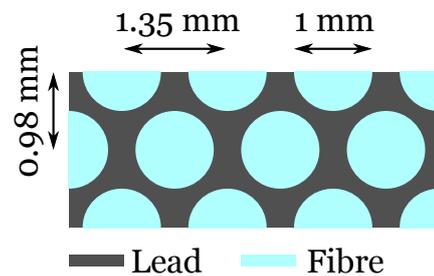


Figure 2.16: Schematic layout of scintillating fibres and extruded lead in the KL

(thus also providing very coarse y position information) read out on both sides by Hamamatsu R1355 PMTs. The KL's main purpose is to distinguish between muons and decay electrons, though at lower energies, the discrimination is relatively poor.[51, 55]

2.5.6 Electron Muon Ranger

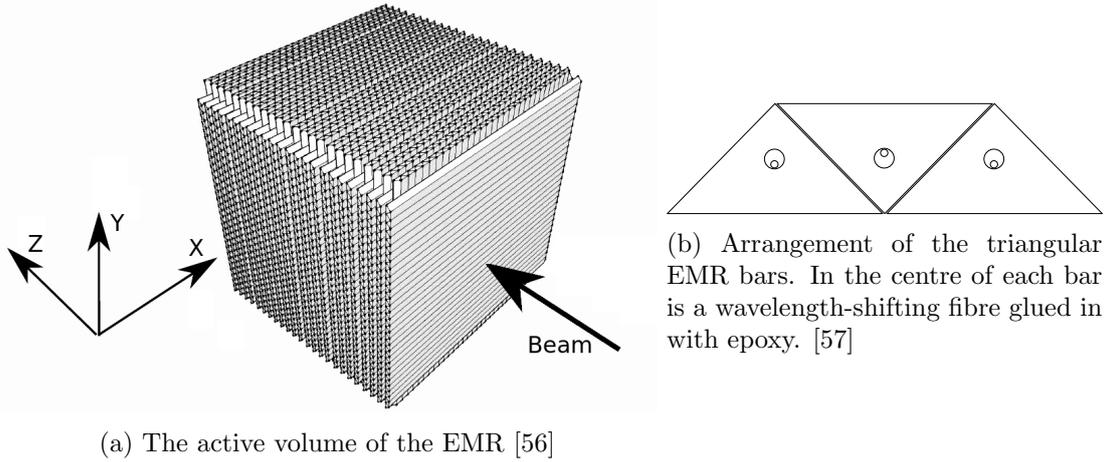


Figure 2.17

The Electron Muon Ranger (EMR) is a fully active scintillating tracking calorimeter which constitutes the downstream end of the MICE beamline. Its main purpose is to distinguish between electrons and muons, though for low- and medium-energy muons which are stopped in the detector, it also provides an energy estimate from the range. Two main discrimination variables are used, plane density (i.e. what proportion of EMR planes between the first and the last hit contain hits from a particle) and shower spread; with especially the former displaying excellent performance. In a multivariate analysis of both, the averaged muon selection efficiency over typical MICE beam settings was found to be $> 99\%$ with a purity $> 99.8\%$. [57] The EMR has a total active volume of nearly 1 m^3 , consisting of 48 planes, alternating between horizontal and vertical orientation, each made up of 59 tessellated triangular scintillator bars. A pixel can thus be reconstructed from hits in two adjacent planes, providing the x and the y coordinate respectively. Each bar contains a wavelength-shifting (WLS) fibre and is read out on one end by a pixel of a 64-channel multi-anode photo-multiplier tube (MAPMT), on the other end jointly with all other bars in the plane by a conventional PMT. [57]

2.6 Software Framework

The MICE Analysis User Software (MAUS) is a bespoke software framework used by MICE and replaces the previously used G4MICE.

2.6.1 Software Elements

MAUS is intended as a full-featured package, designed to handle Monte Carlo (MC) simulation and digitisation, detector reconstruction for both MC and real data, as well as analysis:

Beam Generation

Simple beams can be created using a number of different configuration options to specify the constituent particle species, as well as momentum distributions. More complex beams can be created by using an external package such as G4Beamline to generate particles at the target and propagate them all the way to the beginning of the instrumented section of the beamline, at which point the output file can be read in by MAUS for further processing. It should be noted that measurements in MICE are mostly decoupled from the large uncertainties of current hadronic models, as particles in the beam are individually measured before and after the cooling channel, making the initial pion distribution at the target largely irrelevant.

Geometry

MICE geometries are stored in an online configuration database (CDB) from which specific geometries can be retrieved by run number or by a combination of geometry ID and tags to specify the magnet currents and cooling channel. The geometries themselves are generated from CAD drawings and survey data (with the detector geometries created manually), and are stored in Geometry Description Markup Language (GDML), a subset of XML developed specifically for particle physics geometries. Internally, the geometries are also converted into a MAUS-specific geometry format.

Simulation

Simulation of beams within MAUS is performed by GEANT4, which is included as a third party library. Particles are propagated individually through the beamline geometry and MC truth data stored for every propagation step within a detector volume and at virtual planes which can be defined in the geometry. Both the physics model for particle

interactions with beamline material and the physics processes particles will be subjected to during propagation are configured using MAUS settings.

Detector responses are simulated to convert energy deposition into an analogue signal equivalent to that produced by real data, taking into account physics processes occurring both in the optical material (photon yield, attenuation) and in the readout electronics.

Reconstruction

At the reconstruction step, detector output – whether from simulations or real data – is converted into space points (containing a subset of parameters defining the hit depending on the detector) and – in the case of Trackers and EMR – tracks made up of several points including momentum parameters derived from the properties of the track. Global track reconstruction – described in detail in chapter 3 – then determines which space points and tracks belong to the same particle and produces tracks spanning the entire instrumented beamline.

During data-taking, online reconstruction processes the raw data as it is being recorded and produces histograms to enable short-turnaround diagnostics of possible issues.

Analysis

Analysis on processed datasets can be performed either directly in MAUS or with external scripts which selectively load only the required MAUS components.

2.6.2 Framework Design

Data Structure

Data in MAUS is grouped into spills where for real data, one spill corresponds to one dip of the target. Depending on the type of data (real or simulated) and to what degree it has been processed, a spill can contain DAQ data or a number of MC events, as well as reconstructed events. Each of these then branches off into events for the different detectors, hits, coordinates, etc., as appropriate.

Modular Architecture

At its core, MAUS is made up of modules and uses a dataflow similar to the map-reduce [58] model. There are four types of modules in MAUS:

Inputters create spills, either from existing data in an input file, or empty spills to be filled by MAUS' internal beam generation.

Mappers modify spills on a spill-by-spill basis (allowing parallelised reconstruction). MC simulation, as well as local and global reconstruction are all performed via mappers.

Reducers act on a collection of spills to produce histograms and other aggregated data sets.

Outputters output processed data in the chosen data format.

Programming Languages

MAUS is written in Python and C++. Python provides the frontend (configuration files, executables, etc.) as well as some high-level algorithms, while the bulk of processing code is written in C++ to maximise performance.

Data Formats

Two main data formats are supported by MAUS: ROOT (binary) and JSON (ASCII). Conversion is handled in such a way that all references between objects are maintained. Originally, JSON was used to pass data between modules as well, though this was phased out in favour of the internal C++ data format in order to avoid costly string conversions.

2.6.3 Detector Reconstruction

Following is a description of the local reconstruction performed by the various detectors along the MICE beamline. A detailed description of global track reconstruction can be found in chapter 3.

TOFs

First, the TOF reconstruction evaluates slab hits and discards any hits that do not have a readout from the PMTs on both sides. For hits that pass this condition, readouts from both sides are averaged to determine the time and charge deposit associated with the hit. All possible combinations of hits in x and y slabs are then considered as space point candidates and calibration corrections used to determine which of these will pass as valid space points.

Cherenkov Detectors

For each charge deposition pulse, the Cherenkov reconstruction identifies the times t_1 and t_2 where the ADC value crosses a threshold, i.e. the approximate beginning and end times of a pulse, and then integrates the charge deposit for $t_1 - 8 \text{ ns} \leq t \leq t_2 + 16 \text{ ns}$. A photoelectron count is then obtained from the charge deposit by pedestal subtraction and normalisation.[59]

Trackers

Tracker reconstruction proceeds in five steps: First, ADC data is recorded and channel-by-channel calibrations applied to form digits, i.e. hits in individual tracker channels. Digits in neighbouring channels are then combined by a clustering algorithm and averaged to determine a plane coordinate in the corresponding tracker view (u , v , or w , see section 2.5.4).

Next, space points – with a z coordinate provided by the location of the tracker plane and the x and y coordinates calculated from the view coordinates – are reconstructed from clusters in the three views of the tracker. Initially, all possible triplet space points (i.e. formed from clusters in all three views) are determined by using the geometric fact that the sum of channel numbers for any triplet space point will be the same.² Remaining clusters are then formed into doublet space points (which can still provide full spatial coordinates, though at a lower precision).

Pattern recognition is then performed on the created spacepoints by fitting straight lines (in the absence of magnetic fields in the spectrometer solenoids) or helices to the space points and returning the track candidate with the smallest χ^2 . Pattern recognition will first only look for tracks with hits in all five tracker stations, followed by tracks with hits in only four tracker stations.

Finally, a track fit is performed using a custom Kalman fitter, taking into considerations physical processes such as energy loss and multiple Coulomb scattering.[60]

²This can be shown by writing in polar coordinates

$$u + v + w = r \left[\cos \phi + \cos \left(\frac{2\pi}{3} - \phi \right) + \cos \left(\frac{4\pi}{3} - \phi \right) \right] = 0$$

which implies that the sum of the fibre numbers corresponding to any point in the tracker plane is equal to the sum of the fibre numbers of the central fibres in the plane.

KL

KL reconstruction is similar to TOF reconstruction, though due to the KL only being segmented in the y coordinate, a spacepoint is created whenever there is a readout from the PMTs on both sides of a slab. During global reconstruction, if a trigger contains spacepoints in multiple slabs, adjacent spacepoints are combined with the y position of the resulting spacepoint being the average of the constituent points weighted by charge deposition.

EMR

Each hit in an EMR slab provides two spatial coordinates (yz or xz depending on whether the hit is in a horizontal or a vertical plane), given by the centre of gravity of the triangular slab which recorded the hit. A polynomial fit is used to generate tracks from the recorded hits and some additional information is extracted from the hit distribution, most notably the *plane density* which shows the proportion of planes along the path of a particle that contain hits. It is a strong PID variable for distinguishing between electrons and muons, as muons will produce plane densities near one, whereas the shower nature of electron interactions result in significantly lower plane densities.

2.7 Outlook

MICE Step IV has been consistently producing large amounts of data for over two years now. While the most crucial part of the experiment is evaluating the behaviour of the beam inside an absorber, to fulfil its physics goals of demonstrating sustainable muon ionisation cooling, it is necessary to include re-acceleration in the MICE beamline. Two beamline configurations with different associated costs, risks, and performance have been proposed and funding is currently being sought.

Chapter 3

Global Track Reconstruction

“If debugging is the process of removing bugs, then programming must be the process of putting them in.”

Edsger W. Dijkstra

A high precision – $\mathcal{O}(10\%)$ – measurement of a change in beam emittance, which isn’t possible with conventional bulk measurement techniques, requires MICE to operate as a single particle experiment, i.e. each trigger should contain detector hits from only a single particle. This provides significantly more detailed information about the particle behaviour in the beamline, but comes at the cost of having to ensure that detector hits are only produced by a single particle. Since the stochastic nature of the beam results in a nonzero chance of multiple particles within a trigger window, tracks from multiple particles have to be separated from each other before identification and fitting can be performed. Global track reconstruction thus consists of three steps, track matching (section 3.2), particle identification (section 3.3), and track fitting (section 3.4).

3.1 Global Data Structure

As every detector provides a different set of measurements, each detector group devised their own data structure in order to represent the various data provided by the detector and/or generated by the local reconstruction routines. For producing global tracks, it then becomes necessary to insert the locally reconstructed data into a uniform data structure that can contain all the information relevant for analysis. The following is a brief overview of this data structure. Details can be found in the MAUS documentation [61].

Track Points correspond to the state of a particle at a specific location in the beamline. There will typically be one track point for each of TOF0-2, KL, and the Cherenkovs. A track point contains four-vectors for position and momentum, as well as which detector it belongs to and some detector-specific information (e.g. charge deposition).

Tracks are collections of several track points together with some additional information that applies to the track, such as PID, range in the EMR, or the p-value from the tracker fit. They can also contain references to constituent tracks for book-keeping purposes (e.g. an upstream matched track will reference the upstream tracker track whose track points it contains).

Primary Chains were set up to allow tracing back the steps of global track reconstruction, i.e. looking at the state of tracks and track points at every stage of global reconstruction. They contain all the information that is associated with an individual particle in the beamline, including the tracks that are produced by track matching, PID, and track fitting.

3.2 Track Matching

Track matching determines whether hits in different detectors come from the same particle by propagating a particle's position-momentum vector $(t, x, y, z, E, p_x, p_y, p_z)$ between detectors and checking for agreement (in most cases in x and y). Matching tolerances depend on the respective detector granularity, as well as the propagation distance and material in the path, to account for multiple scattering and the uncertainty on the original measurement that is propagated. The MICE trackers (see section 2.5.4) are chosen as the detectors to propagate from, as they are the most indispensable for a cooling measurement (only they can measure the emittance) and provide the most complete, accurate, and precise subset of the particle's phase-space coordinates – namely (x, y, z, p_x, p_y, p_z) , including the finest granularity in position. Propagation is performed outward from the outermost tracker plane – i.e. upstream from the most upstream plane of the upstream tracker and downstream from the most downstream plane of the downstream tracker – towards the detectors on each side. Thus, separate upstream and downstream tracks are obtained which are then matched via a simple time-of-flight cut. Since at the time of track matching the particle ID has not been determined yet, matching is performed separately for all feasible particle hypotheses (typically e, μ, π with the charge corresponding to the run's beamline polarity) and

the correct track is subsequently picked out by the PID algorithm (section 3.3). An illustration of track matching can be seen in figure 3.2, a listing of the default tolerances for matching in table 3.1.

3.2.1 Propagation Method

Propagation itself is performed using a 4th order Runge-Kutta (RK4) numerical integration method. Runge-Kutta methods [62] calculate a step using a weighted average of increments along the step and the 4th order – also called the classical Runge-Kutta method – is the most widely used of those, as it is the highest order one where the required number of evaluations per step does not exceed its order. For an initial value problem specified as

$$\dot{y} = f(t, y), \quad y(t_0) = y_0, \quad (3.1)$$

where t is the independent coordinate, i.e. the coordinate along which stepping occurs, y is the dependent coordinate, and \dot{y} its derivative with respect to t ; with a step-size of h , the value of y at step $n + 1$ is defined as

$$y_{n+1} = y_n + \frac{h}{6} (k_1 + 2k_2 + 2k_3 + k_4) \quad (3.2)$$

where $k_{1,2,3,4}$ are increments based on the slope of the function at the beginning, twice at the midpoint, and at the end of the interval, given by

$$\begin{aligned} k_1 &= f(t_n, y_n), \\ k_2 &= f\left(t_n + \frac{h}{2}, y_n + \frac{h}{2}k_1\right), \\ k_3 &= f\left(t_n + \frac{h}{2}, y_n + \frac{h}{2}k_2\right), \\ k_4 &= f(t_n + h, y_n + hk_3), \end{aligned} \quad (3.3)$$

Track matching uses an implementation of the RK4 method from the GNU Scientific Library (GSL). Electric and magnetic field values are obtained at every step directly from the defined geometry and included in the equations of motion passed to the integration method. Stepping occurs with respect to the z coordinate, with the equations of motion (i.e. $\dot{y} = f(t, y)$ in the RK4 algorithm) given by

$$\begin{aligned}
\frac{dt}{dz} &= \frac{E}{p_z c} & \frac{dx}{dz} &= \frac{p_x}{p_z} & \frac{dy}{dz} &= \frac{p_y}{p_z} & \frac{dz}{dz} &= 1 \\
\frac{dE}{dz} &= qE_x \frac{dx}{dz} + qE_y \frac{dy}{dz} + qE_z \frac{dz}{dz} \\
\frac{dp_x}{dz} &= qc \left(\frac{dy}{dz} B_z - \frac{dz}{dz} B_y \right) + qE_x \frac{dt}{dz} \\
\frac{dp_y}{dz} &= qc \left(\frac{dz}{dz} B_x - \frac{dx}{dz} B_z \right) + qE_y \frac{dt}{dz} \\
\frac{dp_z}{dz} &= qc \left(\frac{dx}{dz} B_y - \frac{dy}{dz} B_x \right) + qE_z \frac{dt}{dz}
\end{aligned} \tag{3.4}$$

where q and c are the particle's charge and the speed of light, respectively. For MICE Step IV, which does not include RF cavities, there are no electric fields in the beamline which simplifies the bottom 4 equations, noticeably yielding $\frac{dE}{dz} = 0$, though energy loss from interactions with materials is treated separately. At each step, the energy loss in the material is calculated from the Bethe-Bloch equation[5]

$$-\left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \tag{3.5}$$

where $K = 4\pi N_A r_e^2 m_e c^2 = 0.307075 \text{ MeV g}^{-1} \text{ cm}^2$, N_A is Avogadro's number, m_e and r_e are the electron mass and classical electron radius, Z and A are atomic and mass numbers of the traversed material, I is the material's mean excitation energy in eV, $\delta(\beta\gamma)$ is the material's density correction (dependent on $\beta\gamma$), and T_{\max} is the maximum kinetic energy that can be imparted to a free electron in a single collision, defined by

$$T_{\max} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e/M + (m_e/M)^2} \tag{3.6}$$

where M is the mass of the propagated particle.

The inclusion of energy loss requires care when approaching material boundaries – as the material is only sampled once in the middle of each step – as well as when traversing materials with high stopping power – as a large step in a strong magnetic field would result in underestimating the step distance due to the discrepancy between linear and curved distance between start and end point (see figure 3.1). The step-size behaviour in such situations was fine-tuned to ensure that errors from the propagation algorithm itself remain negligible while minimising the number of steps – and therefore the computing time – required.

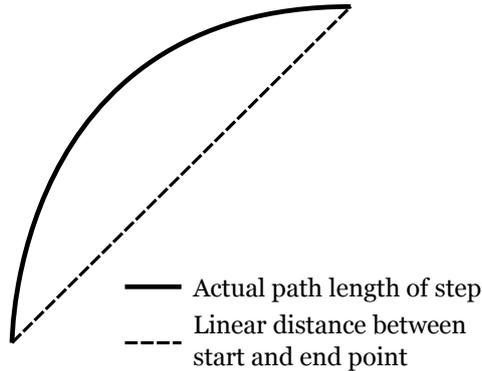


Figure 3.1: Illustration of the difference between true path length and distance between start end end point of a step for curved tracks

3.2.2 TOF1 & TOF2

For TOF1 and TOF2, the tracker track is propagated to the centre (in z) of the respective detector and matched by agreement in the x and y coordinates. Δt between trackers and TOFs cannot be used for this, as the trackers currently do not produce useful timing information. Successful matching is very important here, as both are required in order to match tracks through the cooling channel (section 3.2.7) and a TOF1 match is also required to match to TOF0 (section 3.2.3).

3.2.3 TOF0

Due to the distance and beam optics between TOF0 and TOF1, TOF0 cannot be matched in the same way as TOF1. The uncertainties on the phase-space coordinates from the tracker expand to such a degree during propagation that suitable matching allowances would include a majority of the detector, and furthermore the propagation distance would cause a significant performance penalty. Instead, the particle is propagated to the upstream end of TOF1 and an approximate expected travel time between TOF0 and TOF1 is calculated and compared to the actual difference between the timestamps at the two detectors. To account for energy loss in the air between the two detectors, the travel distance is estimated as $\frac{p}{p_z} \times d$ (where p is the total momentum, p_z the momentum in beamline direction, and d the z distance between the detectors) and energy loss for half that thickness of air calculated (to obtain an average energy between the two detectors). The material in the Cherenkov detectors is not taken into account as it is located very close to TOF0 and therefore – when propagating backwards from

TOF1 – has a very small impact on total travel time. The expected average velocity between the two detectors is then calculated as $\frac{E}{p_z} \times c$, where E is calculated from the momentum reconstructed in the tracker and the mass of the particle type for which matching is performed.

3.2.4 Cherenkov Detectors

As the Cherenkov detectors do not have sufficient timing resolution (each signal window is tens of nanoseconds long), as well as no spatial segmentation to separate multiple hits within a trigger, they never produce more than a single point each. Hence, these are added to tracks without checks, which can result in (unavoidable) errors in the Cherenkov photoelectron count associated with a track if the photoelectrons originated from multiple particles.

3.2.5 KL

Track matching to the KL detector works just like propagation to TOF1 and TOF2 (section 3.2.2), except that as the detector is only segmented in the y coordinate, track matching only compares positions in a single spatial dimension.

3.2.6 EMR

The EMR produces tracks, so matching is performed with the most upstream point of the primary track produced. While the local reconstruction of the EMR produces variable position uncertainties, matching is still performed against a fixed threshold, as uncertainties in the tracker momenta as well as scattering in the lead of the KL dominate over the relatively fine detector granularity.

3.2.7 Through the Absorber

Matching through the absorber is performed by a cut on the travel time between TOF1 and TOF2, so it can only be done if both detectors have been successfully matched to their respective tracker tracks. For the purposes of an emittance measurement, the absorber must essentially be treated as a black box, therefore it is important to not make any assumptions about the physical processes within during track matching. Hence, the cut only verifies that a particle passing through both detectors would have travelled subluminally but not significantly slower than would be realistic in the beamline.

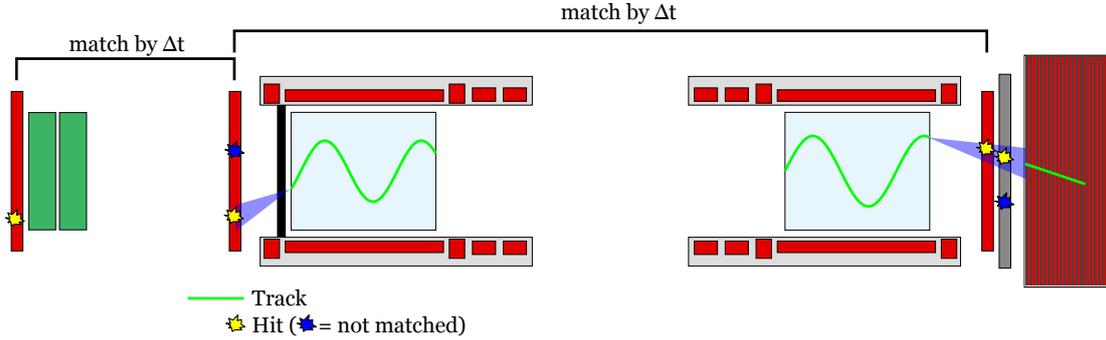


Figure 3.2: Illustration of the track matching process. The pictured detectors are (from left to right): TOF0, Cherenkovs, TOF1, Upstream Tracker, Downstream Tracker, TOF2, KL, EMR.

Detector	Matched by	Detector Granularity	Tolerance
TOF0	t	negligible*	2 ns
TOF1	x, y	60 mm [†]	60 mm
TOF2	x, y	60 mm [†]	50 mm
KL	y	44 mm [†]	50 mm
EMR	x, y	varies [‡]	75 mm
US/DS	β_z	—	$0.5 \leq \beta_z \leq 1.0$

Table 3.1: Default tolerances for track matching (can be modified via configuration datacards)

If matching through the absorber fails (or isn't possible e.g. due to failed TOF1 or TOF2 matching), upstream and downstream tracks are still stored for analysis. Furthermore, if multiple upstream tracks are matched to a single downstream track (or vice versa), the resulting tracks are marked as being mutually exclusive so that the analyst can choose whether to remove them from the sample or subject them to further checks.

3.2.8 Performance

The performance of track matching is influenced by a range of factors:

- Precision of the algorithm itself

*The TOFs have timing resolutions of ~ 50 ps but this is negligible compared to the uncertainties arising from the matching method (see section 3.2.3)

[†]Granularity given is the slab width, i.e. twice the maximum difference between real and reconstructed hit assuming perfect reconstruction

[‡]Uncertainties on EMR track points are calculated from a track fit but are of secondary importance for matching (see section 3.2.6)

- Non-deterministic processes, such as multiple scattering (which is not accounted for in the propagation algorithm at all) and energy loss (which can only be accounted for in terms of the mean energy loss)
- Quality of the local reconstruction in all involved detectors
- Accuracy of the geometry files, especially of the magnetic field maps
- Particle decays (i.e. $\pi^+ \rightarrow \mu^+ + \nu_\mu$ and $\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$)

Additionally, there is an interdependence between efficiencies of some of the detectors:

- No matching can occur without a tracker track on the corresponding side (upstream or downstream) of the beamline, as hits are always matched to a tracker track (with the exception of TOF0, see above)
- TOF0 matching can only occur with TOF1 successfully matched
- Matching through the absorber can only occur with TOF1 and TOF2 successfully matched

While the strong sensitivity to the accuracy of the virtual geometry is a potential source for errors, it also makes track matching a valuable tool for debugging purposes, as residual plots (as seen in figures 3.5–3.10) produced by this module have helped identify a number of problems including geometry misalignments and rotations in the tracker reconstruction.

Due to the aforementioned factors, some of which will differ from run to run, matching tolerances (default values for which are given in table 3.1) are not hardcoded, but can be varied as required. This also means that figures for the efficiency of the algorithm are merely illustrative of the performance of the matching algorithm, rather than figures that would factor directly into measurement results. For run 8373, the default tolerances yield an efficiency of 98% for matching through the absorber (which requires prior matching to TOF1 and TOF2) running over production MC data, which should be representative of high quality runs at 170 MeV/c. Specifically, this means that 98% of particles that produced tracks and space points in the local reconstruction of both trackers as well as TOF1 and TOF2 had at least those two TOFs correctly matched and also succeeded in through-matching. Since the particle ID needs to be known in order to determine whether matching should have occurred, efficiency can only be calculated for MC data.

Propagation Verification Residuals

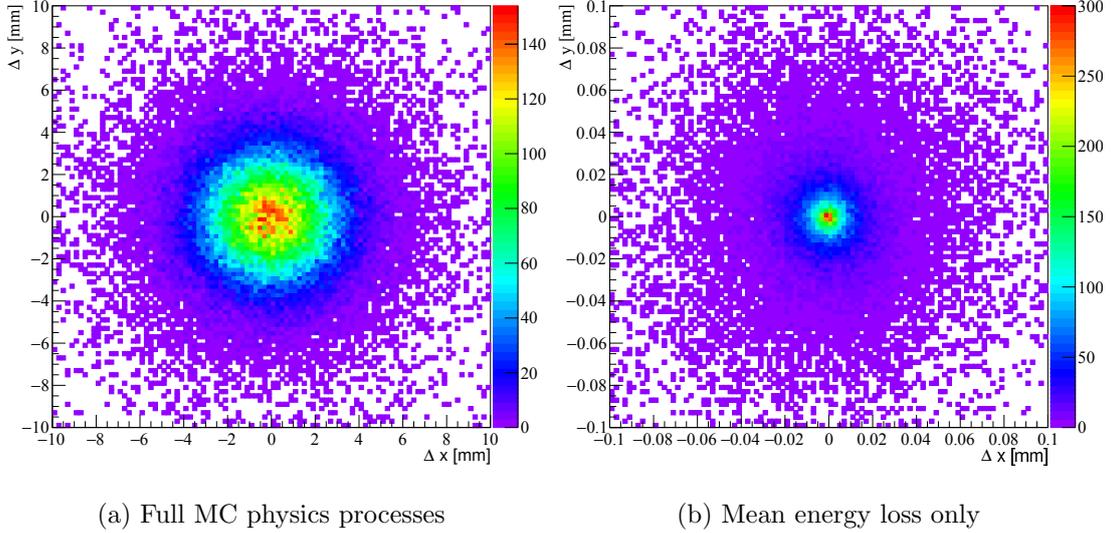


Figure 3.3: Position residuals for MC truth propagation in the downstream tracker

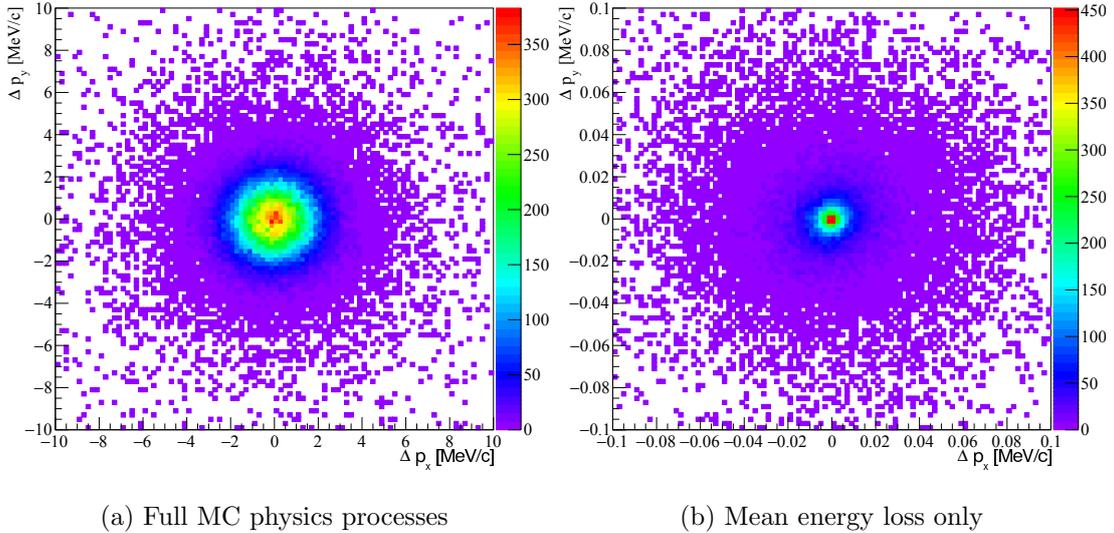


Figure 3.4: Transverse momentum residuals for MC truth propagation in the downstream tracker

The propagation method itself can be verified by propagating MC truth rather than reconstructed data, thereby eliminating several of the factors listed above, most notably the local reconstruction and the accuracy of the geometry files. Figures 3.3 and 3.4 show

residuals for propagation of MC truth data from one end of the upstream tracker to the other. The residuals measure the difference between propagated and MC truth (in the following section between propagated and reconstructed) coordinates. The left hand figures, which are generated from production MC (i.e. the “official” simulation used in analyses), show a fairly small but still quite noticeable residual spread, which is primarily driven by scattering and the stochastic nature of energy loss. The figures on the right are generated from MC where interaction with materials is restricted to mean energy loss only, and as would be expected (assuming that the propagation implementation is physically accurate), the residuals shrink to almost zero (tens of microns and tens of keV).

Track matching Residuals

Figures 3.5–3.10 show track matching residuals for run 8373 in both MC and data with red lines denoting the matching tolerances. Note that the MC residuals were produced with a PID cut, so that only muons are contained in them whereas the data residuals were not, so there is contamination from non-muons present (as can be seen for example in the pion peak in figure 3.5b), hence a larger spread is expected. It should be noted that for this run, there is a severe discrepancy between MC and data for the KL and even more so for the EMR, which suggests that there are still (as of this writing) noticeable problems with the downstream section of the geometry used, likely in the downstream fringe fields of the downstream tracker.

3.3 Particle Identification (PID)

Particle Identification was implemented by Celeste Pidcott and is described in detail in [63]. It is PDF-based, meaning that it works by running Monte-Carlo simulations of the beam in question and generating probability density functions (PDFs) for each particle species for a number of different PID variables (such as TOF0 to TOF1 time-of-flight or plane density in the EMR). For each particle in data, the same PID variables are then calculated, and log-likelihoods (LL_X) assigned by comparing the values to the PDFs. A PID is then assigned to the track if the confidence level (CL_X) for a certain PID,

$$CL_X = \frac{\exp LL_X}{\sum_i \exp LL_i} \times 100 \quad (3.7)$$

exceeds a configurable confidence level threshold. Since global track matching produces separate tracks for each PID hypothesis, a track is then only passed on as being success-

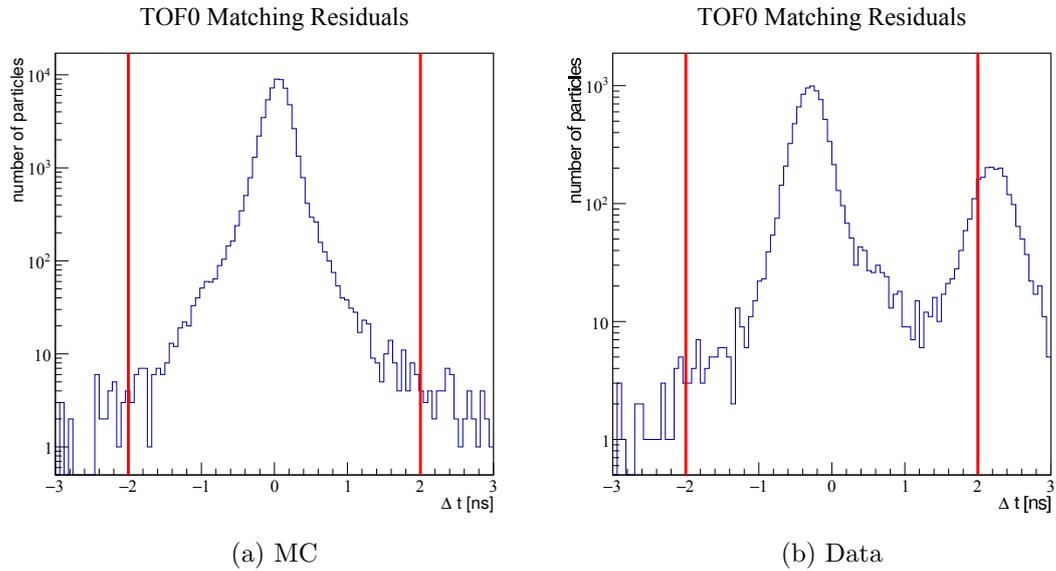


Figure 3.5: Track matching residuals for TOF0. The secondary peak in the figure for data corresponds to a pion population which in MC was filtered out.

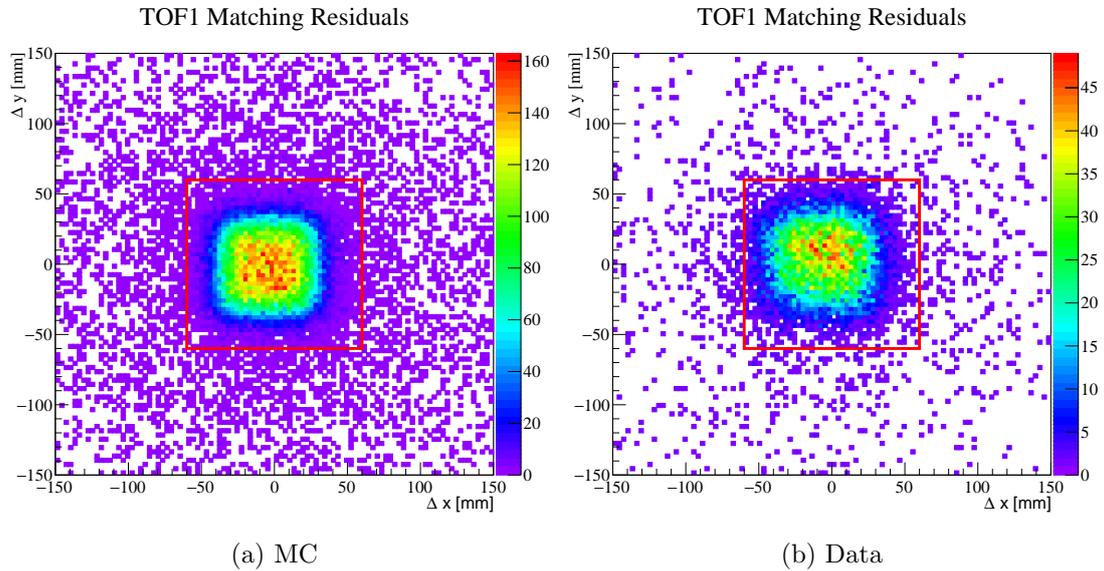


Figure 3.6: Track matching residuals for TOF1. Note the different z scale due to higher statistics in the MC.

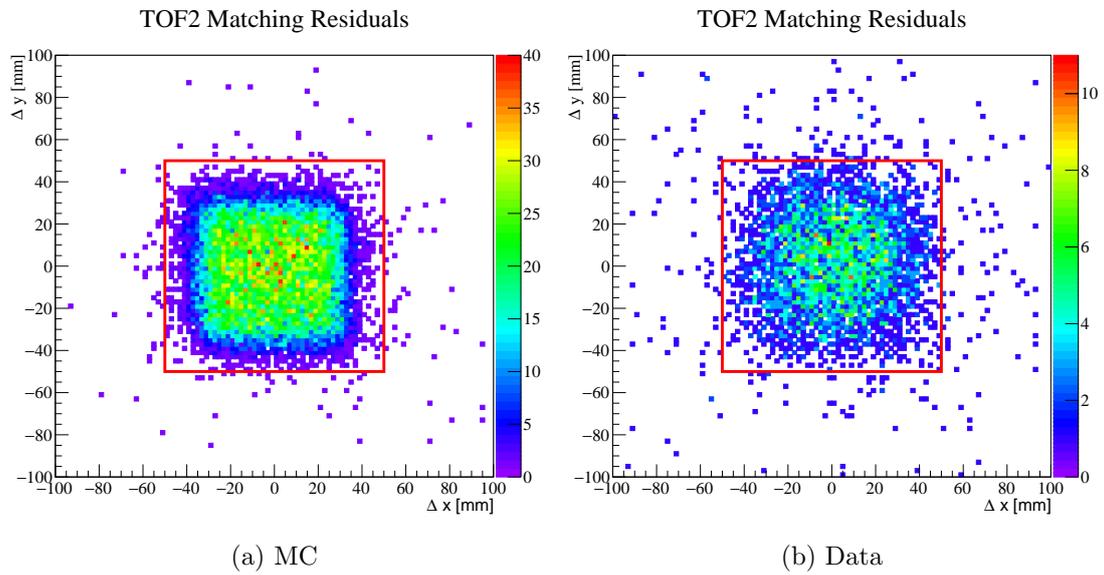


Figure 3.7: Track matching residuals for TOF2

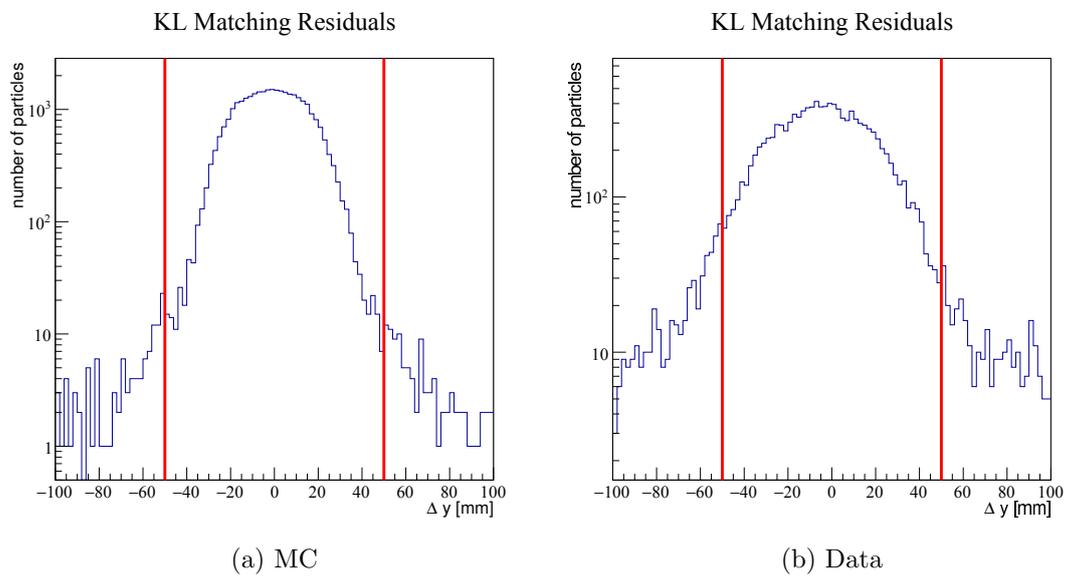


Figure 3.8: Track matching residuals for the KL

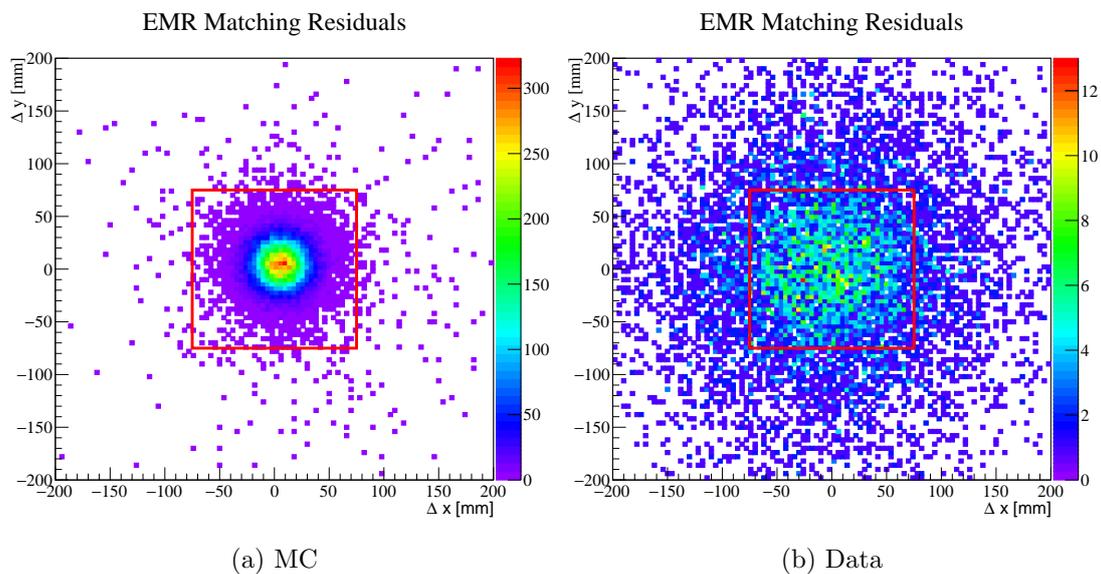


Figure 3.9: Track matching residuals for the EMR

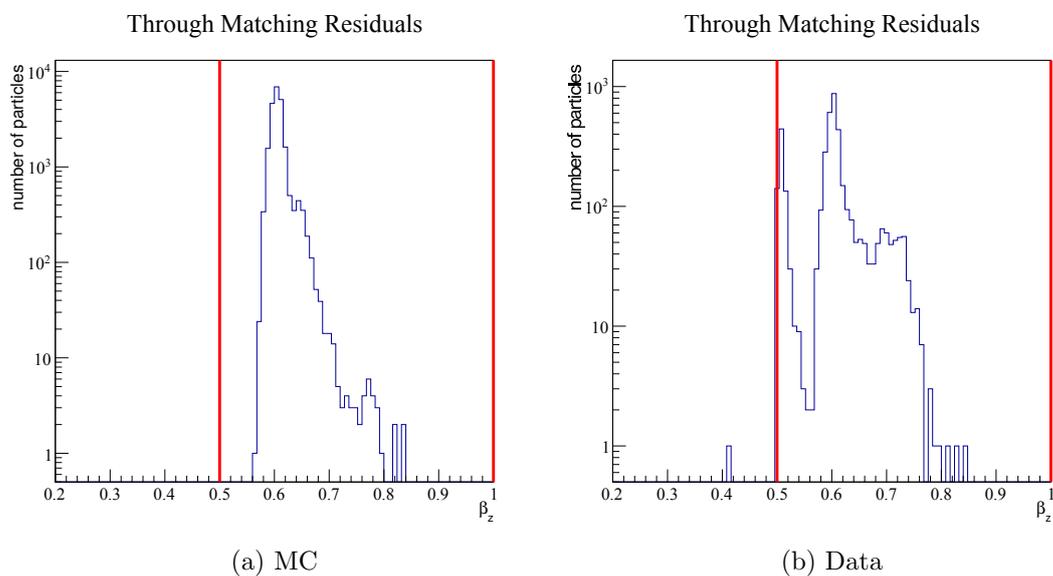


Figure 3.10: Track matching residuals through the absorber. In data, the electron and pion peaks can be seen.

fully PID'd, if the assigned PID matches the PID hypothesis provided by track matching. The PID algorithm was not used in the analysis in this thesis, as the low-level nature of the analysis made a cut-based PID (which sacrifices efficiency but reduces overhead and makes significantly fewer model assumptions) more appropriate.

3.3.1 PID Variables

A wide range of PID variables was implemented, probing the various detectors as well as some combinations of detectors. These variables are grouped into two sets, one for field-on (i.e curved) tracks, and one for field-off (straight) tracks, as the available information that is useful to particle identification differs between the two cases. The field-off PID was primarily intended for commissioning, when the magnets were not yet trained and no absorber was present, and was thus not subject to the constraint that particle identification should be independent between upstream and downstream parts of the beamline. Therefore, PID variables that cross the cooling channel were implemented to compensate for the lack of tracker momentum measurement. A list of the PID variables, as well as their associated efficiencies and purities can be found in [63].

3.4 Track Fitting

Track fitting is currently in development and is being implemented by Chris Rogers. It is based on the Kalman filter that was originally developed for the local tracker reconstruction (see section 2.6.3). A Kalman filter can improve the accuracy and precision of data points in a set by combining the measurement (incl. uncertainties) with external constraints (such as the laws of physics) into a weighted average. The global track fitter uses a similar propagation method as the one used in track matching, with the main difference being that it propagates not only the position and momentum, but an entire covariance matrix in order to determine the weighting that should be used. This allows the weighting to be influenced not only by the uncertainties in the measurements themselves, but also by how for instance passage through materials will increase uncertainties in position and momentum due to random scatters, with a prediction that has been propagated through significant amounts of material given a lower weight. While track fitting is not yet fully implemented in MAUS, the global data structure (section 3.1) has been set up with a later inclusion of track fitting in mind.

3.5 Conclusion

Global track reconstruction combines and processes the output from the various MICE detectors and provides a unified data structure for high level analysis. The global data structure as well as track matching were implemented by the author and designed in such a way that PID and track fitting can be included in global processing in a modular manner. Track matching shows excellent performance (subject to an accurate virtual geometry) as outlined in section 3.2.8.

Chapter 4

Multiple Scattering in the MICE Absorbers

“Research is so unpredictable. There are periods when nothing works and all your experiments are a disaster and all your hypotheses are wrong.”

Francis Collins

4.1 Physics of Multiple Scattering

When passing through a medium, charged particles will experience deflection due to scattering off the Coulomb potentials of nuclei in the medium. Scattering off electrons does also occur, though it only forms a considerable contribution when the mass of the incident particle and the atomic number Z of the target material are low. Accurate modelling of the potential is one of the challenges in devising scattering models and forms one of the key distinctions between scattering models, for example in whether and in what way nuclear screening is considered. Over the last century, many models have been devised, sometimes building on previous models, other times approaching the derivation from quite different angles (no pun intended).

Treatment of scattering can fall into three different domains, *single*, *plural*, and *multiple* scattering:

4.1.1 Single Scattering

Single scattering describes the behaviour for very thin materials (i.e. with a thickness much less than the interaction length), where the probability of more than one scattering

event is negligible. It is a random process that can be described by a scattering cross section, which was first derived by Rutherford[64] as

$$\frac{d\sigma}{d\Omega} = \left(\frac{zZe^2}{8\pi\epsilon_0mv_0^2} \right)^2 \csc^4 \frac{\theta}{2}, \quad (4.1)$$

where z and Z are the charge numbers of the incident and target particle, m and v_0 are the mass and initial speed of the incident particle, and θ is the scattering angle. The Rutherford cross section breaks down, however, at extremely small angles, where it tends to infinity. In 1926, Wentzel[65] published a cross section which includes nuclear screening by the orbiting electrons:

$$\frac{d\sigma}{d\Omega} = \left(\frac{zZe^2}{8\pi\epsilon_0mv_0^2} \right)^2 \left(\sin^2 \frac{\theta}{2} + \frac{1}{4k^2a^2} \right)^{-2}, \quad (4.2)$$

where k is the incident particle's wave number and $a = 0.885a_0Z^{-\frac{1}{3}}$ (where a_0 is the Bohr radius) is the Thomas-Fermi radius – the radius of the electron cloud effectively screening the nucleus. It can be seen that for all but the smallest angles – where the term involving a becomes comparable to the scattering angle – this reduces to the Rutherford cross section.

4.1.2 Plural and Multiple Scattering

Plural scattering applies when the number of interactions in the target material is larger than one but not high enough for the effects to become stochastic, whereas multiple scattering deals with large numbers of interactions. The exact boundary between plural and multiple scattering is somewhat arbitrary, typically taken somewhere around 30 scattering events, but the stochastic nature of multiple scattering makes it feasible to perform significant simplifications. Specifically it is usually sufficient to approximate the central 98 % of the angular distribution as a Gaussian distribution with an RMS width of

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta cp} z \sqrt{x/X_0} [1 + 0.038 \ln(x/X_0)] \quad (4.3)$$

where z is the charge number of the incident particle and x/X_0 is the thickness of the traversed medium in radiation lengths. The equation comes from a fit to the Molière theory – which is also the origin of the 13.6 MeV term in the ionisation cooling equation (equation 2.9) – and according to the Particle Data Group is accurate to within less than

11 % between 0.001 and 100 radiation lengths.[5] It should be noted that the origins of this equation – while it is widely used and cited – are not well-known. A discussion of this can be found in section 5.3 of [66].

4.1.3 The Molière Theory

Molière’s multiple scattering theory [67] forms the basis of many of the scattering implementations used today. Entirely analytical in nature, it reduces the description of scattering to a single parameter, the screening angle, defined as

$$\chi_a = \frac{\lambda}{a} \sqrt{1.13 + 3.76a^2} \quad (4.4)$$

where λ is the particle’s de Broglie wavelength and a is the Thomas-Fermi radius (see section 4.1.1).

The thickness of the scattering medium is described by the unit probability angle,

$$\chi_c = \sqrt{4\pi N t e^4 Z^2 z^2} / p v \quad (4.5)$$

where N is the material’s atomic number density, t is the product of thickness and density of the material, and p and v are the momentum and speed of the incident particle. Physically, it describes the angle for which there is (on average) only a single collision with scattering angle $\chi > \chi_c$.

The angular scattering distribution is then only a function of the ratio of χ_c and χ_a . Specifically (following Bethe’s description[68]), by defining

$$e^b = \frac{\chi_c^2}{1.167\chi_a^2}, \quad (4.6)$$

a parameter B given by the transcendental equation

$$B - \ln B = b, \quad (4.7)$$

and

$$\vartheta = \frac{\theta}{\chi_c B^{\frac{1}{2}}}, \quad (4.8)$$

the distribution can be written as a power series:

$$f(\theta)d\theta = \vartheta d\vartheta \left[f^{(0)}(\vartheta) + B^{-1} f^{(1)}(\vartheta) + B^{-2} f^{(2)}(\vartheta) + \dots \right], \quad (4.9)$$

where $f^{(n)}(\vartheta)$ are functions of ϑ .

The original formulation of the Molière theory ignored scattering off atomic electrons, which for low- Z materials becomes a significant factor. Bethe[68] attempted to correct this by replacing Z^2 with $Z(Z+1)$ in the equation for χ_c . This was criticized by Fano[69] for the implicit incorrect assumption that the low-angle cutoff is identical for elastic and inelastic collisions. Fano instead suggested only using a factor of $Z(Z+1)$ in the case of incident electrons, as well as some smaller tweaks to better account for the the inelastic scattering component.

4.1.4 Other Scattering Theories

The Lewis theory[70] uses a similar formalism as Molière scattering but avoids the small-angle approximation inherent in the latter. It forms the basis of the standard multiple scattering algorithm in GEANT4[71], though details of the implementation beyond the model-agnostic description in section 4.1.5 are not documented.

An alternative scattering model, devised by Goudsmit and Saunderson[72], has recently been added to GEANT4, but is computationally more intensive and currently only implemented for electrons and positrons[73].

4.1.5 Algorithms

The following is a brief description of the multiple scattering algorithm implemented in GEANT4 [71].

Multiple scattering algorithms can be classified depending on whether they simulate collisions one by one, or consider the total effect after a step, i.e. whether they are *detailed* or *condensed*. A hybrid between the two, *mixed* algorithms, simulate large angle scatters individually but use a condensed method for the bulk of scattering events, since for simulating scattering in anything other than thin foils or gases, a detailed treatment tends to be computationally unfeasible. The current default scattering algorithm for muons in GEANT4 as of version 9.3.0 – `WentzelVI` – is mixed, whereas previous versions used the

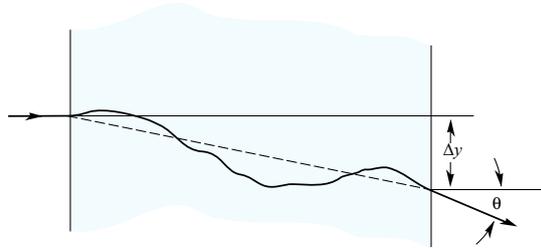


Figure 4.1: Illustration of the bulk effect of multiple single scatters in a material (modified from [5]). Many individual scatters result in a total scattering angle θ and a displacement in where the particle exits the material of Δy .

condensed **Urban** algorithm.

GEANT4 uses transport mean free paths, which are moments of the scattering angle χ , to calculate the scattering parameters. The k -th transport mean free path, λ_k , is defined as

$$\frac{1}{\lambda_k} \equiv 2\pi n_a \int_{-1}^1 [1 - P_k(\cos \chi)] \frac{d\sigma(\chi)}{d\Omega} d(\cos \chi) \quad (4.10)$$

where n_a is the atom number density, $P_k(\cos \chi)$ is the k -th Legendre polynomial, and $d\sigma(\chi)/d\Omega$ is the scattering differential cross section. The first two transport mean free paths are involved in GEANT4's scattering calculations, though λ_2 only as the parameter $\kappa = \lambda_1/\lambda_2$, which – since it varies extremely slowly with energy – is assigned a constant value of 2.5. The inverse of the first transport mean free path is – in analogy with stopping power – sometimes referred to as the scattering power of a material.

As illustrated by figure 4.1, there is a difference between the geometrical path length (which is either a straight or curved line, depending on the presence or absence of magnetic fields), indicated by a dashed line, and the true path length, indicated by a solid line, of a step as the particle traverses the material. The conversion between a given true path length t and a corresponding mean geometrical path length, $\langle z \rangle$, is given by

$$\langle z \rangle = \lambda_1 \left[1 - \exp\left(-\frac{t}{\lambda_1}\right) \right] \quad (4.11)$$

Back-and-forth conversions between $\langle z \rangle$ and t have to be performed to determine the step length in GEANT4, since for physics processes the true path length is relevant whereas the transportation occurs in terms of the geometrical path length.

The mean scattering angle θ after a step of true step length t is given by

$$\langle \cos \theta \rangle = \exp\left[-\frac{t}{\lambda_1}\right] \quad (4.12)$$

and its variance as

$$\sigma^2 = \langle \cos^2 \theta \rangle - \langle \cos \theta \rangle^2 = \frac{1 + 2e^{-2\kappa\tau}}{3} - e^{-2\tau} \quad (4.13)$$

where $\tau = t/\lambda_1$ and κ as described above. The angle after the scattering step is then sampled from a model distribution constructed such that equations 4.12 and 4.13 are satisfied.

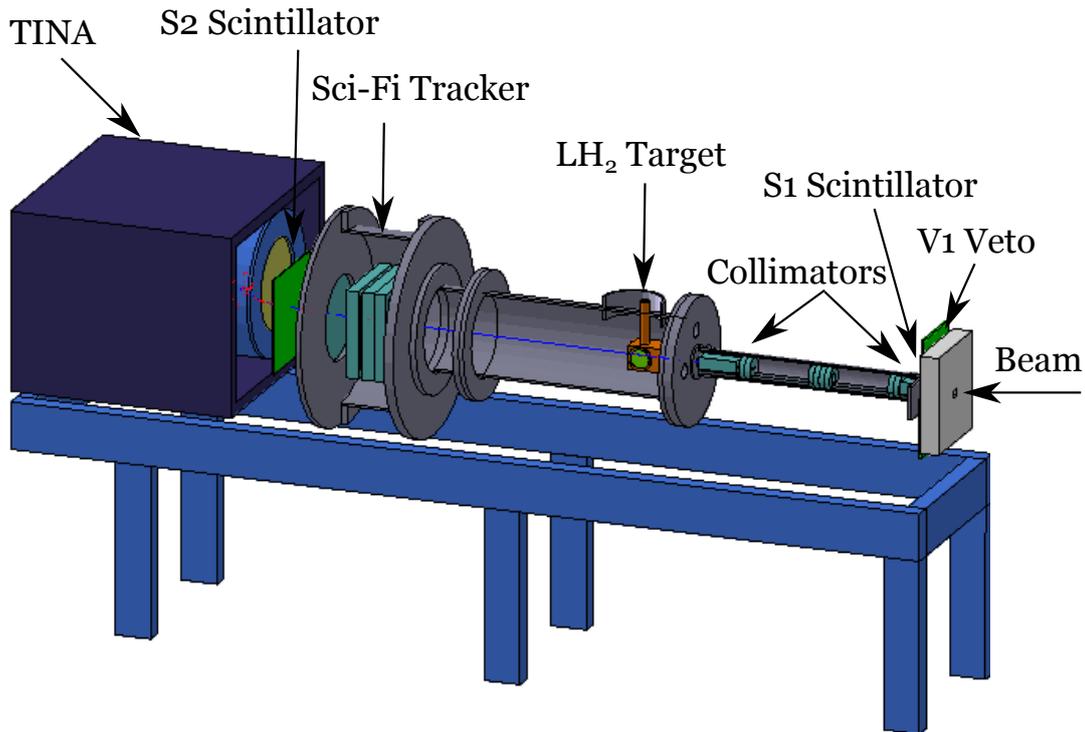


Figure 4.2: The MuScat detector[74]

A number of multiple scattering theories and models have been developed that aim to reduce the (in many cases prohibitively) immense computational burden that would arise from simulating every scattering event individually for particles passing through matter. While generally fairly successful, algorithms have – at least until recently – failed to accurately describe scattering in low- Z materials (see section 4.2). Thus experimental verification remains a crucial factor in the continued development of scattering simulations.

4.2 Previous Scattering Measurements

MuScat[74] was an experiment located at TRIUMF in British Columbia, designed to characterise multiple scattering of moderate momentum ($p_{\text{mean}} = 172 \pm 2.0 \text{ MeV}/c$) muons in several different (mostly low- Z) materials. Scattering distributions were measured by collimating a muon beam before incidence on the target and then measuring the deflection in a finely segmented scintillating fibre tracker on the target’s far side. The target system consists of a cryogenic liquid hydrogen vessel that could be inserted

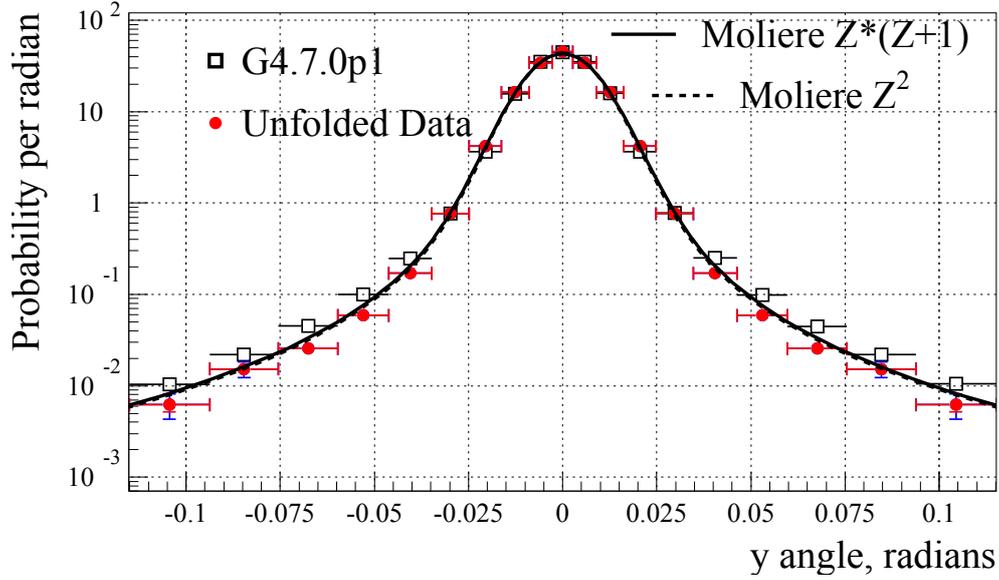


Figure 4.3: MuScat scattering distribution for 0.24 mm of iron [74]. Red error bars include statistical errors only, blue error bars indicate statistical plus systematic errors.

in two different orientations to present either 109 or 159 mm of target material to the beam, as well as a wheel that could be rotated to place any of 12 target disks into the beam. The target materials were lithium (6.40 mm, 6.43 mm, 12.72 mm, 12.78 mm), beryllium (0.98 mm, 3.73 mm), polyethylene (4.74 mm), carbon (2.50 mm), aluminium (1.5 mm), and iron (0.24 mm, 5.05 mm), which together with an empty slot in the disc as well as the possibility of running with an empty liquid hydrogen vessel (again in two orientations) provided a total of 16 different targets.

Examples of the scattering distributions found by MuScat for 0.24 mm of iron and 159 mm of liquid hydrogen can be seen in figures 4.3 and 4.4. The figures show unfolded data from measurements with two curves superimposed based on the Molière theory, one using a Z dependence of Z^2 , the other of $Z \times (Z + 1)$. This is compared with simulated data points from GEANT4 (G4.7.0p1) and – in the case of liquid hydrogen – from the ELMS model (a recent scattering model derived from photo-absorption spectra[75]). For iron, where the difference between the Z terms is $< 4\%$, the curves are practically identical and there is good agreement between theory, simulation, and data, though a small excess at large angles is visible in the GEANT4 simulation. For liquid hydrogen however, the difference between the two Molière curves is quite pronounced. The data follows the Z^2 line for most of the large angle ranges which are dominated by large single scatters and the $Z \times (Z + 1)$ line for the centre of the distribution which is dominated by

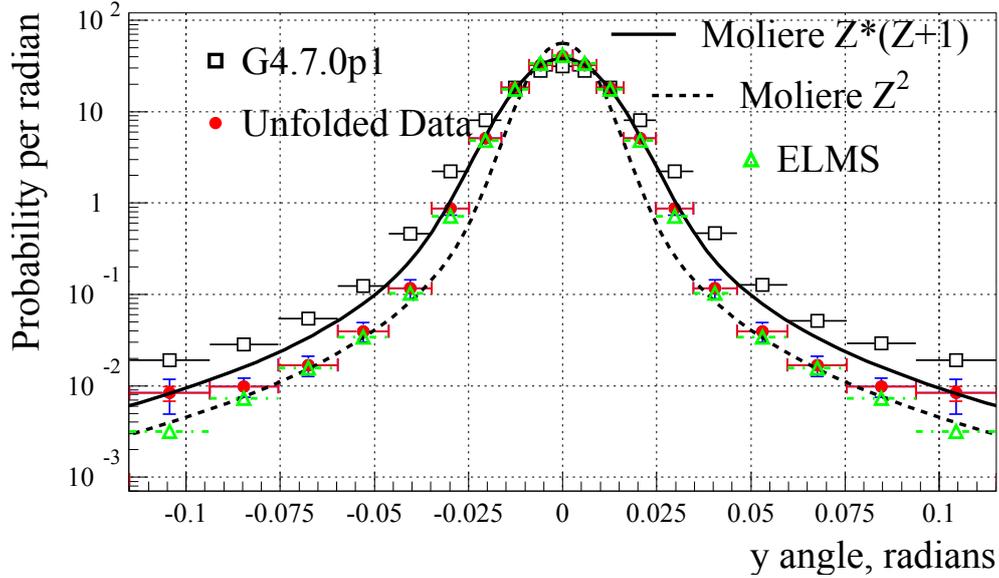


Figure 4.4: MuScat scattering distribution for 159 mm of liquid hydrogen [74]. Red error bars include statistical errors only, blue error bars indicate statistical plus systematic errors.

multiple scattering. Furthermore, the GEANT4 model strongly overestimates scattering, with a factor of 4 discrepancy in the large angle bins. In all but the largest angle bins where the errors are largest, there is a very strong agreement between data and the ELMS model.

These measurements demonstrate that in the low- Z region – which is of particular interest for muon cooling – either form of the Molière theory is inadequate to explain the scattering distributions and that the simulation algorithms used by GEANT4 at the time vastly overestimate large angle scatters.

4.3 Scattering Measurements in MICE

While muon scattering measurements are not the primary aim of MICE, they do play an important role in the experiment, both in their own right and in order to better characterise the cooling performance of the MICE absorbers.

MuScat provided valuable data on muon scattering at 172 MeV/c and how it compares with different models, but it is not necessarily a given that agreements and disagreements will hold equally at other momentum ranges. In contrast to MuScat, MICE – though it will only run with liquid hydrogen and lithium hydride as absorbers

– can scan across a fairly broad range of muon momenta and thus fill in an image that MuScat sampled at a single point, as agreement between model and data at one specific momentum does not necessarily translate to other momentum ranges.

4.4 Measurement Strategy

Different scattering analyses have been performed with MICE, described e.g. in [66], [63], and a MICE paper currently in preparation. These analyses have generally approached the calculation of scattering angle distributions using deconvolution methods, where a “true” distribution is obtained from a measured one by an iterative process using a so-called “response matrix” generated from simulation. One weakness of such an approach is that the deconvolution methods used are heavily affected by a parameter called the regularisation strength – a measure for the weight put on data vs. prediction in order to prevent overfitting – the choice of which is difficult to empirically justify [76]. Additionally, it introduces a model dependence due to assumptions about the input beam.

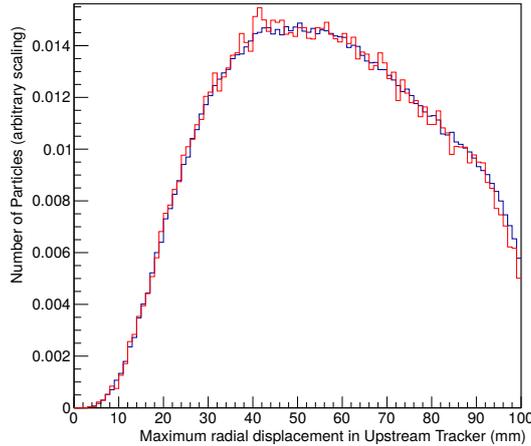


Figure 4.5: Normalised distributions of the maximum radial displacement from the beamline centre in the upstream tracker for 140 MeV/c beams with (blue) and without (red) absorber, demonstrating the equivalence between input beams.

For this analysis, a convolution approach was taken, which though it cannot produce a “true” scattering distribution, enables a direct comparison of modelled and measured distributions without the involvement of regularisation parameters. Rather than attempting to remove effects such as detector responses from the data, it measures them directly and then adds them to the prediction. It also makes only one assumption

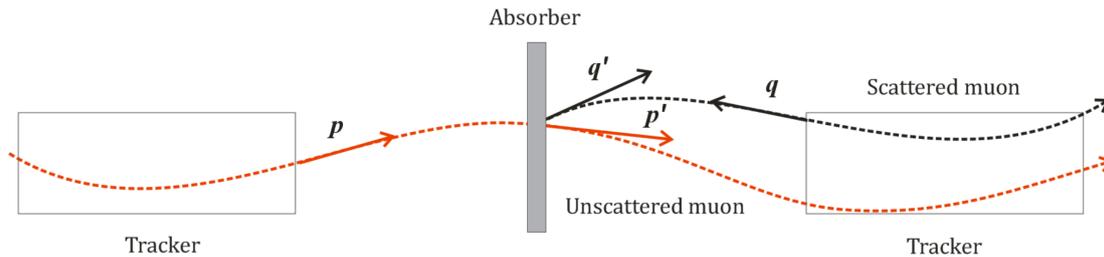


Figure 4.6: An illustration of the propagation required to calculate scattering angles in MICE in the presence of a magnetic field[66].

about the input beams, namely that they are equivalent between runs with and without absorber (see figure 4.5).

This analysis uses data taken in September and October 2016, both with and without the lithium hydride absorber, at 4 different momentum settings for each configuration (subdivided for analysis into 3 momentum slices each). For data-taking without an absorber (hereinafter referred to as “empty” data), the decay solenoid could not be used due to technical problems, resulting in samples of significantly fewer events.

Since both the trackers and the absorber are located within a magnetic field, it is not possible to simply calculate the angles between the momentum vectors of the tracker tracks as it would be when analysing field-off data. Instead, the muons’ phase-space coordinates in either tracker have to be propagated (using the method described in section 3.2.1) to the centre of the absorber. The two-dimensional scattering angles are then calculated as

$$\theta_{x,y} = \arctan\left(\frac{q'_{x,y}}{q'_z}\right) - \arctan\left(\frac{p'_{x,y}}{p'_z}\right), \quad (4.14)$$

where p' and q' are the momentum vectors at the centre of the absorber propagated from the upstream and downstream end, respectively (see figure 4.6).

For each momentum slice analysed (120, 130, 140, 150, 160, 170, 180, 190, 200, 220, 230, 240 MeV/c, with each slice having a momentum range of the stated momentum ± 5 MeV/c, see also figure 4.14 to illustrate the momentum selection), three scattering distributions were produced: From both absorber and empty data, distributions were extracted as described above. The third distribution came from a Monte-Carlo simulation including the absorber, where rather than taking reconstructed phase-space vectors at the trackers, MC truth information just upstream and downstream of the absorber was used for the propagation, yielding a scattering angle distribution for only the absorber itself, without any effects from scattering in other beamline elements or the finite

reconstruction resolution of the trackers. A convolution between this absorber-only distribution with the “empty” distribution is then compared to the distribution generated from the absorber data. Whether a discrepancy exists between the two exceeding what would be expected from statistical variation and systematic errors then indicates how reliable the GEANT4 scattering algorithms are in LiH for the momentum ranges in question. An analysis of liquid hydrogen data was planned, but had to be excluded from this thesis since data-taking was delayed until (expected) October 2017. ELMS predictions were not included, as they are not currently available for LiH.

4.4.1 The Convolution Approach

The convolution approach used in this analysis differs substantially from the deconvolution methods used in previous scattering measurements in MICE as well as MuScat. A discussion is presented below, both from a mathematical point of view and using a toy Monte Carlo simulation.

Mathematical Background

For clarity, this section assumes a two-dimensional system $x - z$ where z is the direction along the beamline. Assume, for simplicity, that the input beam is a pencil beam with N particles as given by

$$f^{\text{in}}(x, \theta) = N\delta(x)\delta(\theta - \theta_i), \quad (4.15)$$

where the incoming particles interact with the absorber at a single point ($x = 0$) and θ_i is the angle with respect to the z axis of the incoming beam. For a particle scattering by an angle θ_s , the input and output angles are then related by $\theta_o = \theta_i + \theta_s$.

The angular distribution in the downstream detector can then be described as the convolution of the scattering model with the input beam (where $\theta_i = 0$ by construction in the equation below as the scattering angle is measured from the input particle direction on a particle-by-particle basis):

$$\begin{aligned}
f^{\text{out}}(\theta_o) &= f^{\text{in}}(x, \theta) * h(\theta_s) \\
&= \int \int h(\theta_o - \theta) f^{\text{in}}(x, \theta) d\theta dx \\
&= \int \int h(\theta_o - \theta) N \delta(x) \delta(\theta) d\theta dx \\
&= N \int \delta(x) dx \int h(\theta_o - \theta) \delta(\theta) d\theta \\
&= Nh(\theta_o) \\
&= Nh(\theta_s)
\end{aligned} \tag{4.16}$$

i.e. the output distribution ($Nh(\theta_o)$) is a measure of the scattering model ($h(\theta_s)$), which is the impulse response of the absorber. In a real-world experiment, there is of course an additional distribution arising from the finite detector granularity, but due to the commutativity of convolutions, this can be added into the equations at any point. The same is true of small scattering effects in minor beamline elements such as windows or even air. The spatial distribution of a real (non-pencil) beam can be integrated out, as long as the input beam distributions are identical in shape.

A complication arises due to the finite aperture size in the downstream detector, which introduces an acceptance effect $\epsilon(x, \theta_o)$, i.e. the probability of a particle being observed at the downstream detector depends on both the transverse position where the scatter occurred and the output scattering angle. This effect, which was first observed in toy Monte-Carlo studies (see section 4.4.2), is strongly dependent on the ratio between upstream and downstream fiducial volume cut: the tighter the downstream cut with respect to the upstream one, the more pronounced the effect becomes. This is due to the fact that particles with a high transverse amplitude, which in the “empty” dataset are still captured in the downstream tracker, may be scattered out of the fiducial volume when an absorber is present. Hence the efficiency in large scattering angle bins is suppressed in the absorber dataset with respect to the “empty” one. In order to compensate for this effect, a transmission-derived correction was applied to the absorber-only MC truth data before the convolution.

Assuming, again, that the input beams are identical in shape (which should be valid given that installation / removal of the absorber does not alter the upstream beam-line), we can integrate out the spatial dependence, and write the acceptance efficiency as $\epsilon(\theta_s)$, i.e. as only a function of *scattering* (rather than output) angle.

Now consider a general form of the input beam f^{in} . Including the efficiency effect,

we now have

$$f^{\text{out}}(\theta_o) = f^{\text{in}}(x, \theta) * [h(\theta_s)\epsilon(\theta_s)] \quad (4.17)$$

and more specifically

$$f^{\text{LiH}}(\theta_o) = f^{\text{in}}(x, \theta) * [h^{\text{LiH}}(\theta_s)\epsilon^{\text{LiH}}(\theta_s)] \quad (4.18)$$

$$f^{\text{empty}}(\theta_o) = f^{\text{in}}(x, \theta) * [h^{\text{empty}}(\theta_s)\epsilon^{\text{empty}}(\theta_s)] \quad (4.19)$$

We can furthermore write $h^{\text{LiH}}(\theta_s)$ as $g * h^{\text{LiH}'}$ (θ_s), where g is a scattering distribution that includes effects such as scattering in beamline windows etc, and $h^{\text{LiH}'}$ describes the scattering in the absorber disk only. Using the same convention for the empty distribution yields $h^{\text{empty}}(\theta_s) = g$, as $h^{\text{empty}'}$ = 1 (there is no notable amount of scattering in a slab of air the thickness of the absorber disk). The efficiency can also be split into two terms, such that $h(\theta_s)\epsilon(\theta_s) = g\epsilon_0 * h'(\theta_s)\epsilon'(\theta_s)$, where the unprimed and primed terms characterise effects due to everything except for the absorber and scattering in the absorber respectively.

The goal of the analysis is to compare the theoretical scattering prediction inside the absorber $h^{\text{LiH}',\text{th}}$ with the real scattering distribution inside the absorber $h^{\text{LiH}'}$, so we convolute $h^{\text{LiH}',\text{th}}$ with the empty absorber distribution. We can now see, that while ϵ_0 is applied by the detector to both distributions, the fact that $h^{\text{empty}'}$ = 1 implies that $\epsilon'(\theta_s)$ is missing from the theoretical prediction (since $\epsilon'(0) = 1$), i.e. we have

$$f^{\text{LiH}}(\theta_o) = f^{\text{in}}(x, \theta) * g\epsilon_0 * [h^{\text{LiH}'}(\theta_s)\epsilon^{\text{LiH}'}(\theta_s)] \quad (4.20)$$

$$f^{\text{LiH},\text{th}}(\theta_o) = f^{\text{in}}(x, \theta) * g\epsilon_0 * [h^{\text{LiH}',\text{th}}(\theta_s)] \quad (4.21)$$

This can finally be corrected by applying an efficiency correction $\epsilon^{\text{LiH}',\text{th}}$, derived from the same MC as $h^{\text{LiH}',\text{th}}(\theta_s)$, to that distribution, yielding

$$f^{\text{LiH}}(\theta_o) = f^{\text{in}}(x, \theta) * g\epsilon_0 * [h^{\text{LiH}'}(\theta_s)\epsilon^{\text{LiH}'}(\theta_s)] \quad (4.22)$$

$$f^{\text{LiH},\text{th}}(\theta_o) = f^{\text{in}}(x, \theta) * g\epsilon_0 * [h^{\text{LiH}',\text{th}}(\theta_s)\epsilon^{\text{LiH}',\text{th}}(\theta_s)] \quad (4.23)$$

4.4.2 Verification with Toy Monte Carlo

A toy Monte Carlo simulation was implemented from scratch in order to verify the basic principles of the method under exclusion of systematic effects. The geometry for this simulation consisted of a LiH disk sandwiched between two aluminium sheets surrounded by a vacuum, all in a constant 4 Tesla magnetic field. For the empty distribution, the

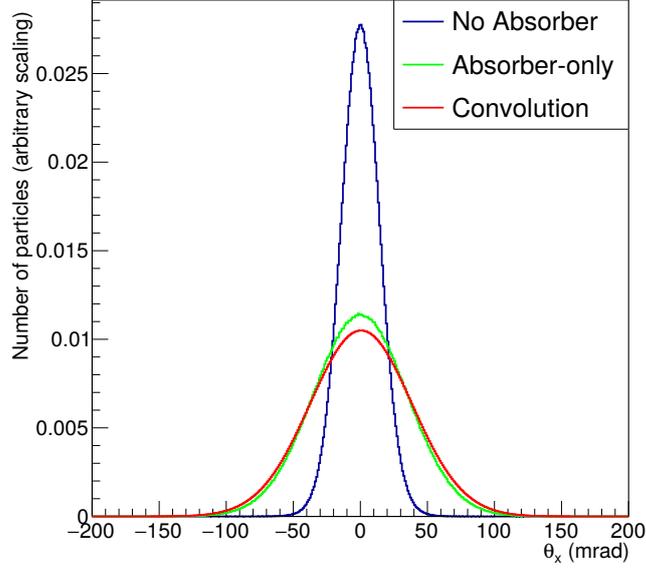


Figure 4.7: Toy MC scattering distributions without absorber, absorber-only, and the convolution of the two.

LiH disk was removed but the aluminium sheets remained. Particles were generated from Gaussian distributions in (x, y, p_x, p_y, p_z) centred around nominal values (all 0 except for p_z which was set to 160 MeV/c to correspond to a typical beamline momentum in MICE). To simulate detector responses, these coordinates were then smeared by Gaussian distributions with widths comparable to the errors reported by the tracker reconstruction in MICE, both at the beginning and at the end of the beamline. The true coordinates were propagated between the nominal detector planes (2 metres upstream and downstream of the absorber). For propagation through material, 1 mm steps were taken, where for every step, the transverse momentum components were modified by a scattering angle randomly drawn from a Gaussian with a width given by equation 4.3.

Scattering distributions were then calculated by propagating coordinates back into the centre of the absorber disk:

- The absorber-only distribution used exact coordinates on each edge of the LiH disk as the input into the propagation.
- “Empty” and reference distributions (incl. the absorber) used smeared coordinates from each end of the beamline from the dataset generated without and with absorber respectively.
- Absorber-only and “empty” distribution were convolved (figure 4.7) and compared to the reference distribution.

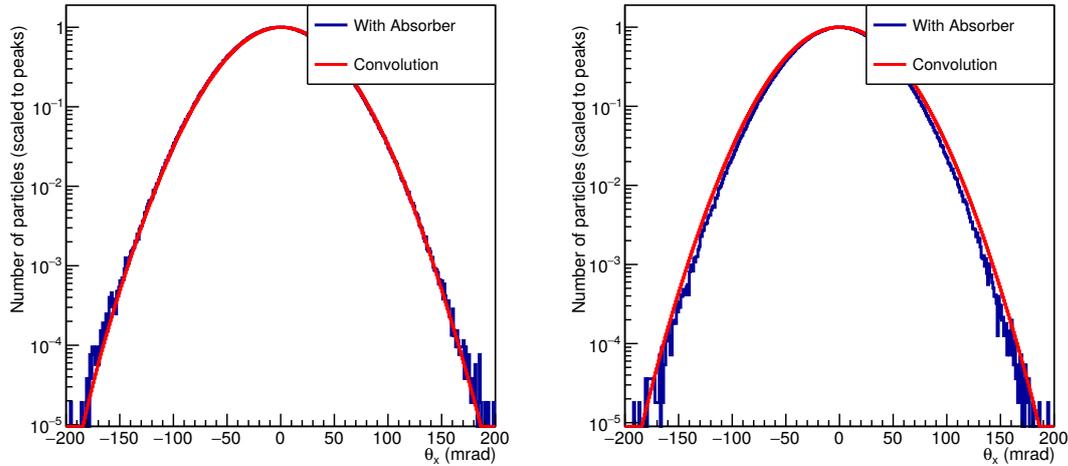


Figure 4.8: Comparison of scattering distributions with absorber and without absorber convolved with absorber-only scattering prediction for wide (left) and tight (right) downstream fiducial volume cut in the toy Monte Carlo. Distributions are scaled to identical peak values.

Fiducial volume cuts were introduced both upstream (75 mm, 150 mm, and 400 mm – the latter being effectively infinite given the beam distribution) and downstream (varied between 50 and 200 mm) of the absorber to evaluate their impact on the result.

As can be seen in figure 4.8, for a downstream fiducial volume cut that is wide compared to the beam, the agreement is perfect within statistical errors, which for the central part of the distribution are tiny, as the Toy MC was run for 5,000,000 accepted particles for most combinations of apertures (note that the red distribution is smoothed out due to the convolution). If the downstream cut is tightened, however, large angle scatters in the absorber distribution are suppressed, which provides the motivation for the efficiency correction described above.

4.5 Selection Criteria

A number of cuts were applied to the data in order to ensure high quality data sets with little contamination by pions and positrons. The different cuts and selections are outlined below. An investigation of possible systematic effects arising from them can be found in section 4.6. Illustrations of the cuts and selections on the parameters in

question using the LiH 170 MeV/c dataset can be seen in figures 4.9–4.15. The number of events remaining after each cut for the different datasets can be seen in tables 4.1–4.4.

4.5.1 Number of Tracker Tracks and TOF0/1 Space Points

In order to calculate the scattering angle, it is necessary to have tracks in both trackers which can then be propagated to the centre of the absorber.

Since the momentum in the trackers compared to the time of flight between TOF0 and TOF1 is the strongest PID indicator available (as the particle mass can be calculated from speed and momentum), it is also required to have space points in both.

Due to the low-level nature of the analysis, using track matching in order to pull apart events with multiple particles would add significantly to the systematics. Hence, exactly one track in each tracker and one space point in each of TOF0 and TOF1 was required for inclusion of an event.

4.5.2 Tracker Track Fit χ^2/NDF

Tracks with a poor track fit in the tracker were excluded by requiring that the reduced χ^2 from the fit be below 5.0 (see figure 4.9, the value was chosen by recommendation from one of the tracker experts).

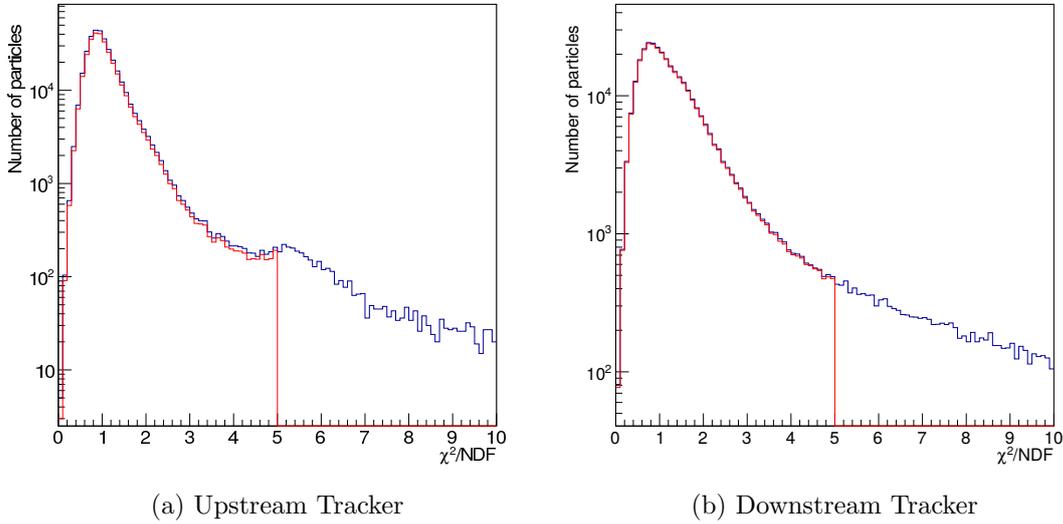


Figure 4.9: χ^2/NDF distribution before (blue) and after (red) the cut for the 170 MeV/c absorber dataset. Note that the slight reduction in the area that is not cut is due to events that are removed due to the cut in the respective other tracker.

	Remaining Events					
	With Absorber			Without Absorber		
Before Cuts	1290076			210835		
Require Tracks in both Trackers	486953			77823		
Exactly one SP in TOF0/1	386979			43933		
Exactly one Track in each Tracker	376926			42803		
χ^2/NDF Cut	343703			39377		
Minimum Number of Tracker Clusters	306442			36050		
Minimum p_t	214442			23922		
No Energy Gain in Absorber	212897			22444		
Fiducial Volume Cut	147092			15818		
Momentum Selection (MeV/c)	120	130	140	120	130	140
Momentum Selection	4950	47101	68384	1644	5697	5174
PID Cut	2901	38137	52826	798	4629	4511

Table 4.1: Remaining events after cuts for runs with 140 MeV/c nominal beam momentum

	Remaining Events					
	With Absorber			Without Absorber		
Before Cuts	1359825			260264		
Require Tracks in both Trackers	441012			98405		
Exactly one SP in TOF0/1	353823			50740		
Exactly one Track in each Tracker	347621			49617		
χ^2/NDF Cut	315533			45499		
Minimum Number of Tracker Clusters	285342			41462		
Minimum p_t	207564			26968		
No Energy Gain in Absorber	206421			25327		
Fiducial Volume Cut	138402			17465		
Momentum Selection (MeV/c)	150	160	170	150	160	170
Momentum Selection	25981	42814	37777	2731	5647	5200
PID Cut	8173	30956	32326	1099	4543	4179

Table 4.2: Remaining events after cuts for runs with 170 MeV/c nominal beam momentum

	Remaining Events					
	With Absorber			Without Absorber		
Before Cuts	2452656			320666		
Require Tracks in both Trackers	342441			49175		
Exactly one SP in TOF0/1	260006			30418		
Exactly one Track in each Tracker	255706			29778		
χ^2/NDF Cut	234063			27454		
Minimum Number of Tracker Clusters	218250			25874		
Minimum p_t	179685			21452		
No Energy Gain in Absorber	177174			18576		
Fiducial Volume Cut	93465			8712		
Momentum Selection (MeV/c)	180	190	200	180	190	200
Momentum Selection	16021	34784	23475	1844	3212	1744
PID Cut	7104	23972	19752	785	2343	1500

Table 4.3: Remaining events after cuts for runs with 200 MeV/c nominal beam momentum

	Remaining Events					
	With Absorber			Without Absorber		
Before Cuts	1517367			305387		
Require Tracks in both Trackers	331984			69568		
Exactly one SP in TOF0/1	248613			37685		
Exactly one Track in each Tracker	246617			37342		
χ^2/NDF Cut	229973			34817		
Minimum Number of Tracker Clusters	221326			33664		
Minimum p_t	185611			28381		
No Energy Gain in Absorber	182345			25201		
Fiducial Volume Cut	83467			10076		
Momentum Selection (MeV/c)	220	230	240	220	230	240
Momentum Selection	12768	27667	24376	1625	3426	2734
PID Cut	5344	14388	16151	612	1697	1880

Table 4.4: Remaining events after cuts for runs with 240 MeV/c nominal beam momentum

4.5.3 Number of Tracker Clusters

Since each tracker has 5 stations, each comprised of 3 planes, a well-reconstructed tracker event will consist of 15 tracker clusters. A lower number of tracker clusters results in lower quality of the reconstructed track, due to the larger position errors on planes with 2 rather than 3 clusters and/or entire planes missing from the input to the track fit. A minimum of 12 clusters for both upstream and downstream tracks is imposed in order to reject tracks with a low reconstruction quality (see figure 4.10)

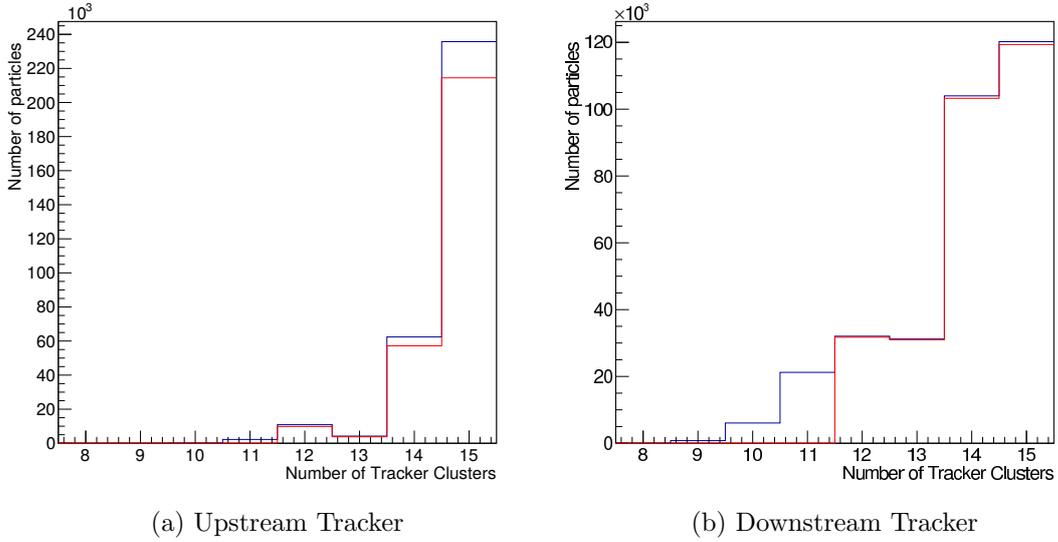


Figure 4.10: Number of Tracker Clusters before (blue) and after (red) the cut for the 170 MeV/c absorber dataset. Note that the reduction in the area that is not cut is due to events that are removed due to the cut in the respective other tracker.

4.5.4 Minimum Transverse Momentum in the Upstream Tracker

The longitudinal momentum of tracks in the tracker is calculated from the pitch of the helix that describes the particles' motion in the magnetic field of the trackers. This results in a strong dependence of the precision of the longitudinal momentum reconstruction on the transverse momentum as shown in figure 4.11. A minimum transverse momentum of 15 MeV/c was chosen to keep the error in the longitudinal momentum component less than half the bin size of 10 MeV/c chosen for the analysis (see figure 4.11).

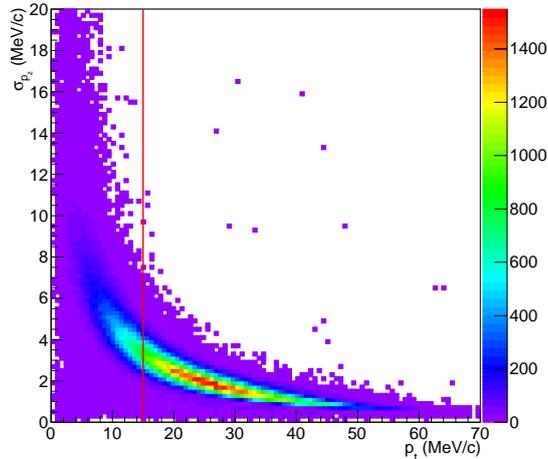


Figure 4.11: Uncertainty in p_z as a function of p_t for the 170 MeV/c absorber dataset. The red line indicates the cut, events to the left of it are removed.

4.5.5 No Energy Gain in Absorber

While the exact amount of energy loss in the absorber for a given particle cannot be predicted due to the stochastic nature of the processes involved, absorber traversal can never result in an energy *gain*, i.e. apparent energy gain in the absorber points to very poorly reconstructed momentum in one or both of the trackers. In order to account for the possibility of very low energy loss in the absorber, a requirement was imposed that the reconstructed total momentum in the downstream tracker cannot be more than 5 MeV/c (approximately the tracker momentum resolution) larger than that in the upstream tracker (see figure 4.12).

4.5.6 Fiducial Volume Cut

There are two main motivations for fiducial volume cuts in the trackers. In both trackers we want to avoid particles that came very close to the edge of the bore (i.e. the tracker housing's inner radius), both to cut out particles that may have scraped the tracker housing, and because the magnetic fields far away from the bore centre are less well understood.

A more stringent cut comes from the diffuser at the entry of the upstream tracker. As particles that passed through the diffuser housing rather than its aperture would be

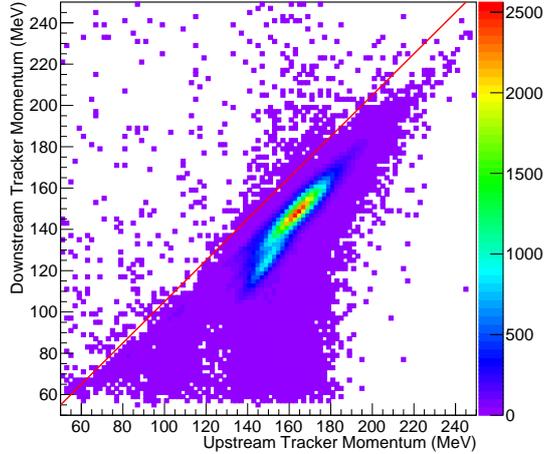


Figure 4.12: Reconstructed momentum in the downstream tracker as a function of reconstructed momentum in the upstream tracker for the 170 MeV/c absorber dataset. The red line indicates the cut, events above it are removed (retained in the plot).

subject to significantly less-predictable energy loss, inclusion of such particles would decrease the performance of the PID cut.

Finally, a tighter fiducial volume cut upstream than downstream mitigates to some degree the efficiency issues mentioned in section 4.4.1, thus requiring a less severe efficiency correction.

The fiducial volume cut is achieved by performing a simple least-squares circle fit on the transverse positions of the tracker track points and determining the maximum distance from the beamline centre of any point on the perimeter of the circle. In the upstream tracker this cut was placed at a radius of 70 mm as at that size, the aforementioned efficiency effect is relatively small (and the cut lies well within the diffuser aperture), while at the same time there is no significant loss of total particle rate after all cuts. In the downstream tracker, the cut was placed at 150 mm, i.e. the tracker’s nominal fiducial volume (see figure 4.10).

4.5.7 Momentum Selection

Since scattering angles are dependent on the particle momentum, ideally a dataset would have a very narrow momentum spread. However, in this case, the error associated with the momentum measurement (especially p_z), as well as statistical considerations, require a relatively coarse momentum binning. A width of 10 MeV/c was chosen for the mo-

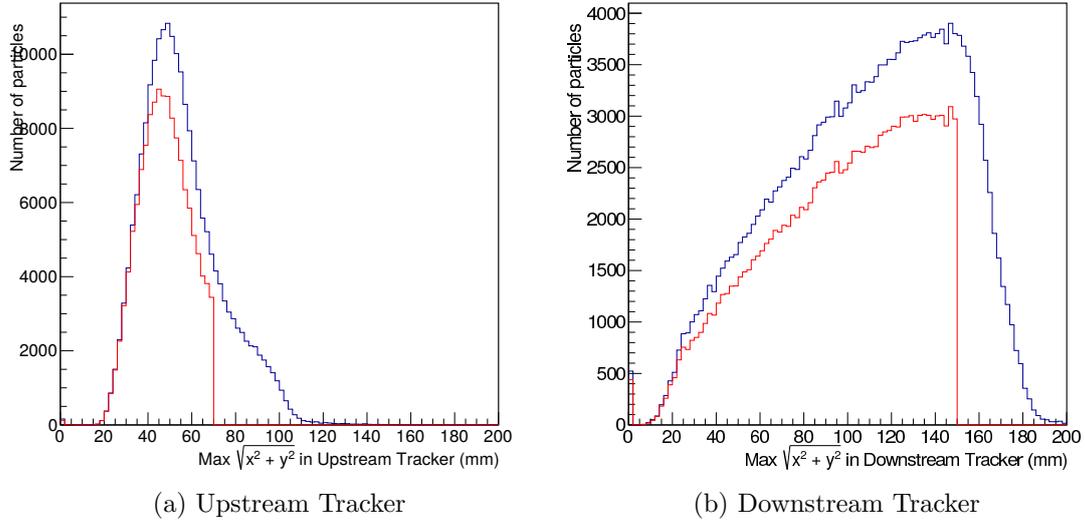


Figure 4.13: Maximum transverse position of the path inside the tracker before (blue) and after (red) the cut for the 170 MeV/c absorber dataset. Note that the reduction in the area to the left of the cut is due to events that are removed due to the cut in the respective other tracker.

momentum selection, as this is significantly larger than the uncertainty on the momentum for most of the particles in the beamline, retains relatively high statistics for each momentum slice, and allows coverage of most of the momentum space of the runs analysed (which had nominal beam momenta of 140, 170, 200, and 240 MeV/c). Figure 4.14 illustrates the selection of a specific momentum range from the distribution of the data set. Note that since the momentum distribution of the input beam is not uniform, the labels used below for a given momentum slice do not necessarily correspond to the mean momentum of that slice, but rather indicate the momentum range (nominal value ± 5 MeV/c) contained within it.

4.5.8 PID Selection

Pions and electrons were removed from the sample by making a cut on the TOF0-TOF1 time-of-flight, which given the narrow momentum range in the selected sample is a strong indicator of particle mass. Due to a slight (currently unexplained) offset in the parameter between data and Monte Carlo, MC truth information could not be reliably used to determine the appropriate values for the cuts. Instead, Gaussians were fitted to the peaks corresponding to the different particle species and the selection range placed 4

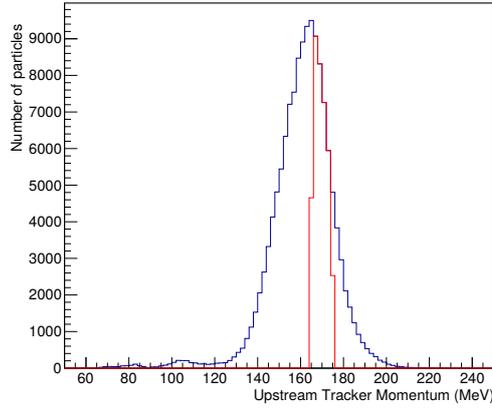


Figure 4.14: Momentum distribution in the upstream tracker before (blue) and after (red) momentum selection for the 170 MeV/c absorber dataset.

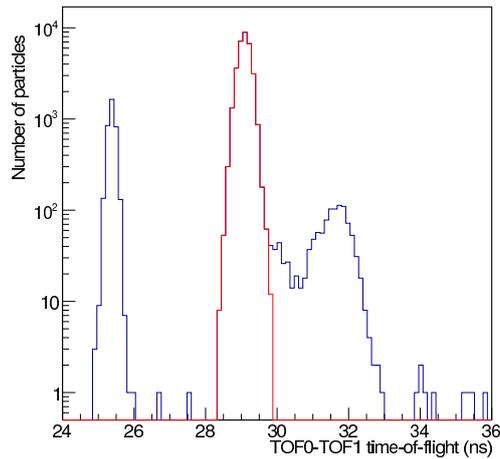


Figure 4.15: TOF0–TOF1 time-of-flight distribution before (blue) and after (red) PID selection for the 170 ± 5 MeV/c momentum slice. The three peaks (from left to right) correspond to positrons, muons, and pions respectively, the population between the muon and the pion peak was found from MC to consist of pions that decayed to muons between TOF0 and TOF1. For the depicted momentum selection, the (Gaussian fit derived) cut used was between 28.39 and 29.79 ns.

standard deviations from the centre of the muon peak. There exists a small population of particles spread between the pion and muon peaks. Using MC truth information (the aforementioned shift between MC and data was not an issue in this case, as this was a qualitative rather than quantitative question), it was found to consist of muons from

pion decays that occurred between TOF0 and TOF1. Since this population is very small compared to the peak itself, and the particles in it were not special other than in the fact that their parent pions decayed after as opposed to before TOF0, the exact location of the cut within that population is essentially irrelevant (see figure 4.15).

4.6 Systematic Errors

A number of potential sources of systematic errors were identified, based on the cuts mentioned above as well as various uncertainties regarding the MICE beamline. For each, the first step was to determine whether the given source has a measurable effect on the scattering distribution. Unless specifically stated otherwise below, this was done by varying the parameter in question in the Monte Carlo and then comparing sets with different values; depending on the type of variation either the standard value and a varied one, or one value varied upwards and the other one downwards (below referred to as A_n and B_n where $n \in \{1, 2\}$). For example, in the case of the cut on the number of tracker clusters, distributions were produced where the the imposed minimum was placed at either 11 or 13 (see section 4.6.2). To achieve statistical independence, the dataset was cut into two halves, which produced two pairs of data, where the n indicates the first or second half of the data. Data pairs were then compared via a χ^2 (using ROOT's `Chi2Test()` implementation based on [77]) and the parameter considered to not have a measurable systematic effect if the p-value calculated from the χ^2 values from comparing those data sets that varied the parameter (i.e. A_1 , B_2 , and A_2 , B_1 , respectively) was not noticeably lower than from those that varied only statistically (i.e. A_1 , A_2 , and B_1 , B_2 , respectively).

If an effect was found, a covariance matrix encoding the variance and correlations between bins of the scattering distribution was generated by running a number of simulations while varying the respective parameter according to a Gaussian distribution with width defined by the value of the systematic error. Correlations between systematic errors (likely existing e.g. between the upstream and downstream magnetic field uncertainties) were not taken into account, making the resulting error estimation conservative. Since this method is highly computationally intensive, the number of samples taken had to be kept as small as possible while still being representative in order to keep the processing time involved (even on a computing cluster) at a reasonable level. For this purpose, rather than directly drawing the variations randomly from a Gaussian with the appropriate width, a much higher number of random values was generated, sorted, and then every n -th entry drawn, yielding a distribution that follows the Gaussian much

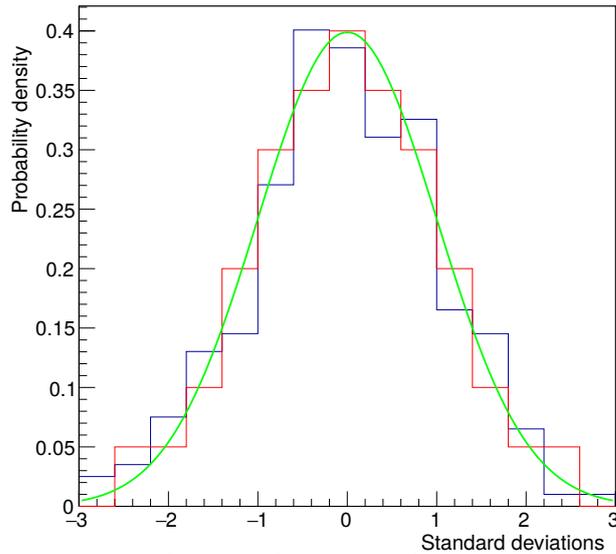


Figure 4.16: Illustration of the sampling smoothing used for evaluation of systematic uncertainties. Sample of size 500 (blue), sample of size 50 (red) drawn from a sample of size 10,000,000 and Gaussian distribution for reference (green).

more closely than would usually be the case for a similarly low number of samples (see figure 4.16). In order to include covariances between the x and y projections of the scattering angle, θ_x and θ_y , the two distributions were appended to each other for calculating the covariance matrix, yielding bimodal distributions with twice the number of bins. The following potential sources of systematic errors were considered:

4.6.1 Tracker Track Fit χ^2/NDF Cut

The χ^2/NDF cut was varied between 4.0 and 10.0 by recommendation from one of the tracker experts. No measurable systematic effect was found.

4.6.2 Number of Tracker Clusters Cut

The minimum number of tracker clusters required to pass the cut was varied between 11 and 13 (nominal value 12, values below 11 would result in a high probability of entire planes missing resulting in significantly degraded reconstruction quality). No measurable systematic effect was found.

4.6.3 Minimum Transverse Momentum in the Upstream Tracker Cut

The value of the minimum transverse momentum cut was varied by ± 1.5 MeV/c, which was representative of the combined error of the transverse momentum components in the vicinity of the cut as reported by the tracker reconstruction (see section 2.6.3). No measurable systematic effect was found.

4.6.4 No Energy Gain in Absorber Cut

The cut was varied by ± 4 MeV/c, which was representative of the combined error of all momentum components of tracks in both trackers as reported by the tracker reconstruction. No measurable systematic effect was found.

4.6.5 Fiducial Volume Cut

Since the fiducial volume cut is based on a circle fit, it was first necessary to evaluate by how much this circle fit could vary based on the detector resolution. The x and y coordinates of the track points in a random selection of tracks were varied by the errors reported by the tracker reconstruction and the spread of the resulting largest distance from the beamline centre calculated. The spread was found to be approximately Gaussian, with a width of less than 1 mm in most of the sample that was considered. The fiducial volume cut was then varied by ± 2 mm. No measurable systematic effect was found for either the upstream or the downstream cut.

4.6.6 Momentum Selection

Similarly to the minimum transverse momentum cut (see section 4.6.3), the selection boundaries were varied by an amount representative of the combined errors of the momentum components (± 2.5 MeV/c). A sample covariance matrix for the systematic effect can be seen in figure 4.17, with matrices for all datasets in appendix A.

4.6.7 PID Selection

As mentioned in section 4.5.8, the exact location of the time-of-flight PID cut has minimal impact. From the uncertainties on the time-of-flight measurements (see section 2.5.2), it was calculated that the resolution of the cut could only vary the number of particles included by around 0.05 % as very few particles are contained within the areas a few tens of picoseconds on either side of the cut. This could not contribute a measurable systematic effect.

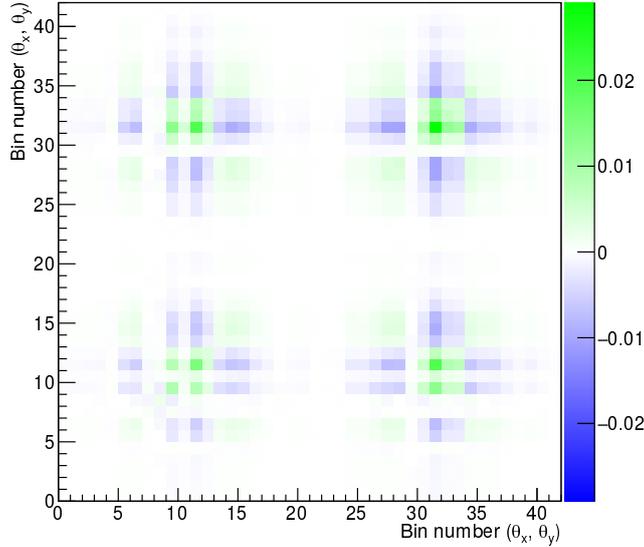


Figure 4.17: Sample covariance matrix for the momentum selection systematic for LiH data at 170 MeV/c. Bin numbers 1–21 and 22–42 represent the θ_x and θ_y ranges respectively between -0.2 and 0.2 radians in this and all following covariance matrices.

4.6.8 Amount of Absorber Material

The amount of absorber material in the beamline is a potential source of systematic errors since this would increase or decrease the width of the scattering distribution for the absorber itself obtained from the Monte Carlo simulation. The errors on mass, diameter, and thickness of the absorber are 0.5 g, 0.05 cm, and 0.05 cm, resulting in an error on the density of 0.72 %. For the actual scattering distribution, the relevant parameter is the product of density and width however, which, since the errors on thickness and density partially cancel, has an error of only 0.44 %. Using this uncertainty with equation 4.3 yields a final error on θ_0 of only 0.056 % – on the order of a few μrad , which is orders of magnitude below the levels that could be measured in MICE.

The atomic makeup of the absorber was measured with an uncertainty of 0.01 %, far too small to have a measurable effect.

4.6.9 Beamline Alignment

The transverse position of the downstream tracker was varied by ± 1 mm and its rotation by ± 1 mrad, which should be conservative estimates for the actual misalignment. Sample covariance matrices for the systematic effects can be seen in figure 4.18, with matrices

for all datasets in appendix A.

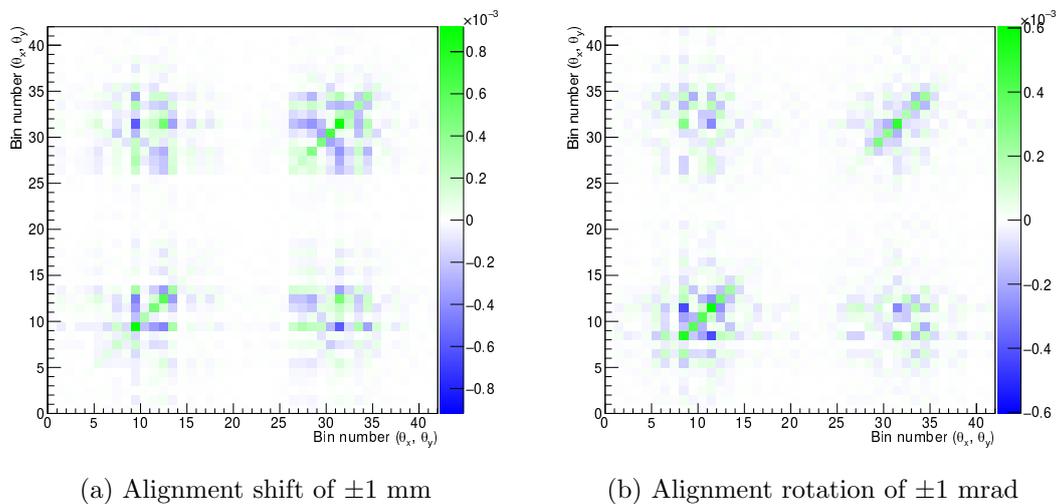


Figure 4.18: Sample covariance matrices for alignment and rotation systematics for LiH data at 170 MeV/c

4.6.10 Alignment Correction

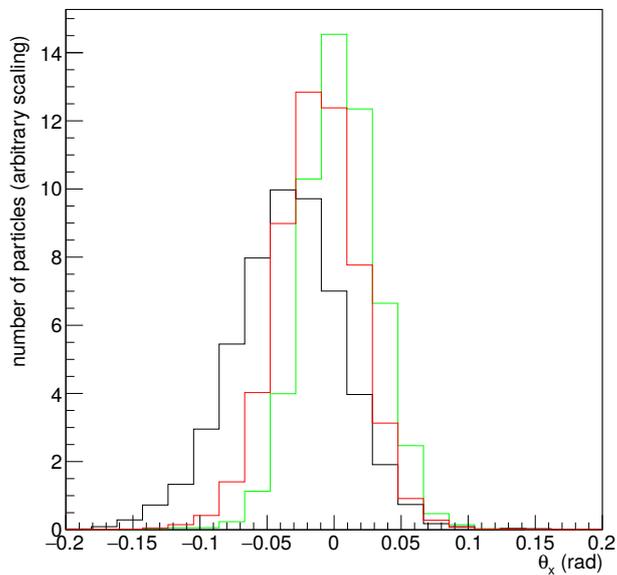


Figure 4.19: Scattering distributions for 180 (black), 200 (red), and 220 (green) MeV/c before shift correction, showing the momentum-dependent shift of the distributions.

When first analysing the data, it was discovered that in general, especially in the LiH dataset, the scattering distributions were not centred at 0. Moreover, the mean of the distributions actually varied as a function of momentum. Since there is no known physics (rather than reconstruction) effect that could cause such a shift, this triggered an investigation regarding the possibility of this being the result of a misalignment, i.e. a discrepancy in position between the actual beamline components and the geometry. To determine whether a simple misalignment could cause such an effect, a Monte Carlo simulation was performed in which a pencil beam varying only in longitudinal momentum was generated just ahead of the most downstream tracker plane in the upstream tracker. Scattering was turned off to eliminate as much random variation as possible and propagation back into the absorber was performed from MC truth rather than reconstructed coordinates. The downstream coordinates were shifted by 5 mm in x before propagation. As can be seen in figure 4.20, this does indeed produce a significant shift in reconstructed scattering angle which varies with input momentum.

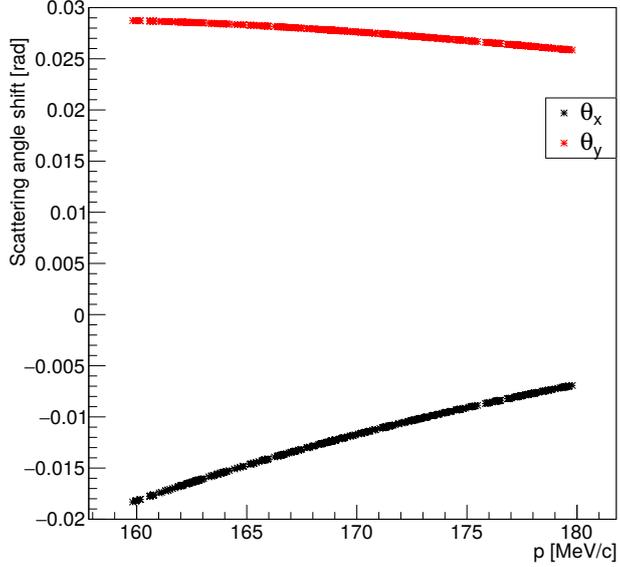


Figure 4.20: MC-based momentum-dependent shift in reconstructed scattering angle as result of a 5 mm spatial misalignment of the downstream tracker

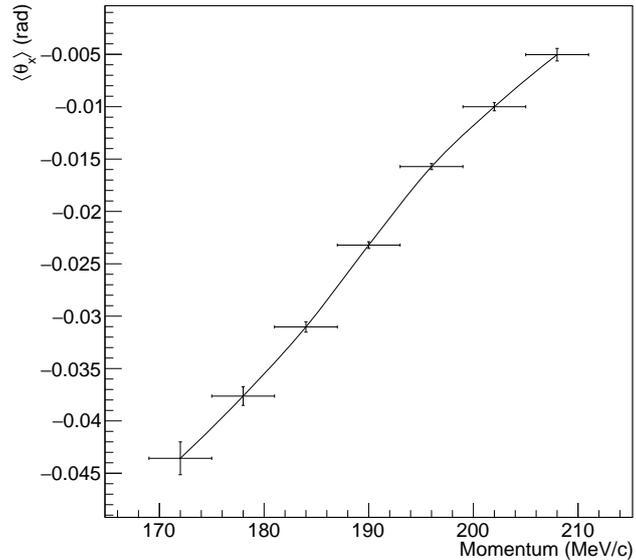


Figure 4.21: Shift correction graph for the dataset with a nominal beamline momentum of 200 MeV/c

It should be noted that since a spatial mismatch between the propagated tracks

does not factor into the calculated scattering angle, this effect can only occur because the magnetic fields are not completely uniform radially, since otherwise the propagation helix would only shift in the direction of the misalignment but not change its shape (and therefore still point in the exact same direction at the point where the scattering angle is calculated).

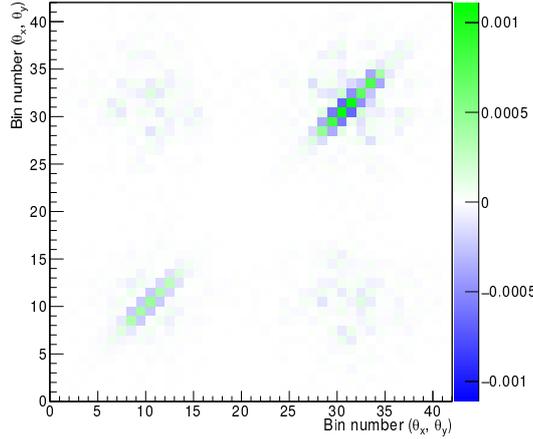


Figure 4.22: Sample covariance matrix for the alignment correction systematic for LiH data at 170 MeV/c

Ensuring that the scattering distributions are appropriately centred is important in order to numerically compare the LiH distribution with the convolution. Additionally, the momentum dependence may introduce a slight broadening. A full alignment study was not possible within the timeframe for this thesis: In its simplest form, it would have involved taking data without magnetic fields in the exact same beamline arrangement as was done for the used data, but since the beamline had been modified since, this was not possible. A field-on alignment study – i.e. using the data from this analysis – would be very involved as displacement and rotational misalignments cannot be treated independently, so a fit would have to be performed in at least 10 dimensions (three for displacement, two for rotation in each of the trackers). Hence the following method was used to mitigate the observed effect: Data was processed in two passes. In the first pass, angular distributions were binned by momentum (at 6 MeV/c per bin somewhat finer than the binning for the final plots) and for each bin, the mean scattering angle (i.e. the mean offset from 0) was calculated. In the second pass, the scattering angle of each event was corrected by interpolating between the previously corrected means in order to determine the appropriate correction at that specific momentum point. To determine

the systematic effect, multiple corrected samples were produced where in each case, for each particle the momentum used to determine the correction was varied randomly by a Gaussian with a width corresponding to the mean error on the momentum to simulate the distribution broadening that can still occur due to the uncertainty in momentum. A constant value for this mean error was used, as while it depends strongly on p_z (see figure 4.11), there is little dependence on the magnitude of the total momentum. A sample covariance matrix for this can be seen in figure 4.22. The fairly low covariances (statistical variances for the dataset pictured in figure 4.22 peak at around 0.017, see also figure A.21a) calculated for both this and the alignment systematics (for which the correction – without random variation – was also used) indicate that this approach is very effective in mitigating alignment problems.

4.6.11 Magnetic Field Strength

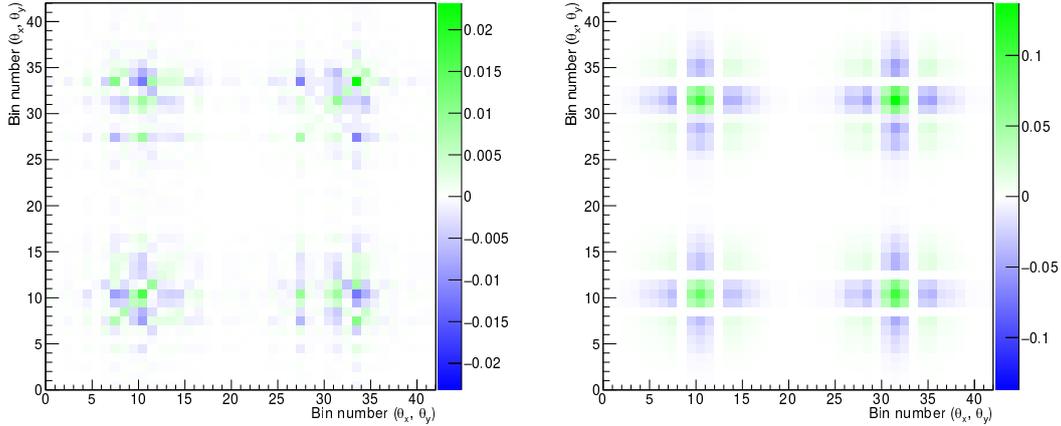
Comparison of the field map model with magnetic probes in the beamline yielded an uncertainty of around 2% in the strength of the magnetic field inside the trackers’ fiducial volumes. Field strengths in both trackers were separately varied by $\pm 2\%$. Uncertainties on the fields produced by the focus coils surrounding the absorber had been measured in flip mode (see section 2.4.1) to be less than 0.2% (no field map fits for solenoid mode have been made to date), so these fields were varied by $\pm 0.2\%$. Sample covariance matrices for the systematic effects can be seen in figure 4.23, with matrices for all datasets in appendix A.

4.7 Results

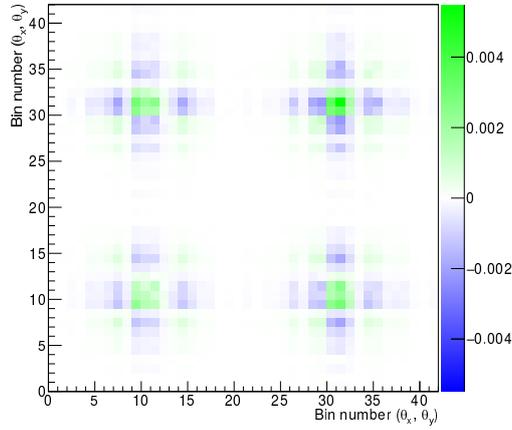
Comparisons between the LiH scattering distributions and those resulting from a convolution between “empty” data and the absorber-only distributions extracted from Monte-Carlo can be seen in figures 4.25 to 4.36, a comparison of RMS widths of the distributions across the entire momentum range in figure 4.24. The errors shown are calculated by taking the square roots of the diagonals of the sum of all covariance matrices (including those for statistical errors, which are diagonal matrices). The match between each two datasets was also evaluated numerically by calculating the χ^2/NDF value as

$$\chi^2 = (\mathbf{V}_{\text{LiH}} - \mathbf{V}_{\text{conv}})^T \mathbf{C}^{-1} (\mathbf{V}_{\text{LiH}} - \mathbf{V}_{\text{conv}}) \quad (4.24)$$

where V_{LiH} is the histogram of the LiH data scattering distribution (with θ_x and θ_y appended together as described above), V_{conv} is the histogram of the convolution of the



(a) Upstream magnetic field strength variation of $\pm 2\%$ (b) Downstream magnetic field strength variation of $\pm 2\%$



(c) Focus coils magnetic field strength variation of $\pm 0.2\%$

Figure 4.23: Sample covariance matrices for magnetic field uncertainty systematics for LiH data at 170 MeV/c

empty absorber scattering distribution with the MC derived absorber-only scattering distribution, C is the sum of all covariance matrices for the dataset, and NDF , the number of degrees of freedom, is the number of bins in the distribution (i.e. twice the number of bins in the θ_x and θ_y distributions) minus 2 to account for the independent scaling of θ_x and θ_y . The χ^2/NDF values for all datasets can be seen in table 4.5.

As can be seen from the covariance matrices in appendix A, the errors for the empty absorber datasets are dominated by statistical uncertainties, whereas for the LiH datasets, uncertainties in the magnetic fields and momentum selection, as well as

Momentum (MeV/c)	120	130	140	150	160	170
χ^2/NDF (stat + sys)	0.3033	0.6516	0.7669	0.3453	0.9097	0.5905
χ^2/NDF (stat only)	3.6426	13.8948	6.9199	1.1663	2.6374	2.4328
Momentum (MeV/c)	180	190	200	220	230	240
χ^2/NDF (stat + sys)	0.2971	0.3724	0.5979	0.4901	0.3560	0.5478
χ^2/NDF (stat only)	1.4401	2.4914	2.2661	1.2298	1.3587	0.9388

Table 4.5: χ^2/NDF values for all momentum selections

statistics contribute comparable amounts.

In general – especially for higher momenta – LiH data and the convolution mostly agree within errors, though the convolutions appear to be shaped slightly narrower, i.e. fewer large angle scatters, especially in the low-momentum regions (see also figure 4.24). This effect may at least partially be due to systematic effects, as especially those related to alignment will generally result in a broadening of the distribution. LiH data, as indicated by the larger covariances from the shift correction, appears to suffer from more severe misalignment issues which therefore causes a more significant broadening compared to the convolution. It should be noted however, that some previous analyses in MICE have produced results consistent with a slight underestimation of large angle scatters in GEANT4 (version 9.6 is used for MICE), while others have not seen such an effect[63, 78].

Only in the very lowest momentum ranges, especially 120 MeV/c (see figure 4.25, this was also the dataset with the lowest overall statistics), was a relatively severe discrepancy observed. Since scattering in this momentum range has never before been measured in MICE (or MuScat), future analyses may want to investigate this end of the momentum spectrum produced by MICE with higher statistics and potentially in the absence of magnetic fields.

The results presented here suggest that the multiple scattering algorithms used by GEANT4 have improved considerably in their treatment of low- Z materials since version 7.0 which was used in MuScat and which strongly overestimated large angle scatters. Furthermore, this analysis has demonstrated the efficacy of using a convolution approach to measure multiple scattering distributions. While this method cannot produce “true” distributions as can be achieved with a deconvolution (since detector effects cannot be removed from the data), it has a significant strength when comparing data to any scattering model, as it has no dependence on the choice of a semi-arbitrary parameter such as a regularisation strength. Additionally, this analysis represents the first time that field-on scattering measurements have been performed in MICE (the only previous

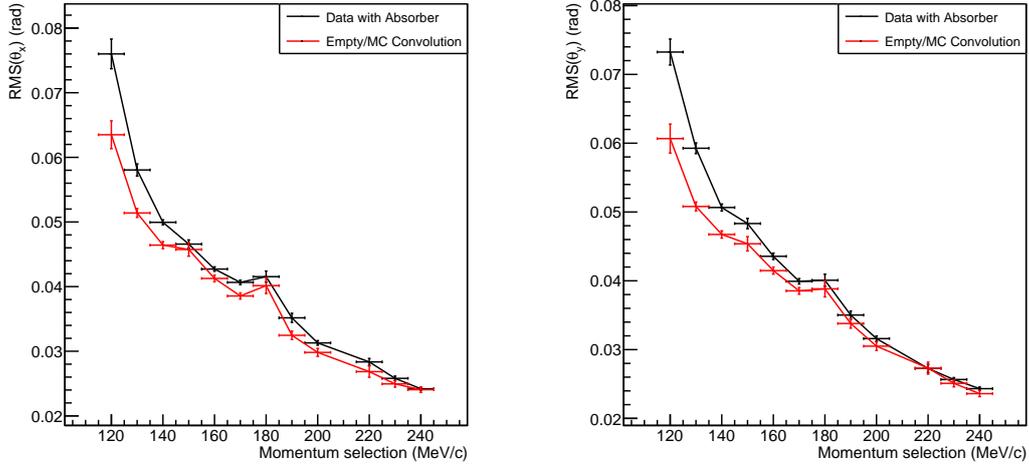


Figure 4.24: RMS widths of the scattering angle distributions for LiH data and convolution between “empty” data and absorber-only MC scattering distribution. The lower momentum range may indicate a slight underestimation of scattering by the GEANT4 scattering model. Note that the kink at 180 MeV/c corresponds to a dataset with large systematic errors, which can broaden the distribution.

field-on scattering analysis[63] was performed using Monte Carlo only, as real field-on data was not yet available at the time). While the presence of magnetic fields in the trackers allows a significantly better momentum resolution than relying on time-of-flight information (or even just beam settings) for the momentum selection, field uncertainties and momentum selection will become the leading source of error if significantly better statistics are available.

It is suggested that future multiple scattering measurements in MICE focus on the low-momentum region, which may show a significant discrepancy between data and Monte Carlo. If a convolution method is used in order to minimise the dependence on beam models and the regularisation strength, such an analysis would require significantly better statistics for no-absorber data. The most suitable approach in order to ensure strong results may be separate field-on (with fine momentum binning) and field-off (with coarse momentum binning) analyses, though the former would benefit from a detailed study of the beamline alignment and magnetic field maps.

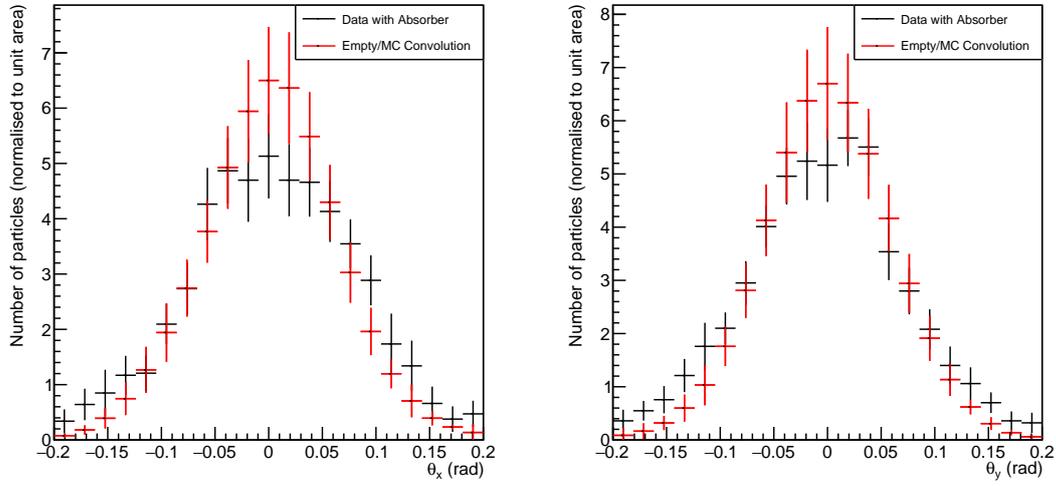


Figure 4.25: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 120 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

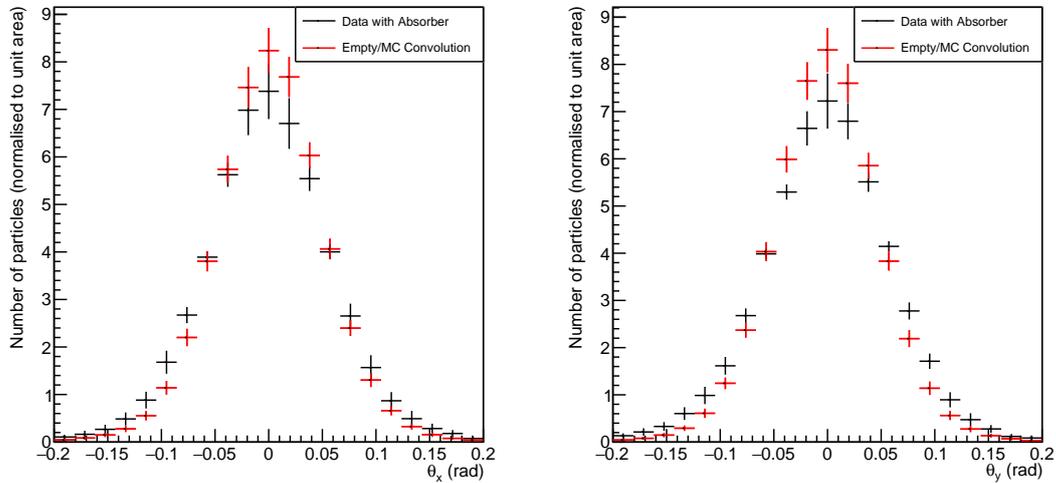


Figure 4.26: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 130 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

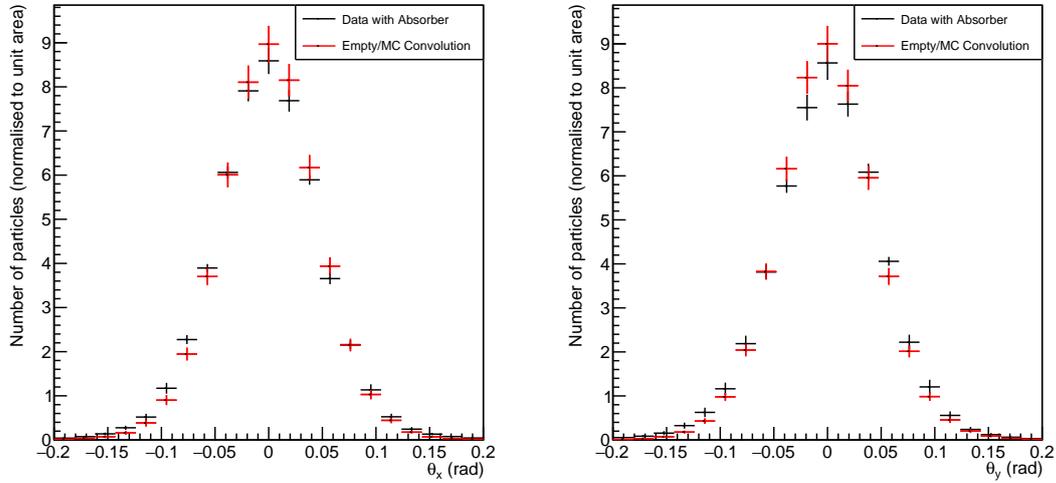


Figure 4.27: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 140 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

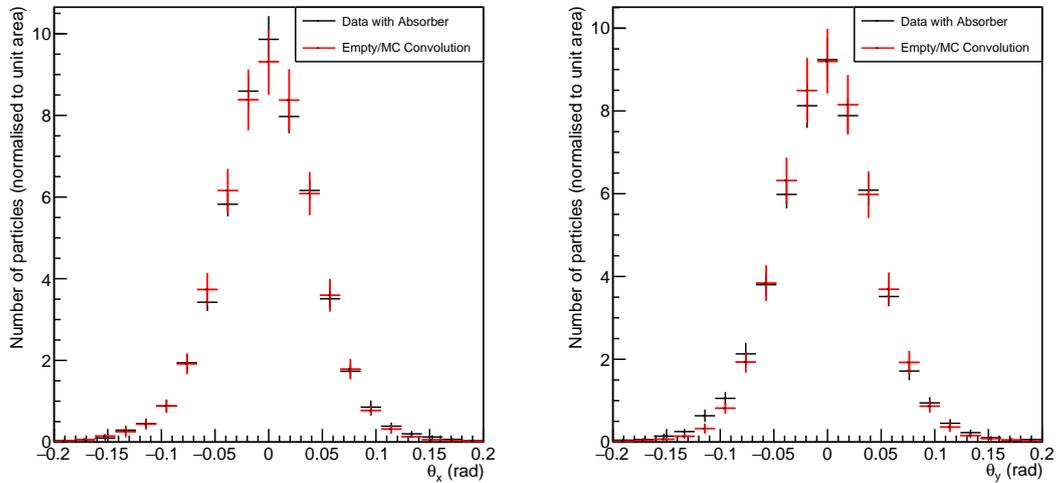


Figure 4.28: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 150 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

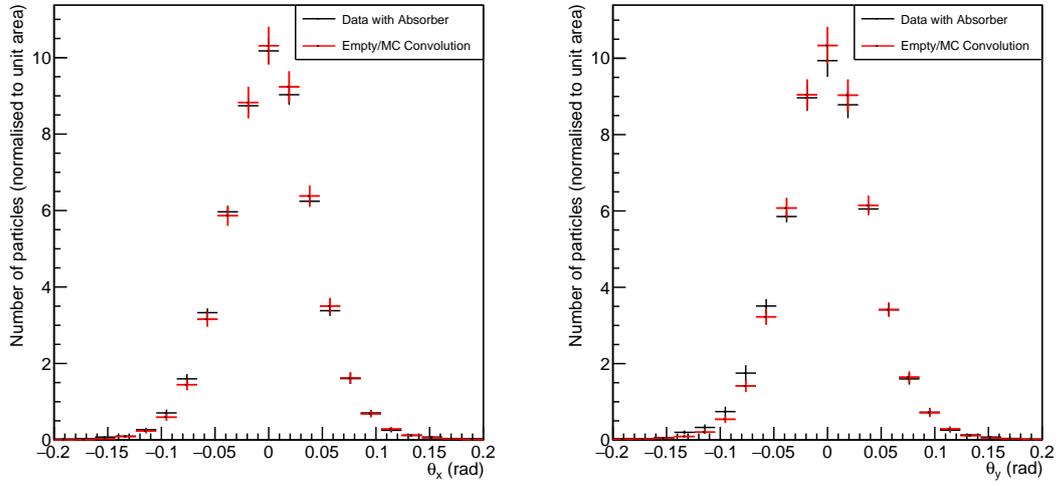


Figure 4.29: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 160 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

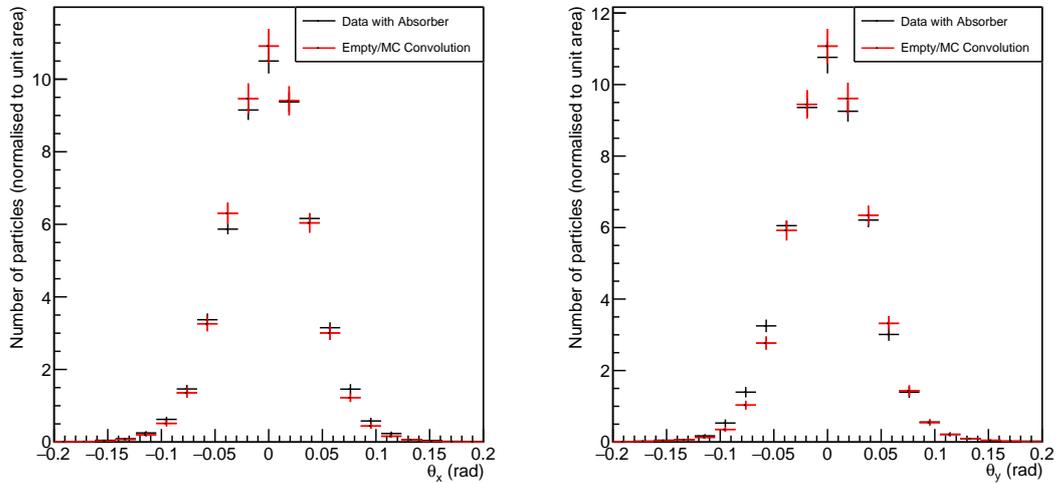


Figure 4.30: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 170 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

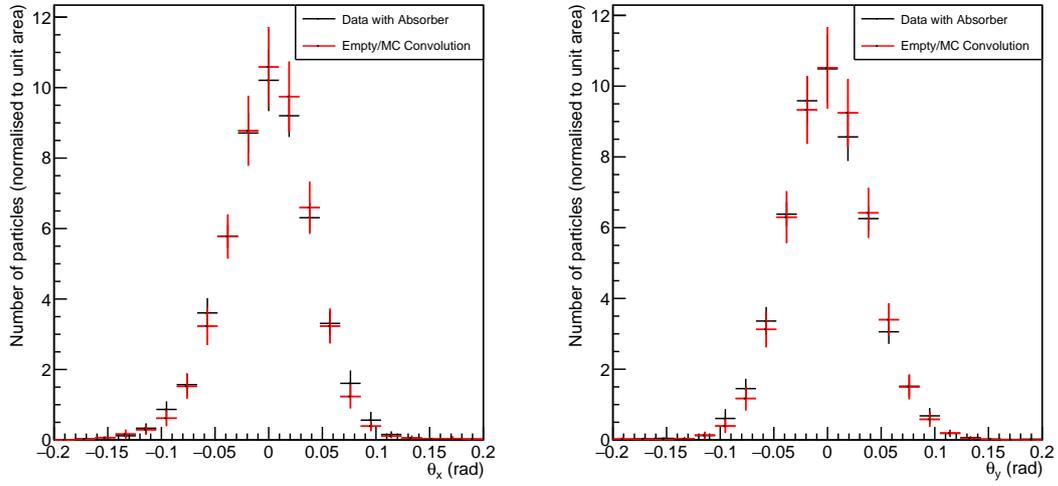


Figure 4.31: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 180 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

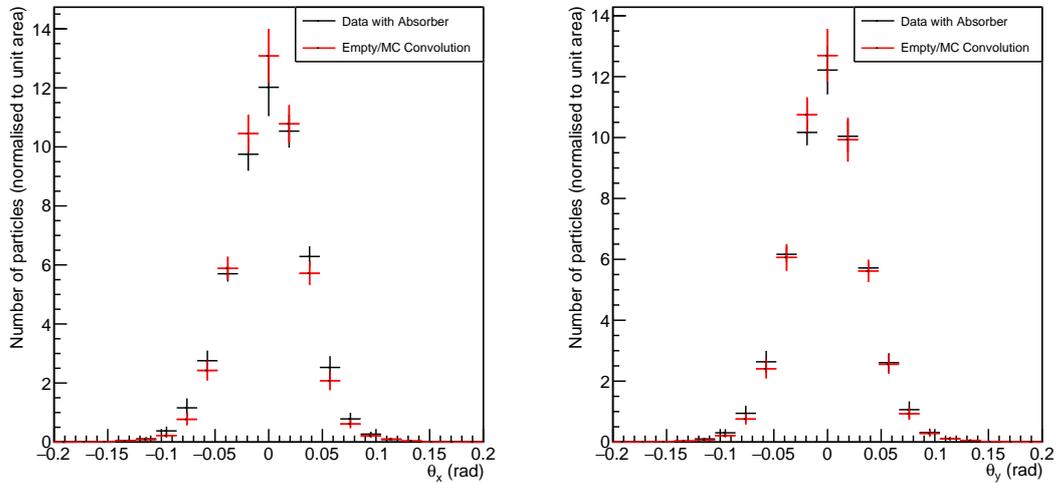


Figure 4.32: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 190 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

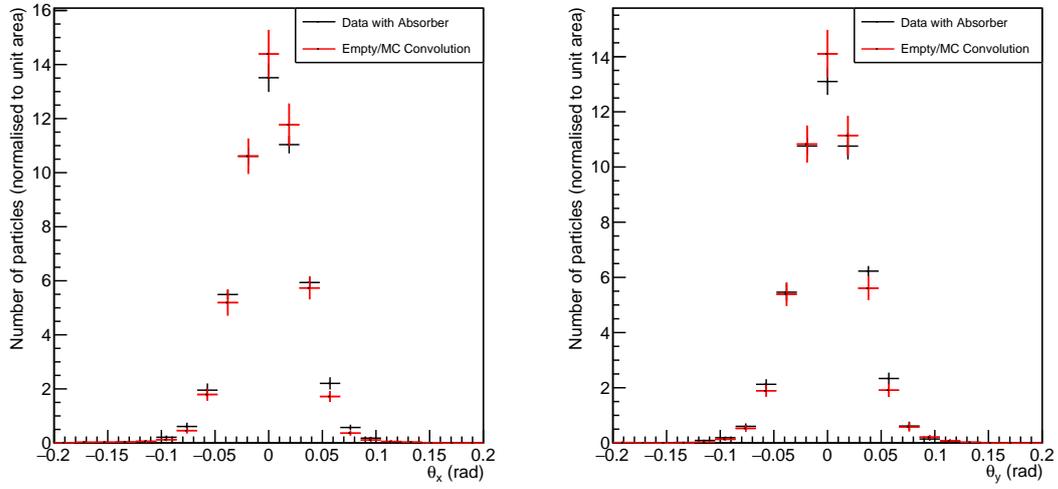


Figure 4.33: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 200 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

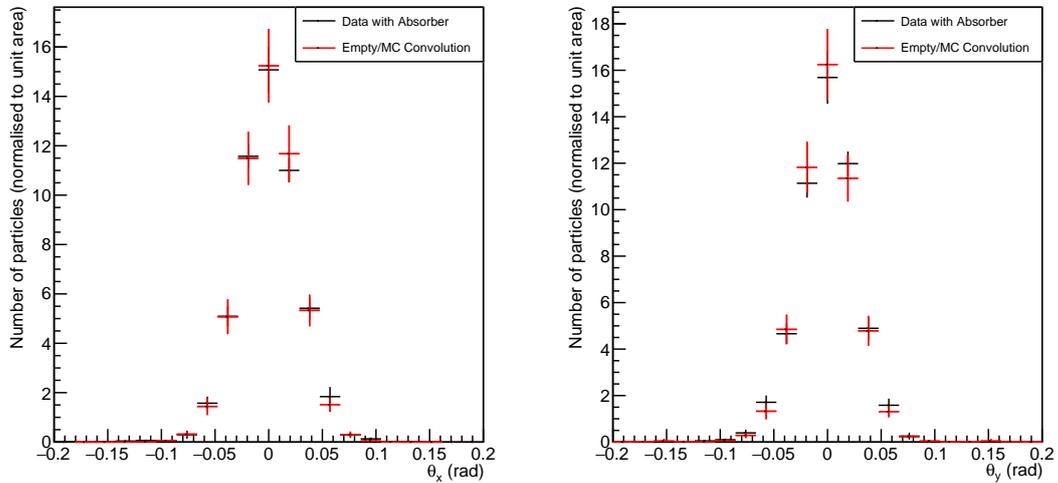


Figure 4.34: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 220 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

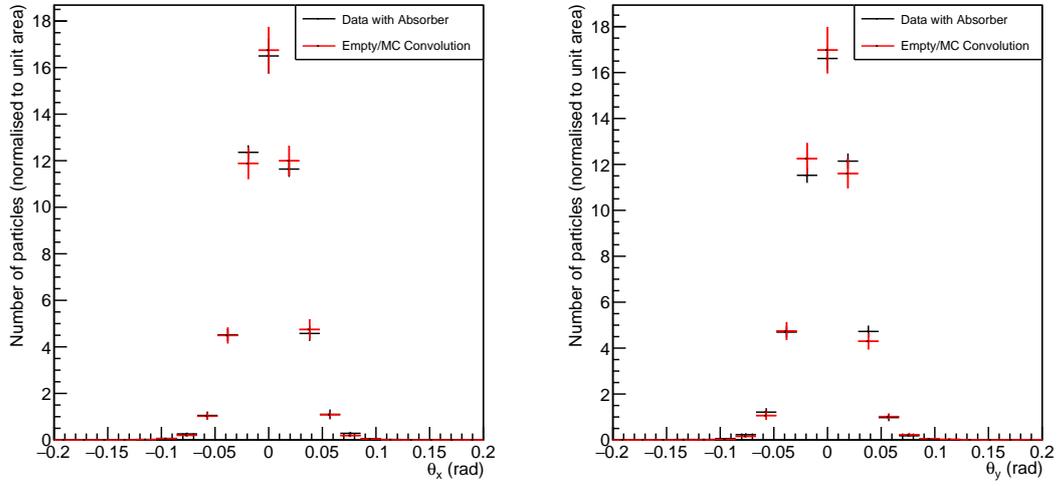


Figure 4.35: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 230 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

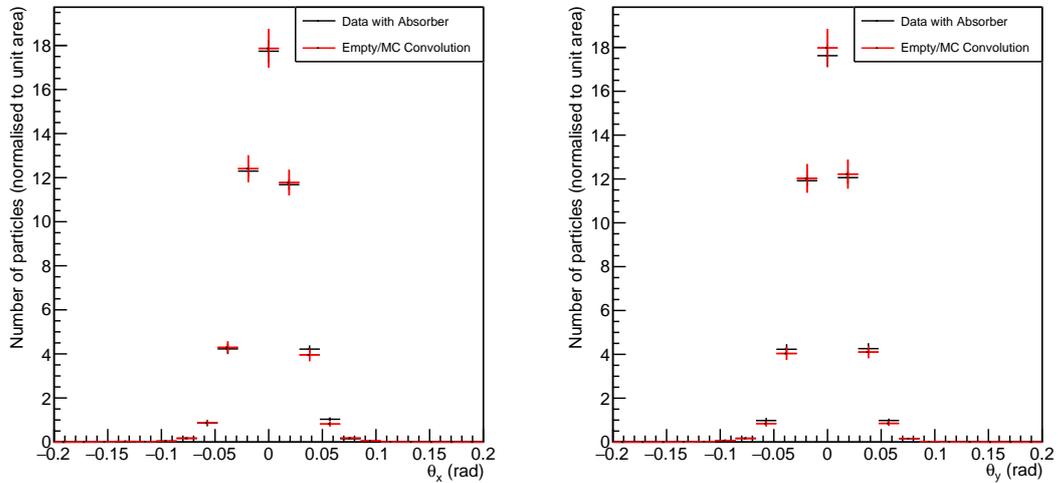


Figure 4.36: Comparison between LiH data (black) and convolution between “empty” data and absorber-only MC scattering distribution (red) at 240 MeV/c. Distributions are normalised to unit area. Total errors (i.e. statistical + systematic) are shown.

Chapter 5

Summary

“We never are definitely right, we can only be sure we are wrong.”

Richard P. Feynman

The Muon Ionisation Cooling Experiment (MICE) was designed as a proof of concept to demonstrate muon ionisation cooling – the reduction of the phase-space of a muon beam by passing it through alternating layers of a low- Z absorber (such as liquid hydrogen or lithium hydride) – which reduce its momentum in all directions – and RF cavities – which increase the momentum in the beamline direction only. Ionisation cooling is the only currently known method for muon cooling, as the short mean lifetime of muons makes typical methods used for electron or proton beams infeasible; and is an essential technology for the construction of a future Neutrino Factory and/or Muon Collider.

A track matching algorithm (presented in chapter 3) was implemented in order to combine locally reconstructed output from the various MICE detectors and present it within a unified global data structure for further processing such as particle identification, global track fitting, and finally analysis. Matching is performed by using a propagation routine based on a 4th order Runge-Kutta method to determine whether a given tracker track likely originated from the same particle as hits in the other detectors. Illustrative efficiency calculations show a good performance of the track matching algorithm on MICE data.

The efficacy of ionisation cooling is determined by the interplay of energy loss and multiple coulomb scattering in the absorber, as the energy loss drives the cooling whereas beam heating due to multiple scattering partially offsets this effect. The relative strengths of these two processes determines the equilibrium emittance of the beam, i.e.

the lowest emittance that can be achieved for a given input beam with a given absorber material. An expression for this equilibrium emittance using a fit to the Molière theory of scattering (section 4.1.3) can be seen in equation 2.10.

In order to be able to accurately model the cooling performance of an ionisation cooling channel, it is vital to validate the accuracy of scattering models in simulations. Results from the MuScat experiment[74] showed that GEANT4 at the time (in version 7.0) vastly overestimated large angle scatters in low- Z materials. Previous scattering measurements in MICE[63, 78] show significant improvement in agreement between GEANT4 and data (version 9.6 is used for MICE), but in some cases suggested that the simulations now slightly *underestimate* large angle scatters.

The analysis presented in chapter 4 of this thesis demonstrates an alternative approach for the scattering angle calculations in order to provide independent confirmation of these results. Whereas previous measurements used a deconvolution method, which is dependent on the choice of the regularisation strength – a measure for the weight put on data vs. prediction in order to prevent overfitting – a convolution method was used, in which the scattering distribution (including all detector effects) from data taken without an absorber present in the beamline is convolved with an MC derived scattering distribution of only the absorber itself. The presented results support the validity of this approach for future measurements.

Furthermore, previous measurements were taken in the absence of magnetic fields, which makes calculation of the scattering angles significantly more straightforward but suffers from very limited momentum resolution (as well as a loss of statistics due to lower transmission rates). This analysis deals with field-on data which allows fairly accurate momentum reconstruction, though at the cost of additional sources of systematic errors, namely the uncertainties in the field maps. Beamline misalignments also become more significant sources of errors in the presence of magnetic fields. Detailed alignment and magnetic field studies may be advisable before further field-on scattering studies are conducted on MICE data.

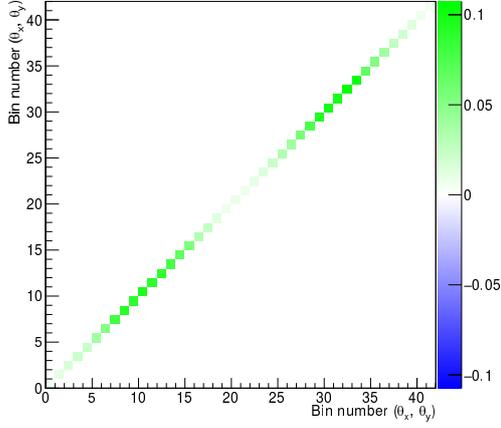
The results presented in section 4.7 support previous findings that the scattering model used by GEANT4 has improved significantly since MuScat results were published and lends tentative support to the presence of a slight bias in the opposite direction from what was identified at the time, though the uncertainties in this analysis – primarily driven by low statistics for the no-absorber data – result in an agreement between data and convolution within errors for most of the momentum range studied. Further study is warranted to confirm or refute the presence of this bias, and it is proposed that such study is performed using a convolution approach on both field-on and field-off data.

Appendix A

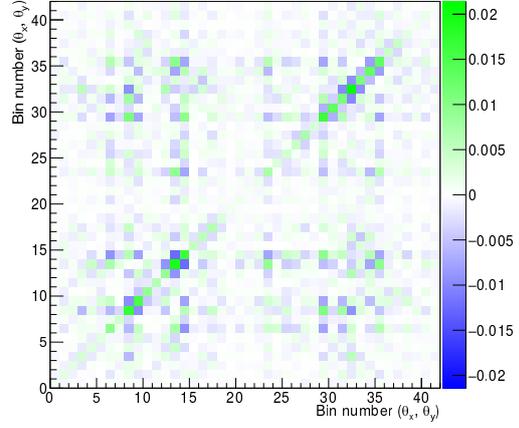
Covariance Matrices

Numerical versions of the covariance matrices listed below are available for download at http://micewww.pp.rl.ac.uk/projects/analysis/wiki/Jan_Greis_PhD_Analysis_Supplementary_Data.

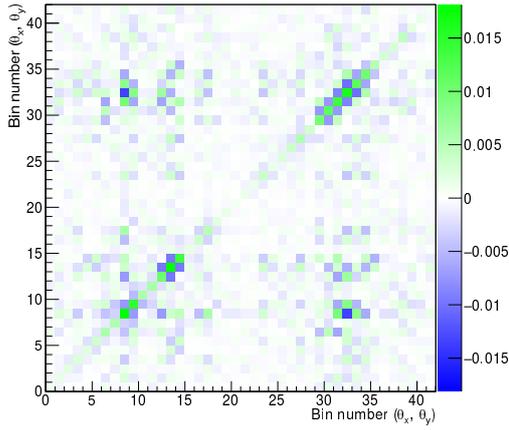
A.1 Covariance Matrices for Individual Systematic Errors



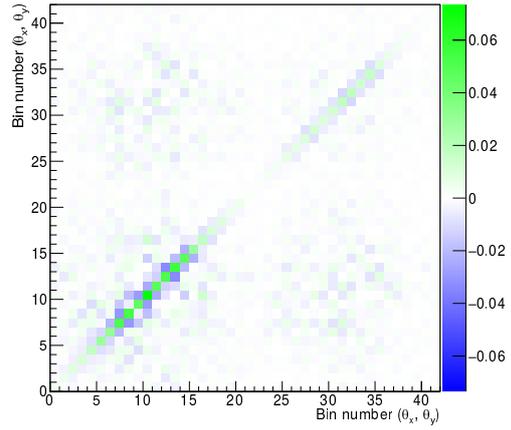
(a) Statistical errors



(b) Alignment shift of ± 1 mm

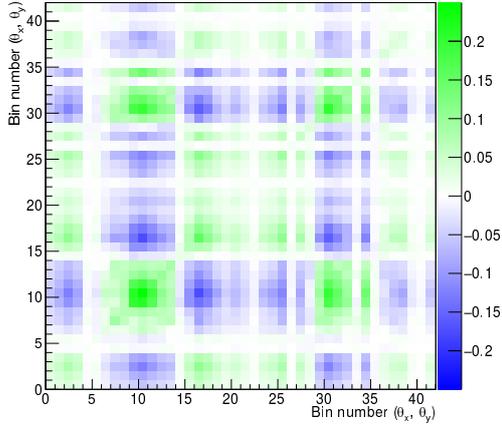


(c) Alignment rotation of ± 1 mrad

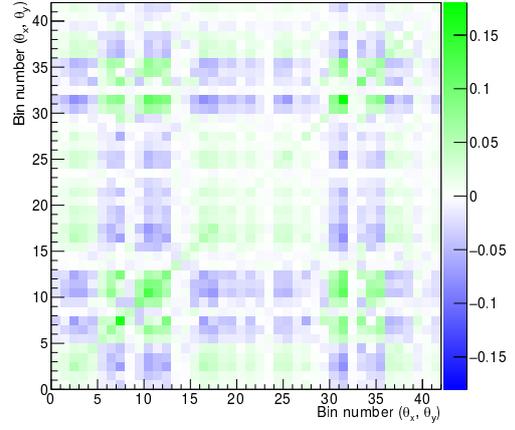


(d) Momentum-dependent shift correction

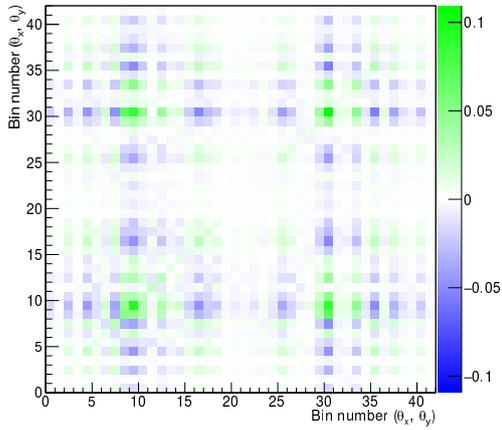
Figure A.1: Statistical and alignment covariance matrices for the 120 MeV/c LiH dataset



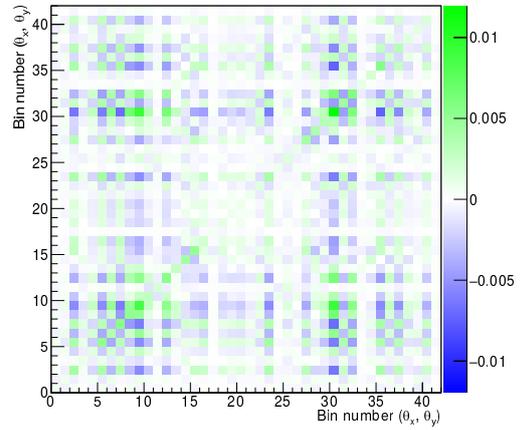
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

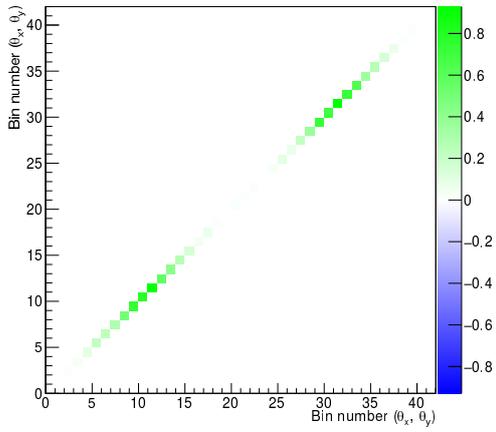


(c) Downstream magnetic field uncertainty

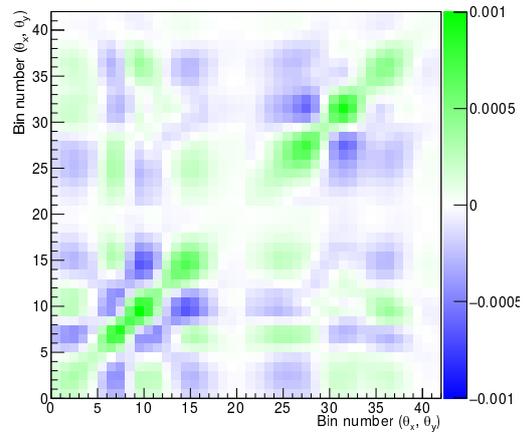


(d) Focus coils magnetic field uncertainty

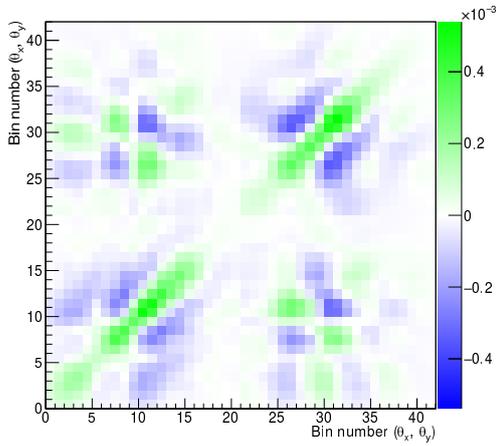
Figure A.2: Momentum selection and field uncertainty covariance matrices for the 120 MeV/c LiH dataset



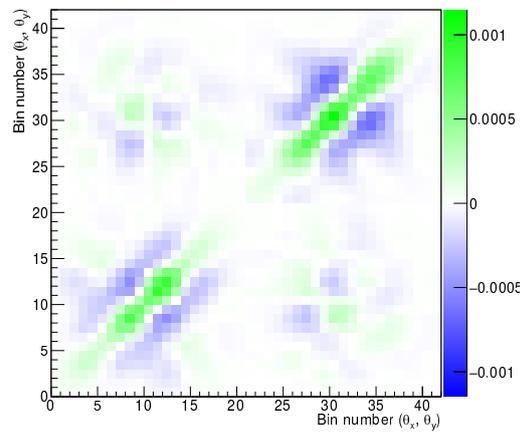
(a) Statistical errors



(b) Alignment shift of ± 1 mm

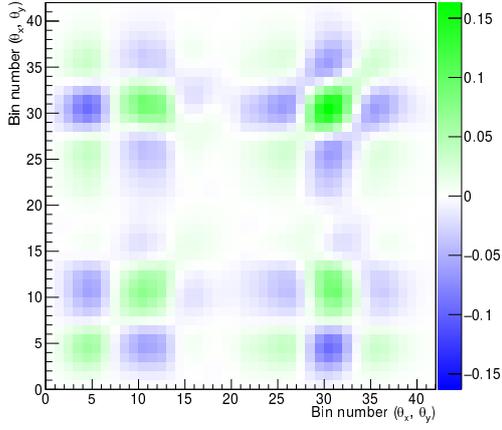


(c) Alignment rotation of ± 1 mrad

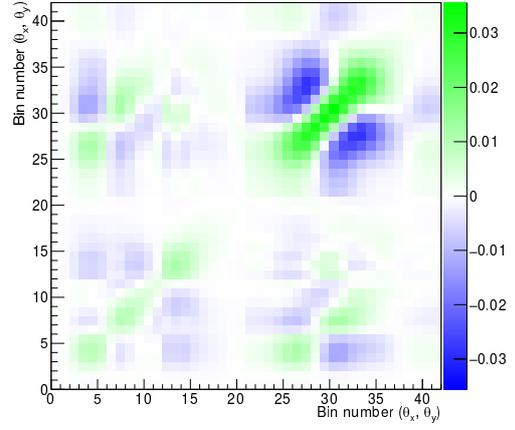


(d) Momentum-dependent shift correction

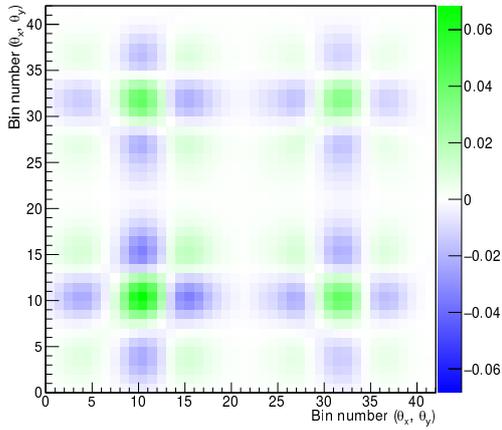
Figure A.3: Statistical and alignment covariance matrices for the 120 MeV/c convoluted dataset



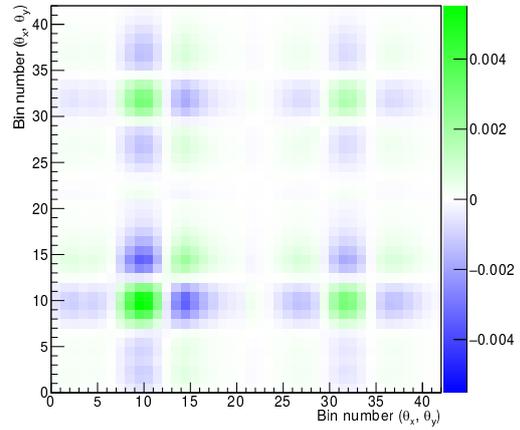
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

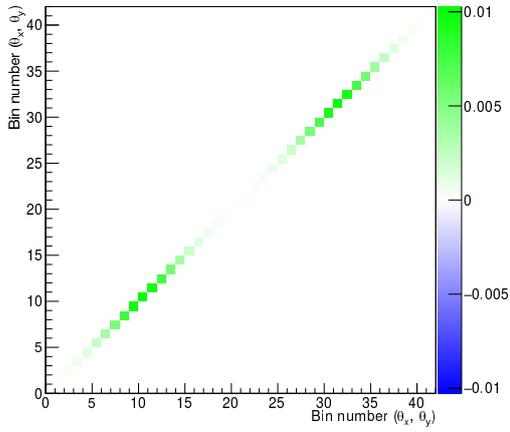


(c) Downstream magnetic field uncertainty

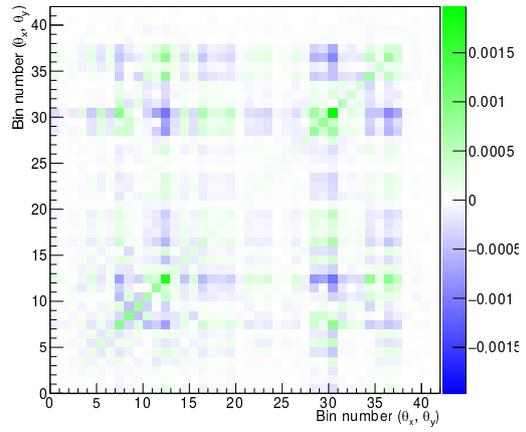


(d) Focus coils magnetic field uncertainty

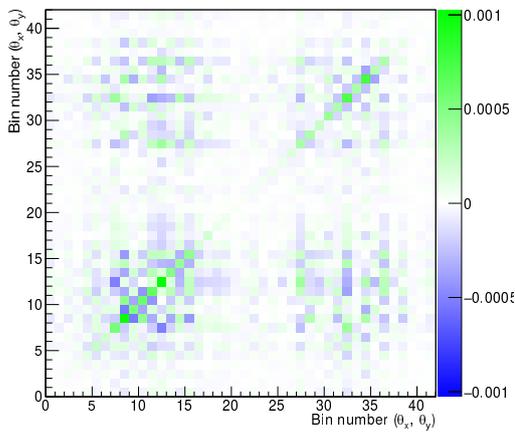
Figure A.4: Momentum selection and field uncertainty covariance matrices for the 120 MeV/c convoluted dataset



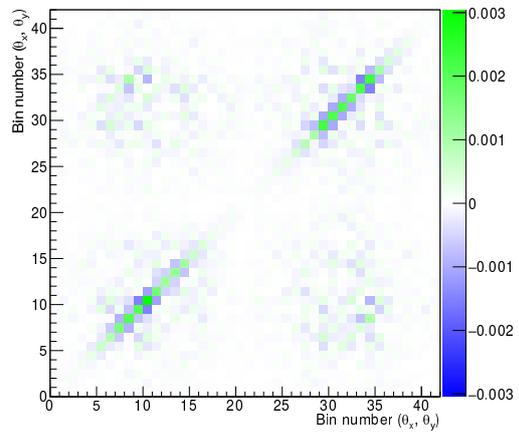
(a) Statistical errors



(b) Alignment shift of ± 1 mm

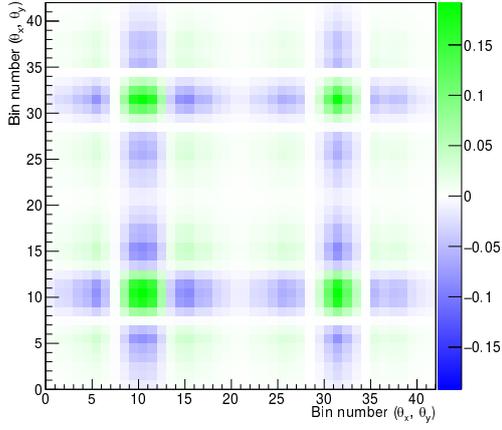


(c) Alignment rotation of ± 1 mrad

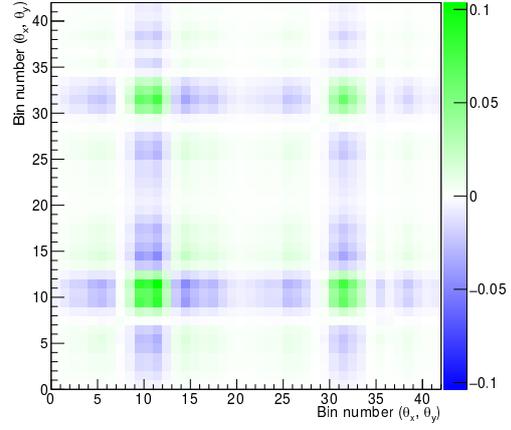


(d) Momentum-dependent shift correction

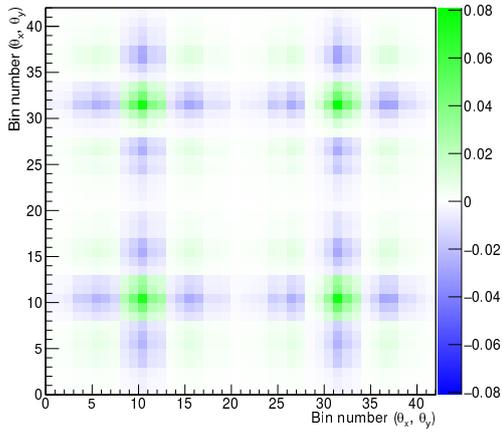
Figure A.5: Statistical and alignment covariance matrices for the 130 MeV/c LiH dataset



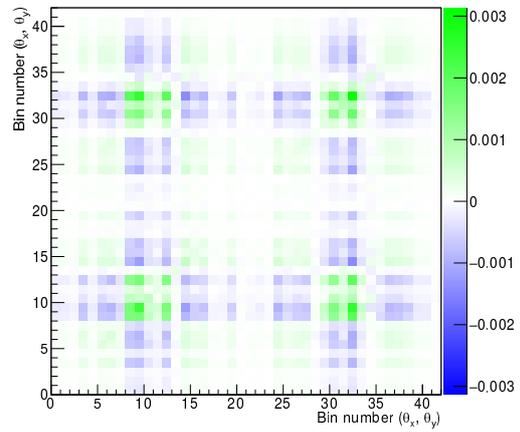
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

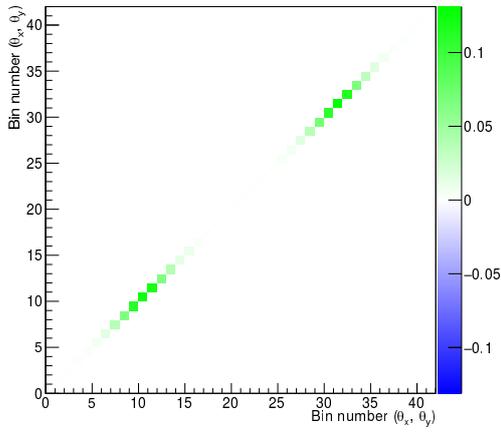


(c) Downstream magnetic field uncertainty

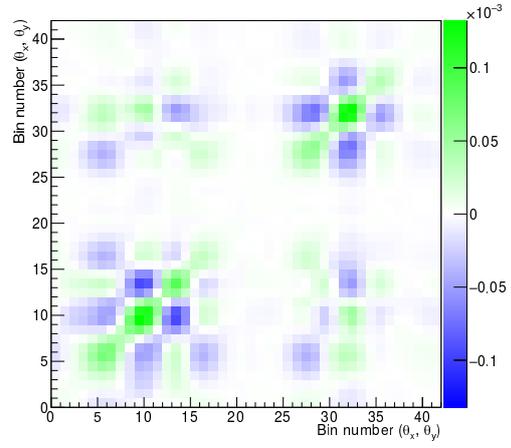


(d) Focus coils magnetic field uncertainty

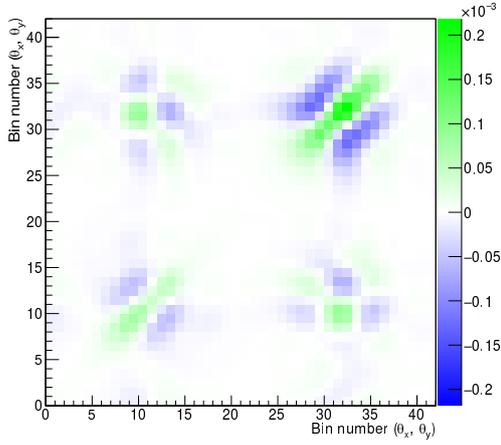
Figure A.6: Momentum selection and field uncertainty covariance matrices for the 130 MeV/c LiH dataset



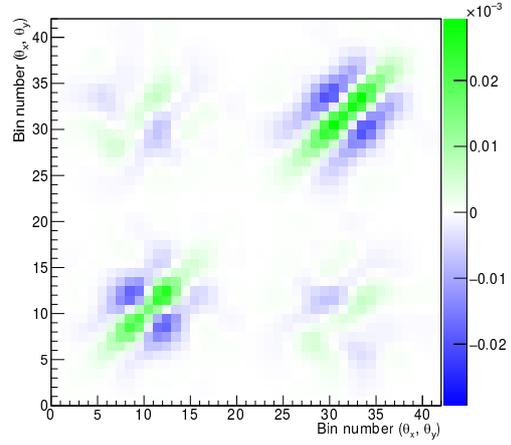
(a) Statistical errors



(b) Alignment shift of ± 1 mm

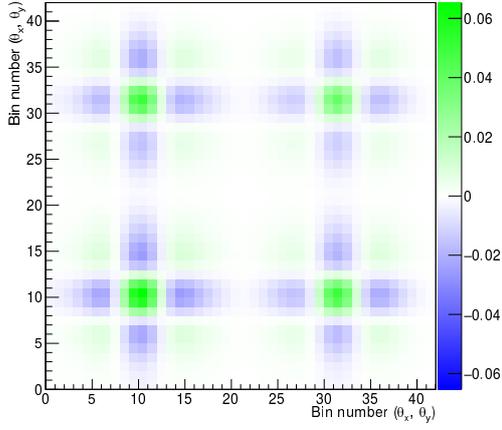


(c) Alignment rotation of ± 1 mrad

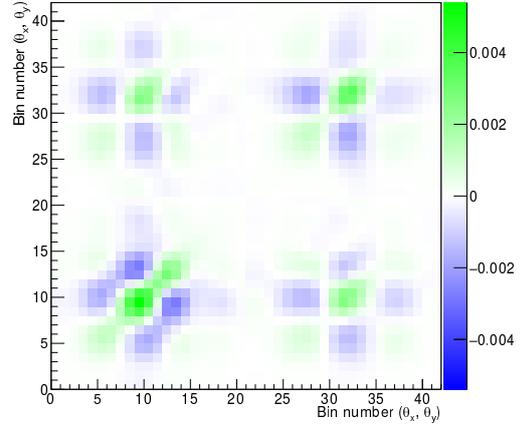


(d) Momentum-dependent shift correction

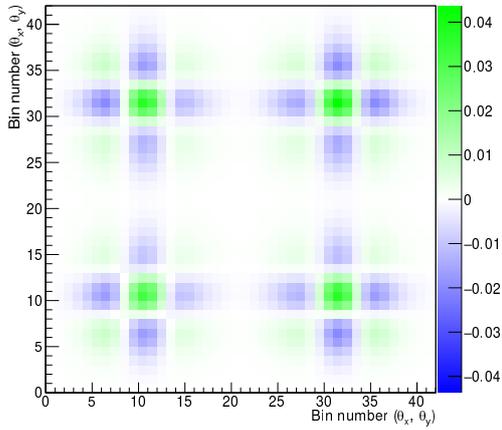
Figure A.7: Statistical and alignment covariance matrices for the 130 MeV/c convoluted dataset



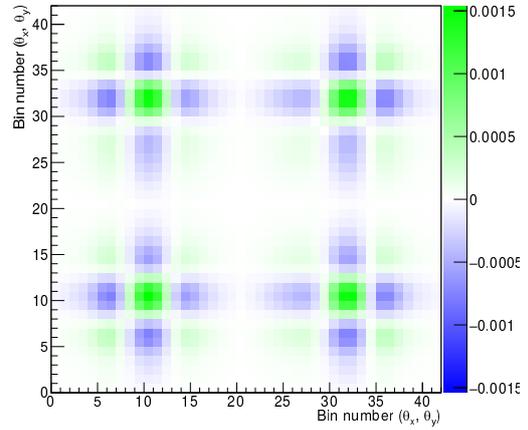
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

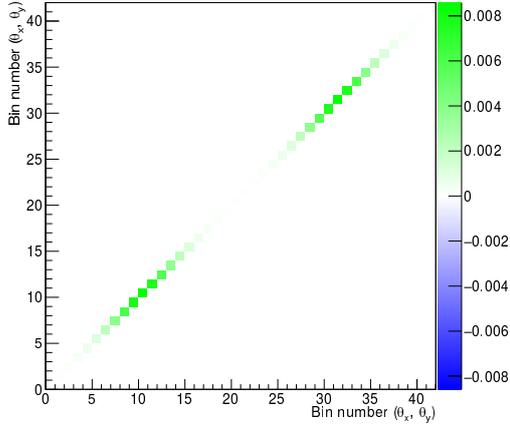


(c) Downstream magnetic field uncertainty

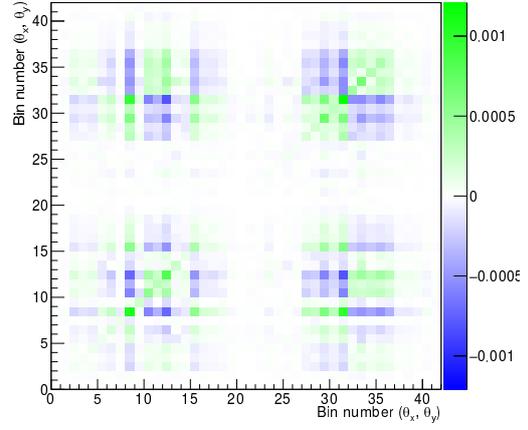


(d) Focus coils magnetic field uncertainty

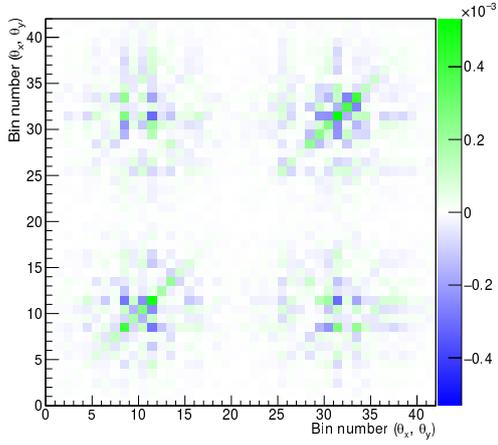
Figure A.8: Momentum selection and field uncertainty covariance matrices for the 130 MeV/c convoluted dataset



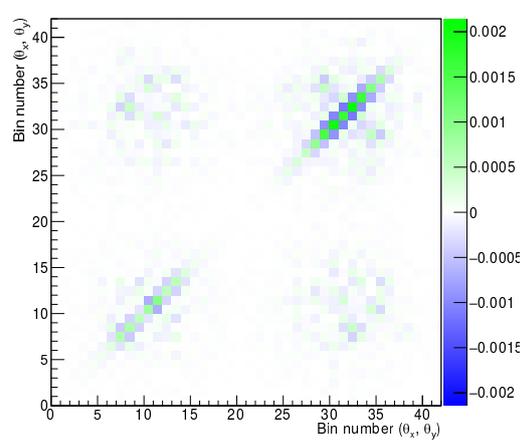
(a) Statistical errors



(b) Alignment shift of ± 1 mm

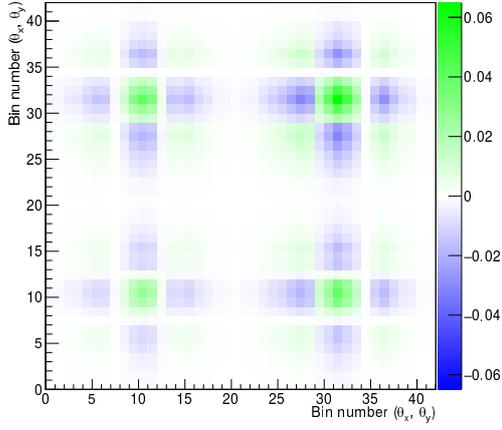


(c) Alignment rotation of ± 1 mrad

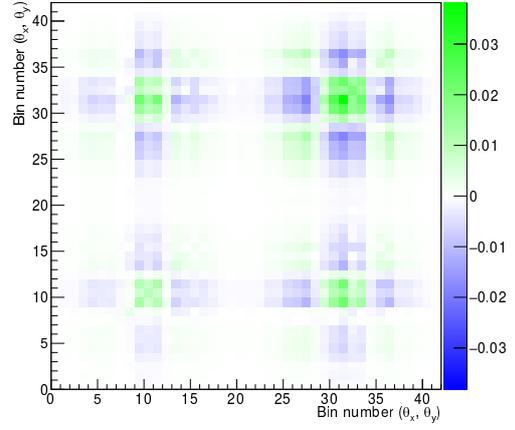


(d) Momentum-dependent shift correction

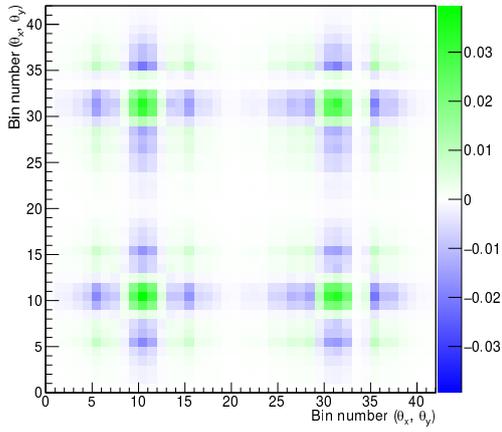
Figure A.9: Statistical and alignment covariance matrices for the 140 MeV/c LiH dataset



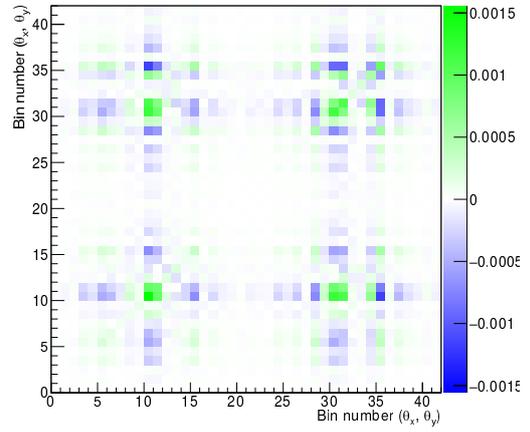
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

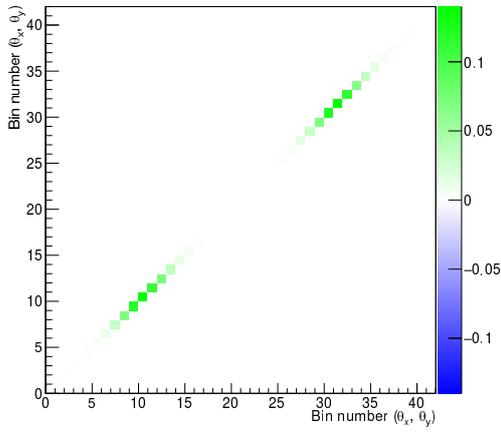


(c) Downstream magnetic field uncertainty

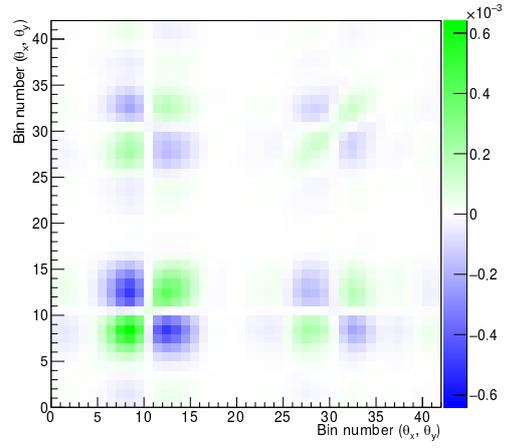


(d) Focus coils magnetic field uncertainty

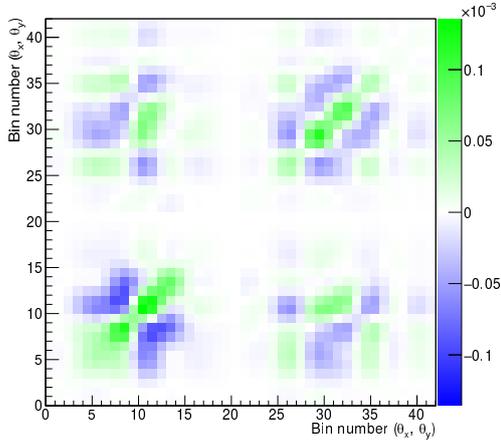
Figure A.10: Momentum selection and field uncertainty covariance matrices for the 140 MeV/c LiH dataset



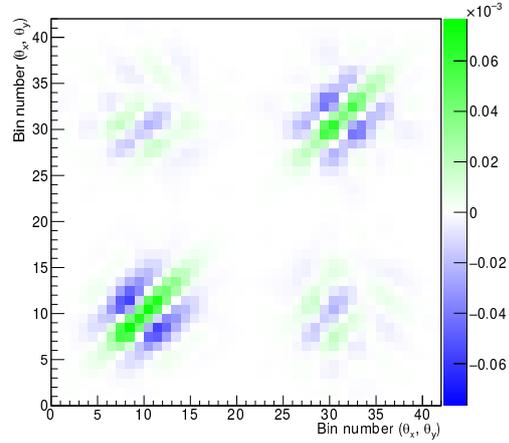
(a) Statistical errors



(b) Alignment shift of ± 1 mm

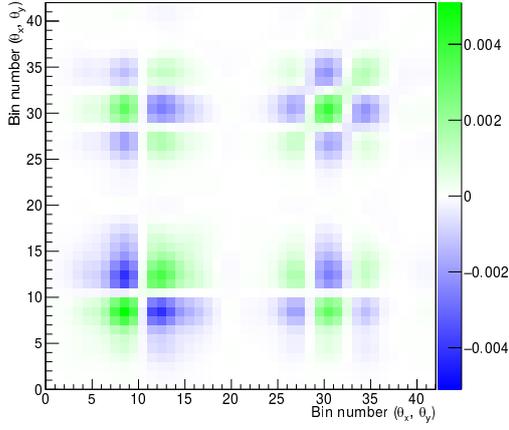


(c) Alignment rotation of ± 1 mrad

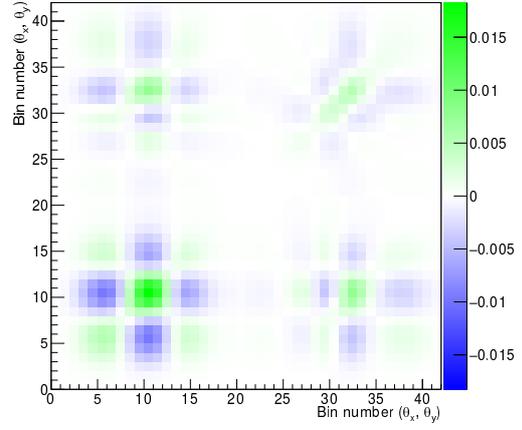


(d) Momentum-dependent shift correction

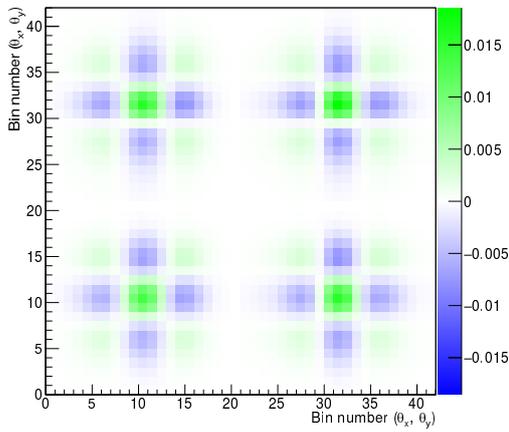
Figure A.11: Statistical and alignment covariance matrices for the 140 MeV/c convoluted dataset



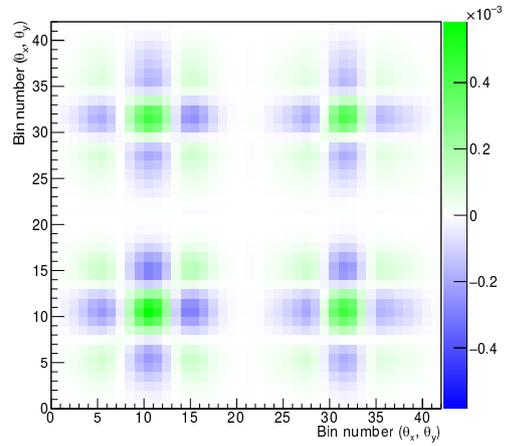
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

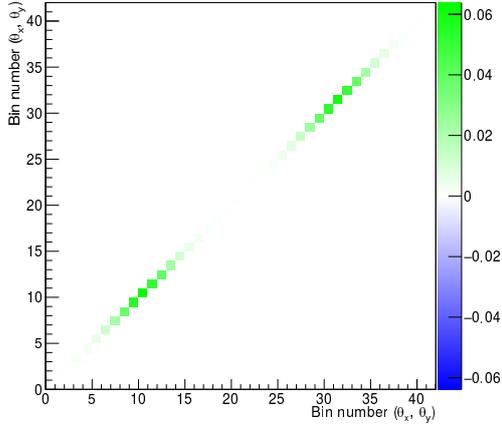


(c) Downstream magnetic field uncertainty

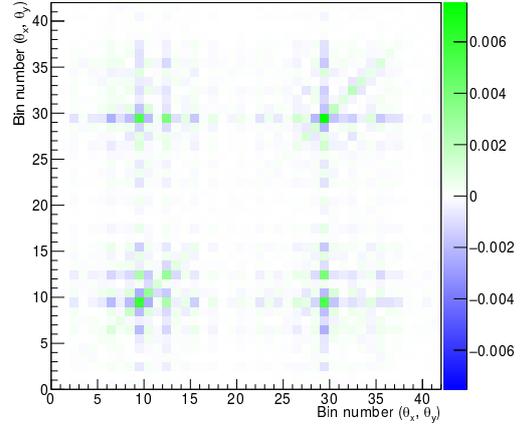


(d) Focus coils magnetic field uncertainty

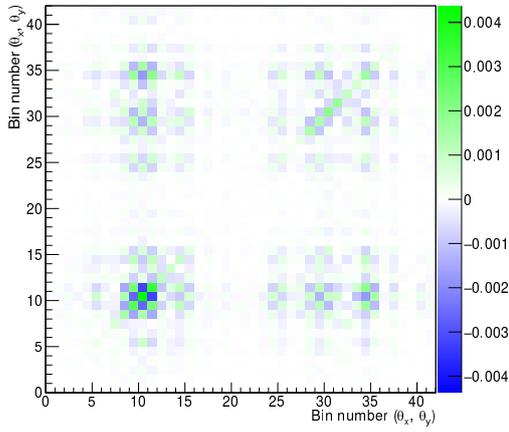
Figure A.12: Momentum selection and field uncertainty covariance matrices for the 140 MeV/c convoluted dataset



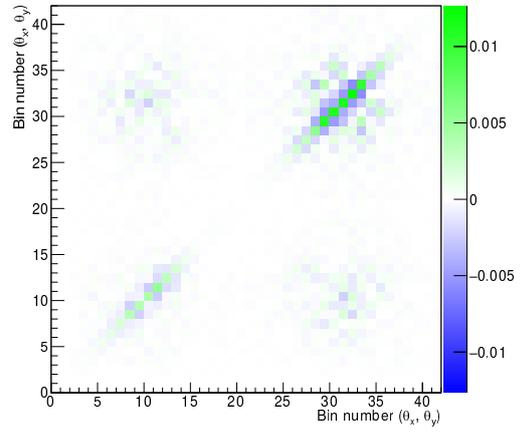
(a) Statistical errors



(b) Alignment shift of ± 1 mm

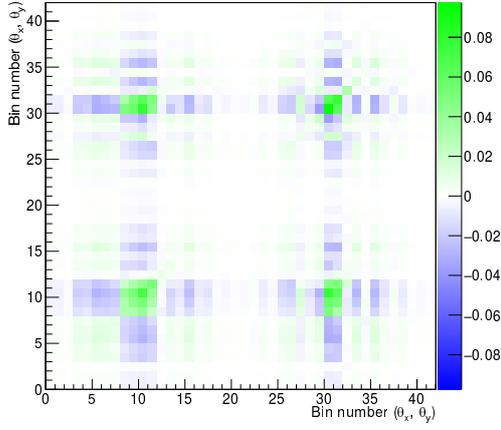


(c) Alignment rotation of ± 1 mrad

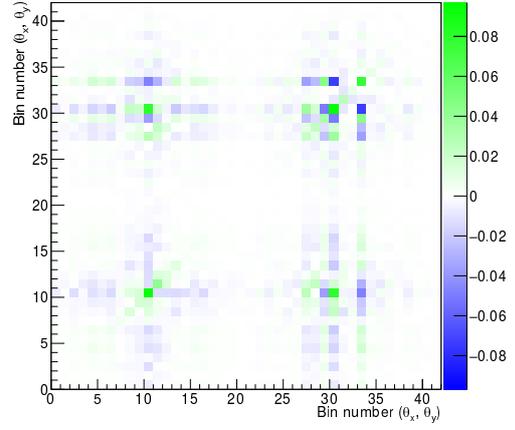


(d) Momentum-dependent shift correction

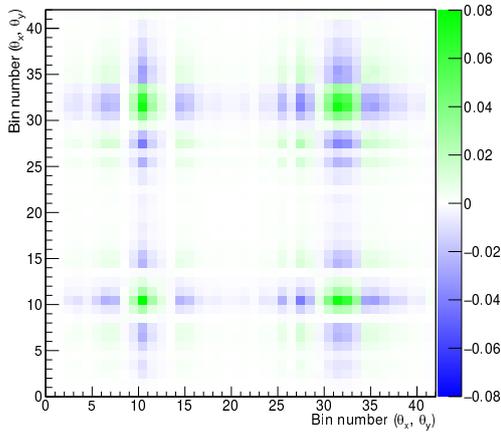
Figure A.13: Statistical and alignment covariance matrices for the 150 MeV/c LiH dataset



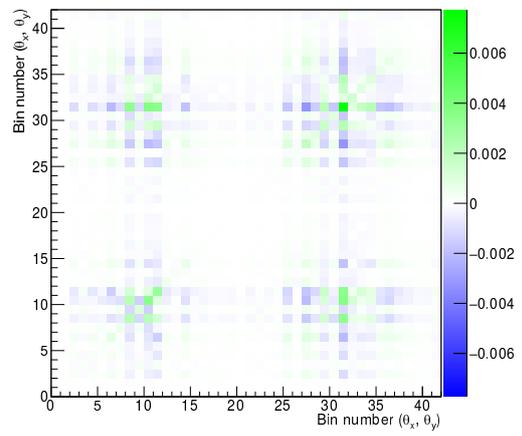
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

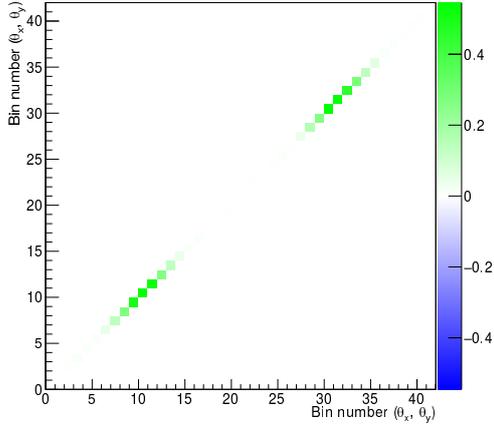


(c) Downstream magnetic field uncertainty

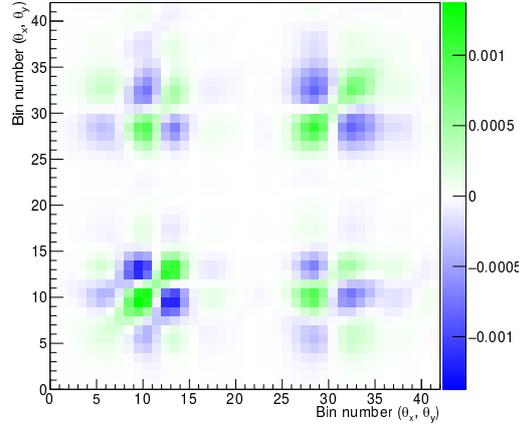


(d) Focus coils magnetic field uncertainty

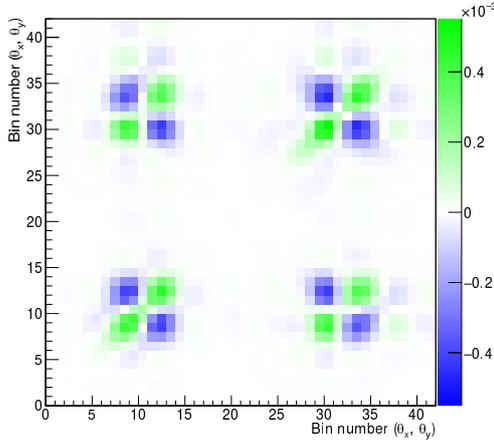
Figure A.14: Momentum selection and field uncertainty covariance matrices for the 150 MeV/c LiH dataset



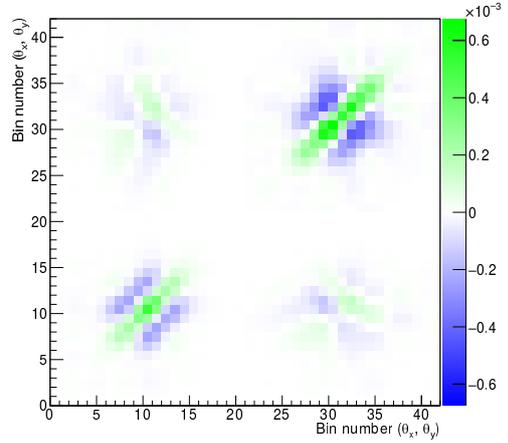
(a) Statistical errors



(b) Alignment shift of ± 1 mm

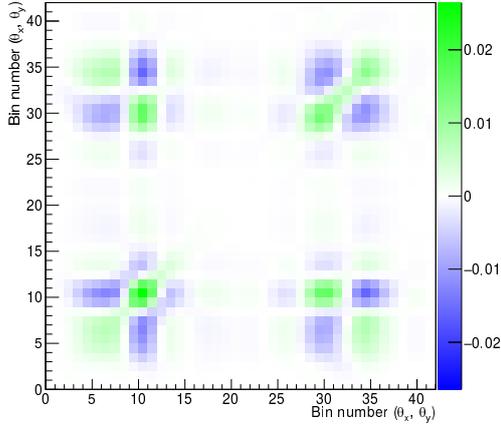


(c) Alignment rotation of ± 1 mrad

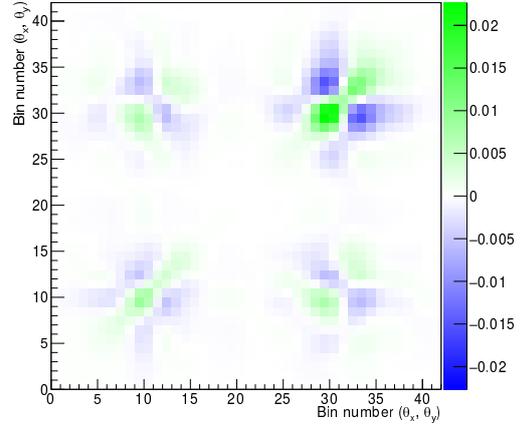


(d) Momentum-dependent shift correction

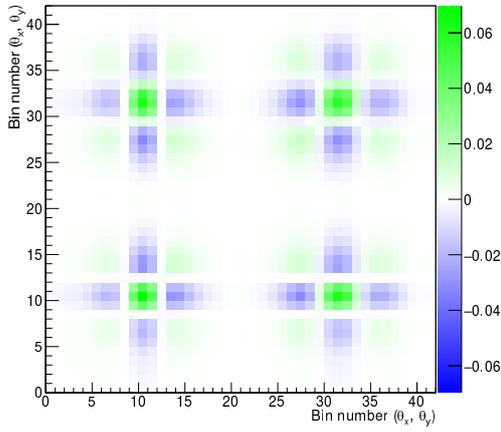
Figure A.15: Statistical and alignment covariance matrices for the 150 MeV/c convoluted dataset



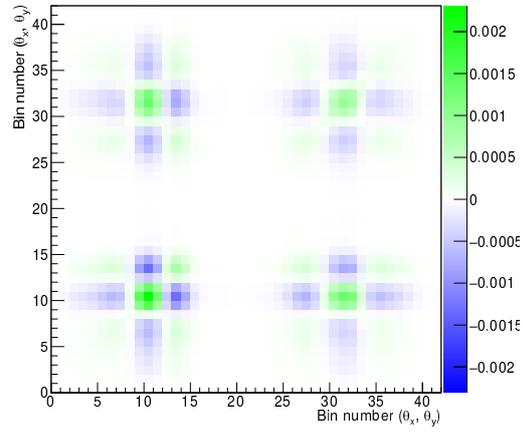
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

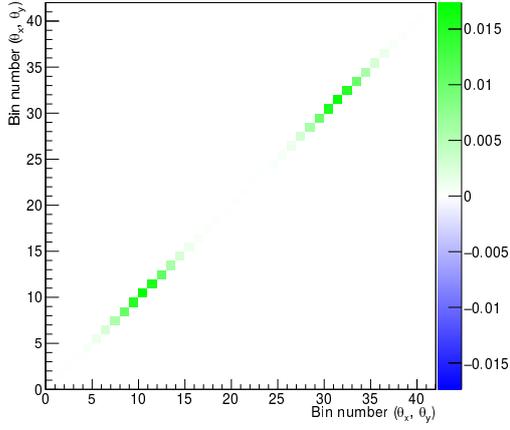


(c) Downstream magnetic field uncertainty

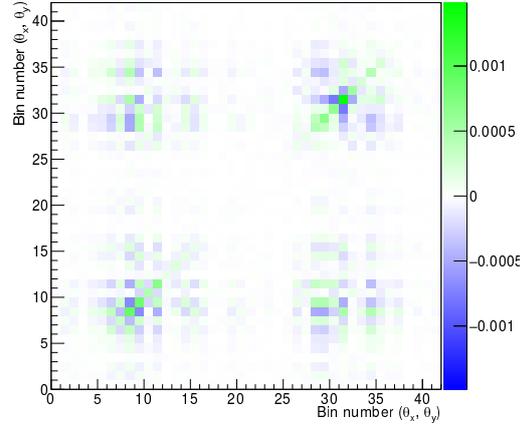


(d) Focus coils magnetic field uncertainty

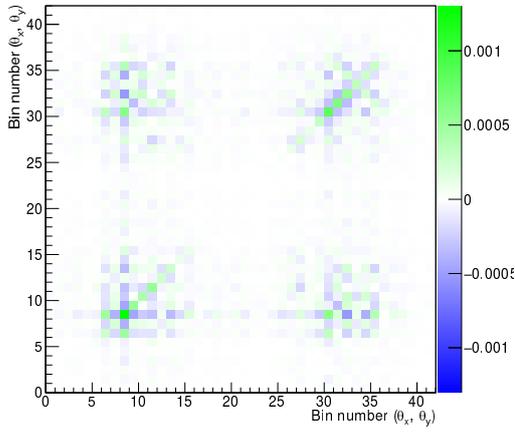
Figure A.16: Momentum selection and field uncertainty covariance matrices for the 150 MeV/c convoluted dataset



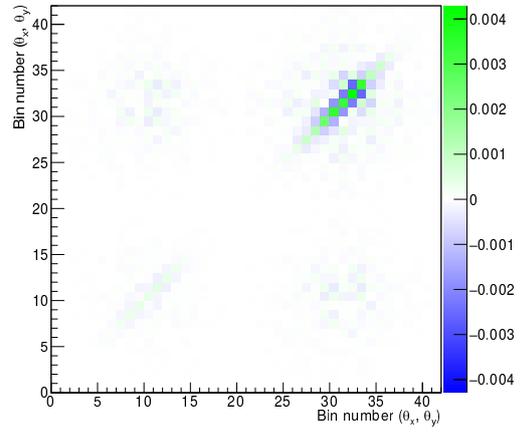
(a) Statistical errors



(b) Alignment shift of ± 1 mm

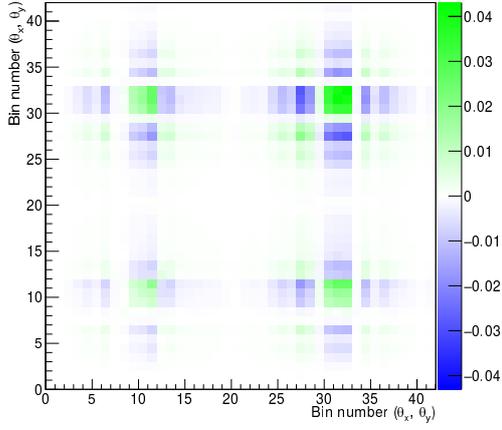


(c) Alignment rotation of ± 1 mrad

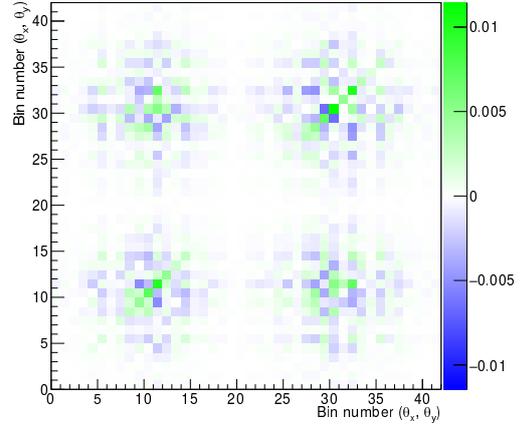


(d) Momentum-dependent shift correction

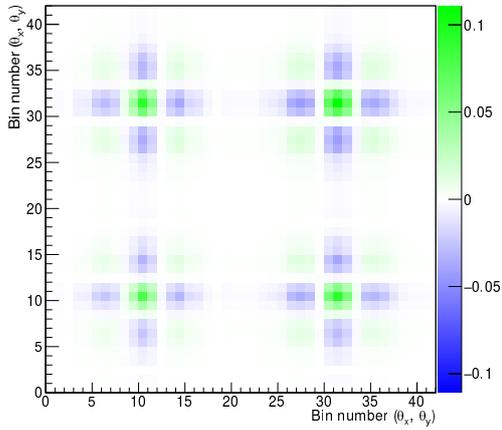
Figure A.17: Statistical and alignment covariance matrices for the 160 MeV/c LiH dataset



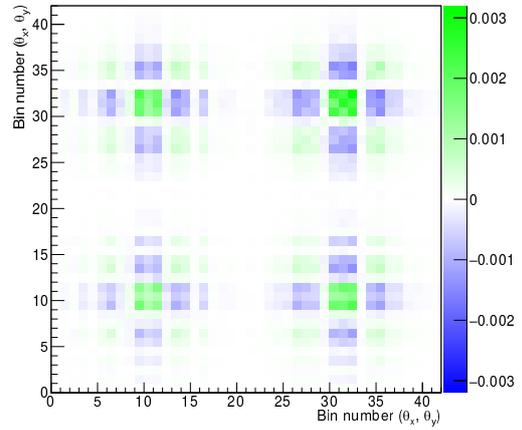
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

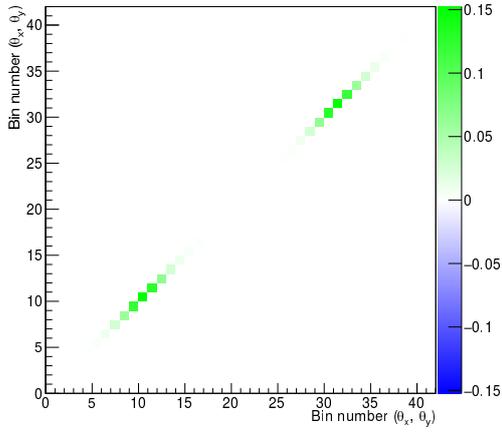


(c) Downstream magnetic field uncertainty

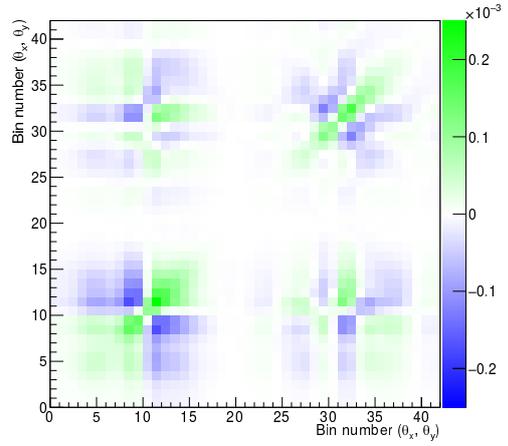


(d) Focus coils magnetic field uncertainty

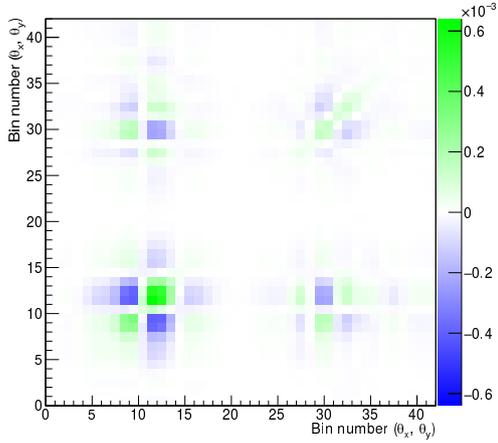
Figure A.18: Momentum selection and field uncertainty covariance matrices for the 160 MeV/c LiH dataset



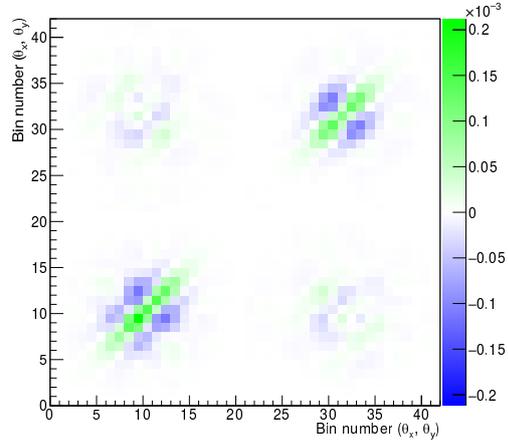
(a) Statistical errors



(b) Alignment shift of ± 1 mm

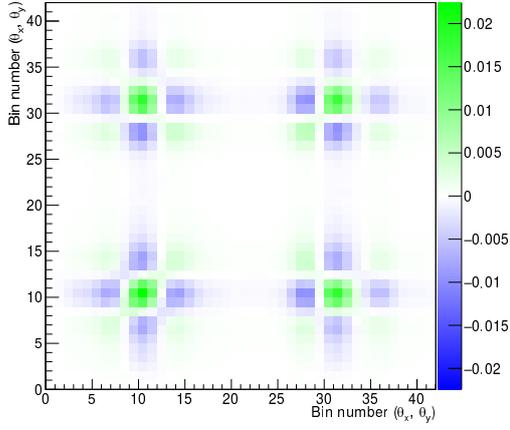


(c) Alignment rotation of ± 1 mrad

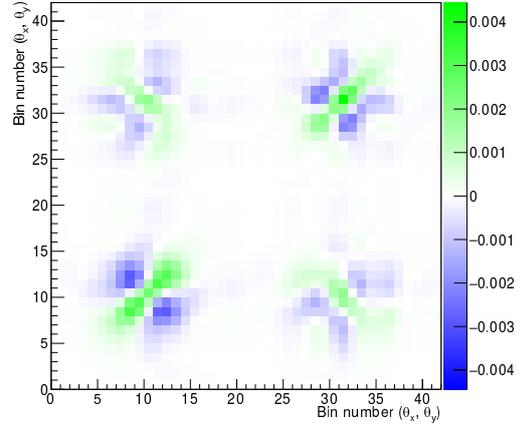


(d) Momentum-dependent shift correction

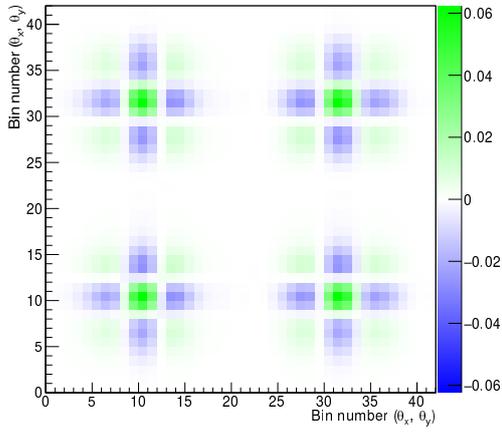
Figure A.19: Statistical and alignment covariance matrices for the 160 MeV/c convoluted dataset



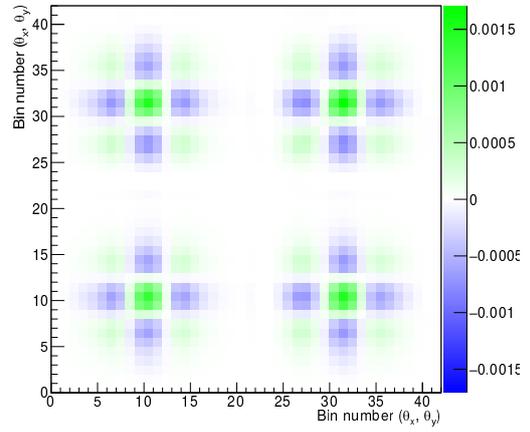
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

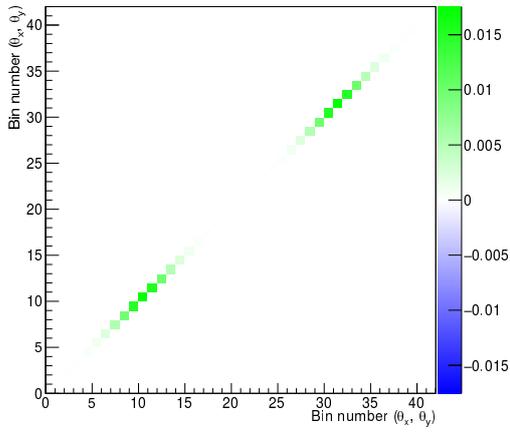


(c) Downstream magnetic field uncertainty

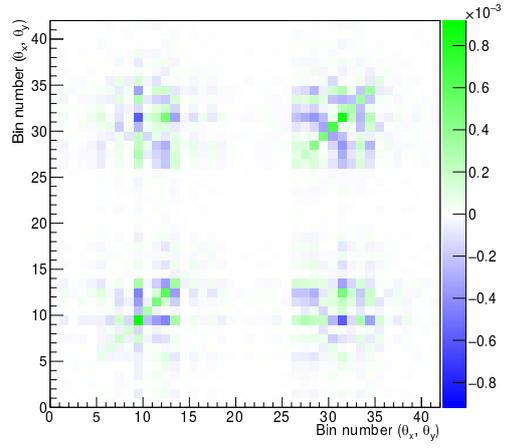


(d) Focus coils magnetic field uncertainty

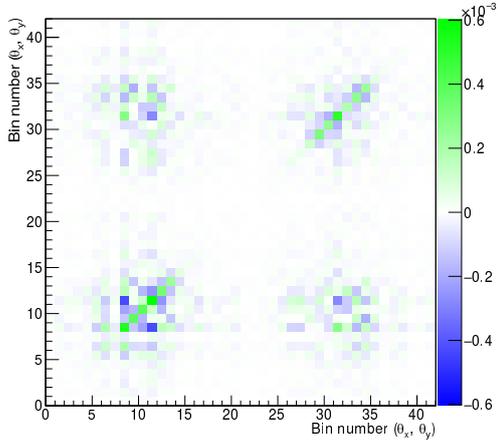
Figure A.20: Momentum selection and field uncertainty covariance matrices for the 160 MeV/c convoluted dataset



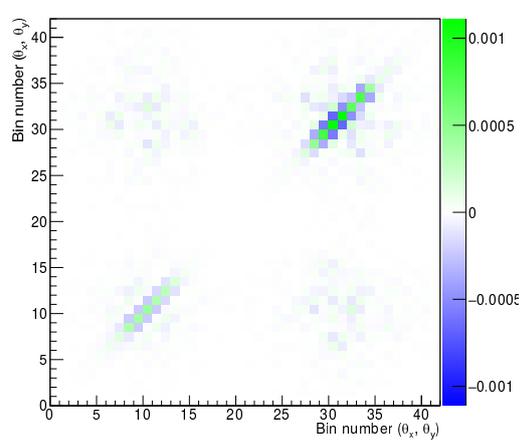
(a) Statistical errors



(b) Alignment shift of ± 1 mm

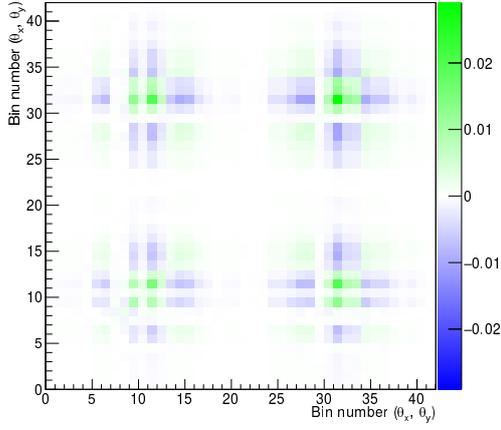


(c) Alignment rotation of ± 1 mrad

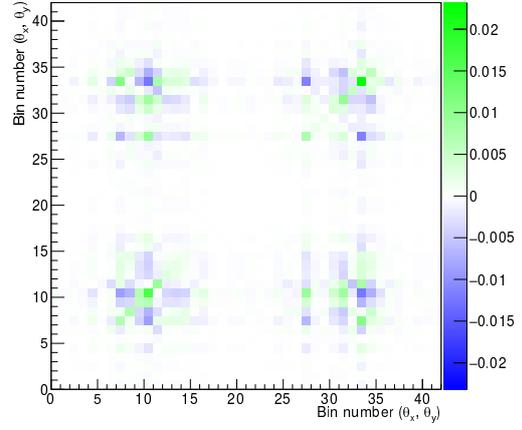


(d) Momentum-dependent shift correction

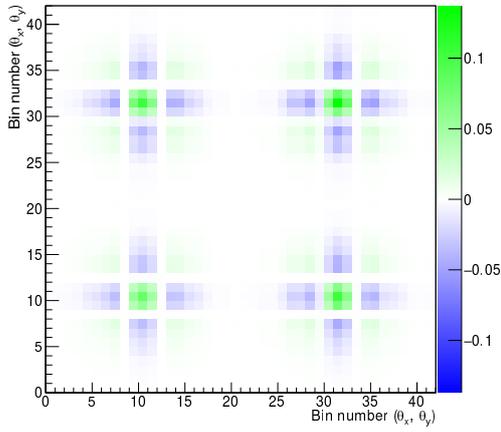
Figure A.21: Statistical and alignment covariance matrices for the 170 MeV/c LiH dataset



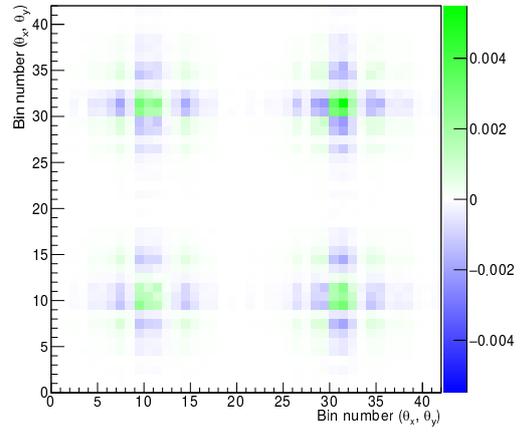
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

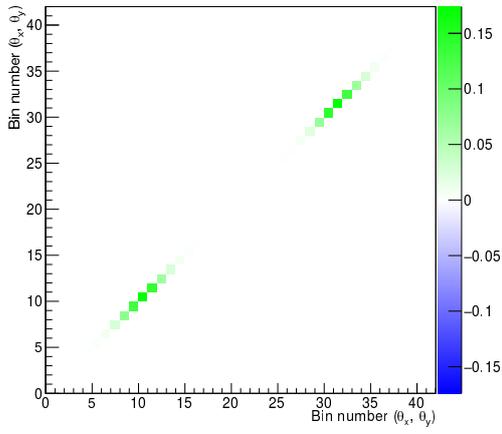


(c) Downstream magnetic field uncertainty

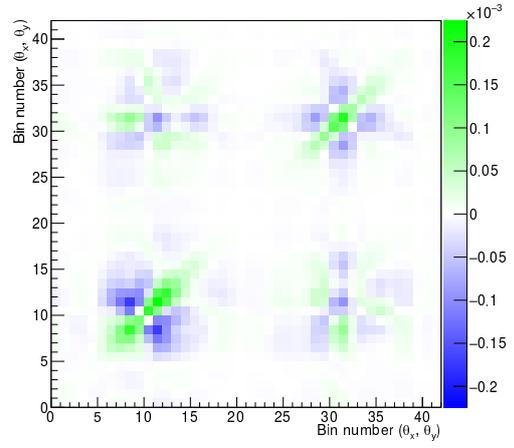


(d) Focus coils magnetic field uncertainty

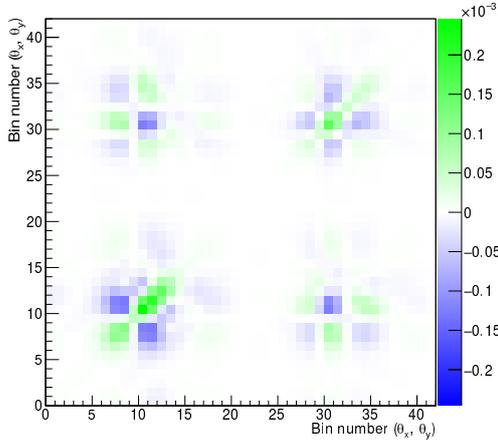
Figure A.22: Momentum selection and field uncertainty covariance matrices for the 170 MeV/c LiH dataset



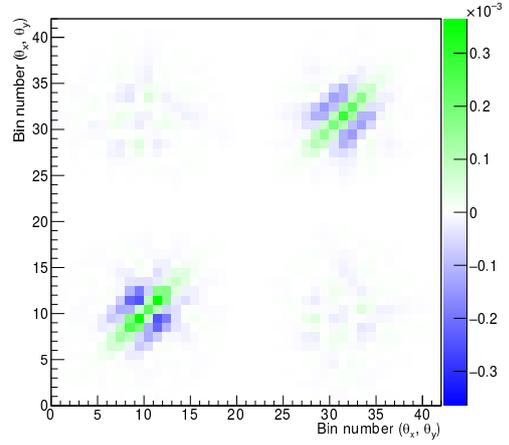
(a) Statistical errors



(b) Alignment shift of ± 1 mm

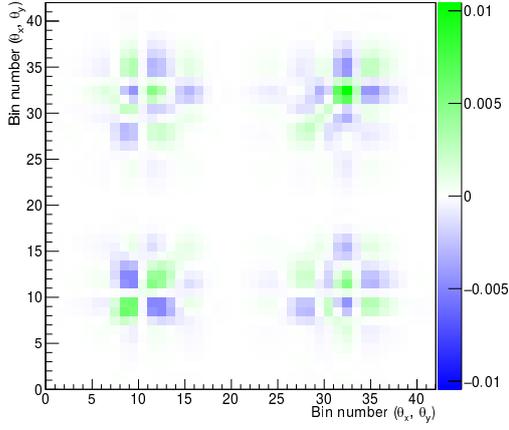


(c) Alignment rotation of ± 1 mrad

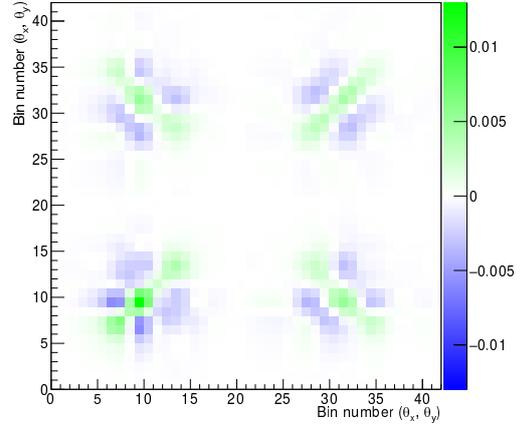


(d) Momentum-dependent shift correction

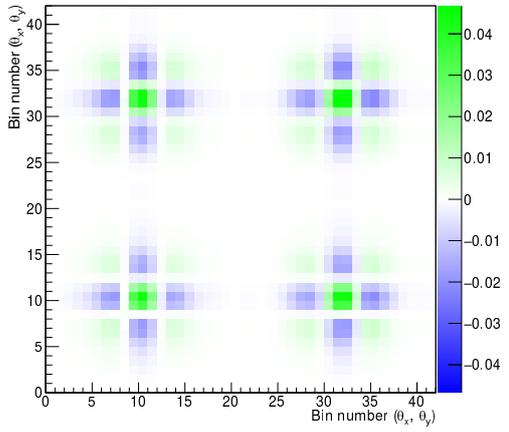
Figure A.23: Statistical and alignment covariance matrices for the 170 MeV/c convoluted dataset



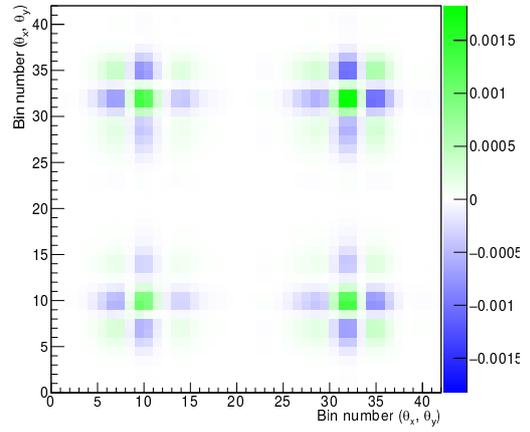
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

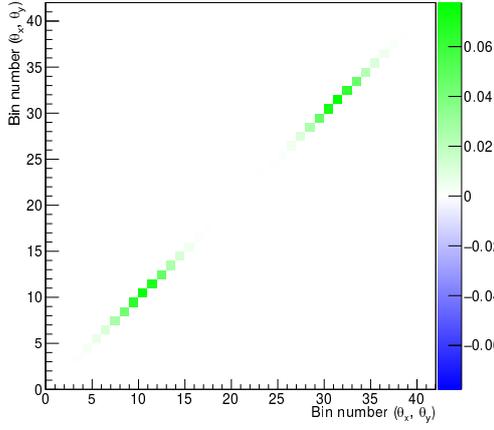


(c) Downstream magnetic field uncertainty

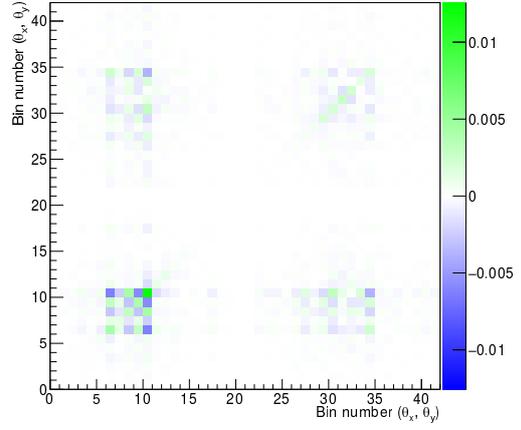


(d) Focus coils magnetic field uncertainty

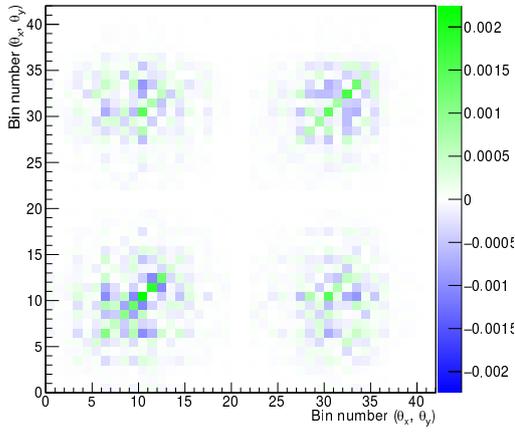
Figure A.24: Momentum selection and field uncertainty covariance matrices for the 170 MeV/c convoluted dataset



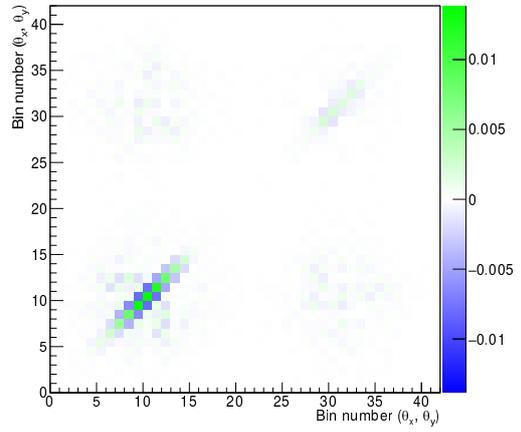
(a) Statistical errors



(b) Alignment shift of ± 1 mm

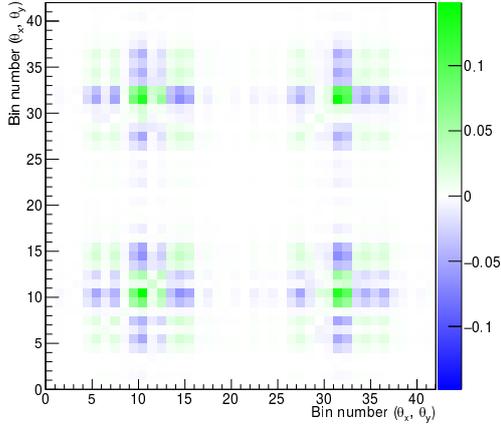


(c) Alignment rotation of ± 1 mrad

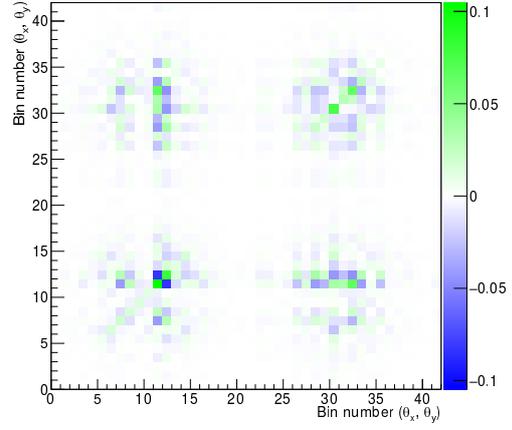


(d) Momentum-dependent shift correction

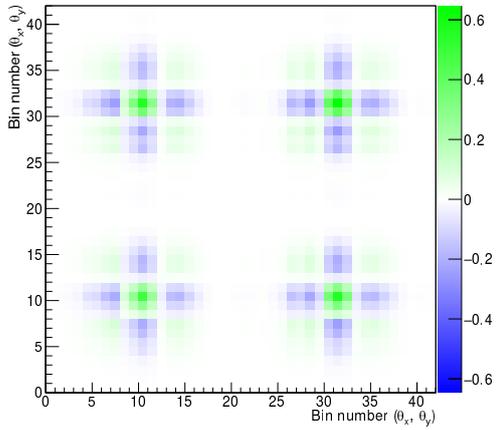
Figure A.25: Statistical and alignment covariance matrices for the 180 MeV/c LiH dataset



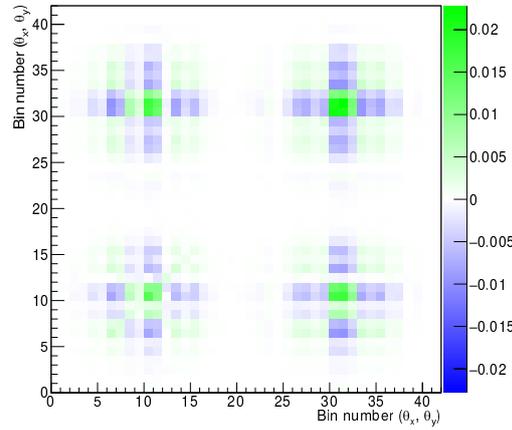
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

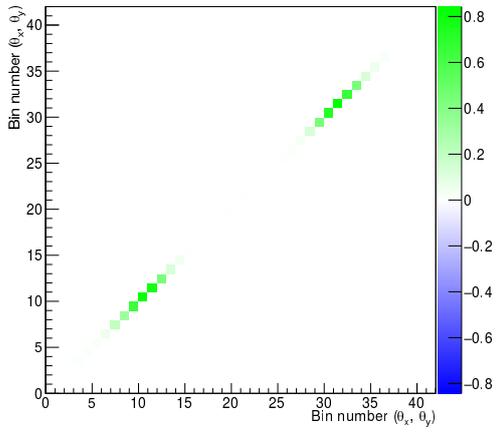


(c) Downstream magnetic field uncertainty

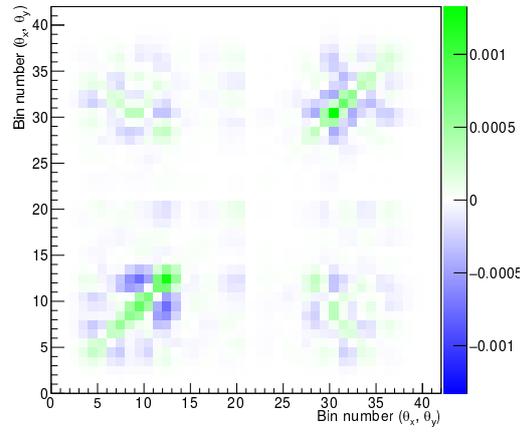


(d) Focus coils magnetic field uncertainty

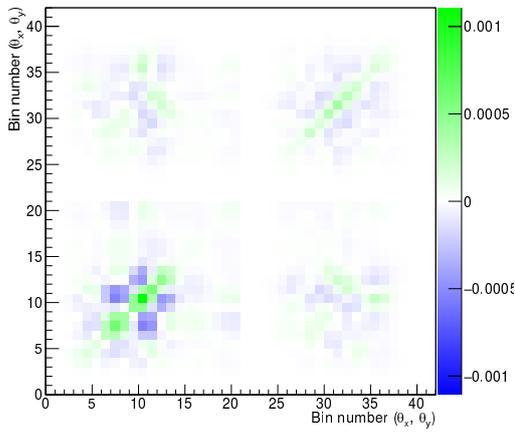
Figure A.26: Momentum selection and field uncertainty covariance matrices for the 180 MeV/c LiH dataset



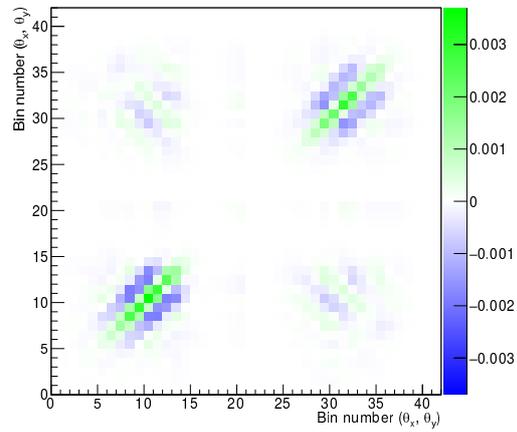
(a) Statistical errors



(b) Alignment shift of ± 1 mm

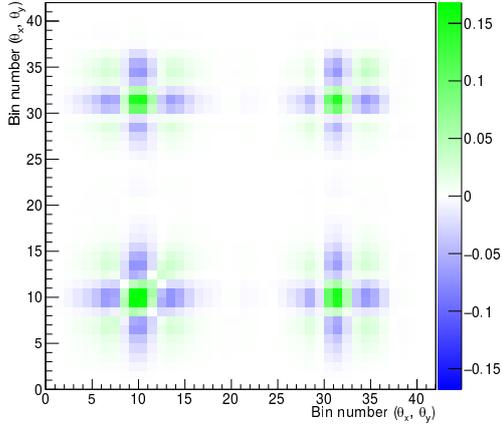


(c) Alignment rotation of ± 1 mrad

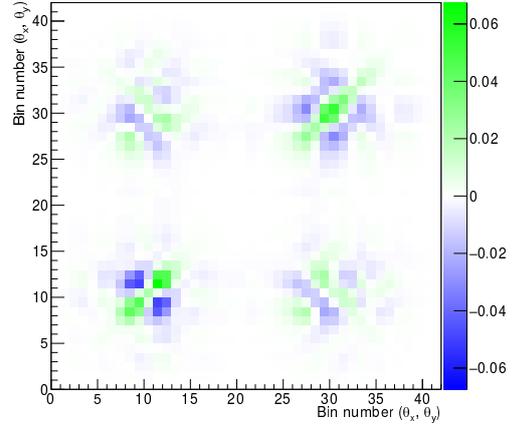


(d) Momentum-dependent shift correction

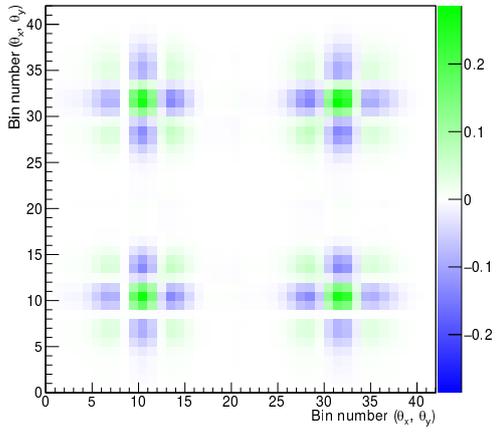
Figure A.27: Statistical and alignment covariance matrices for the 180 MeV/c convoluted dataset



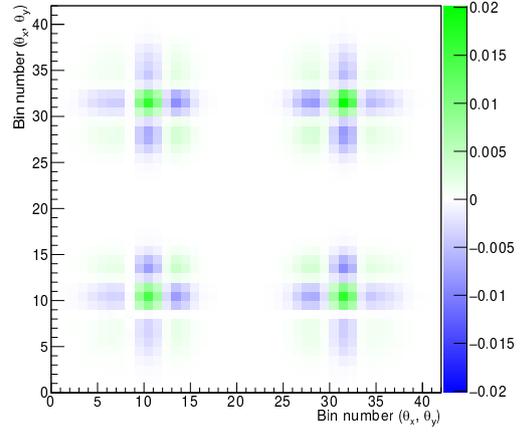
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

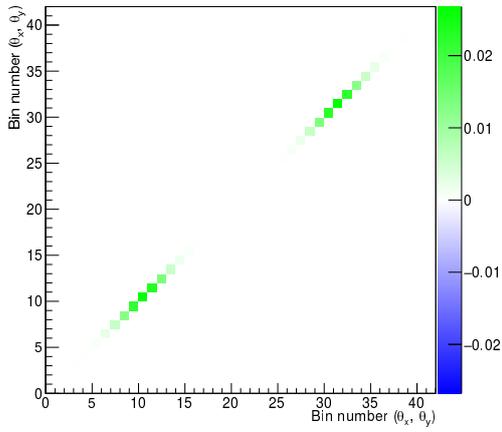


(c) Downstream magnetic field uncertainty

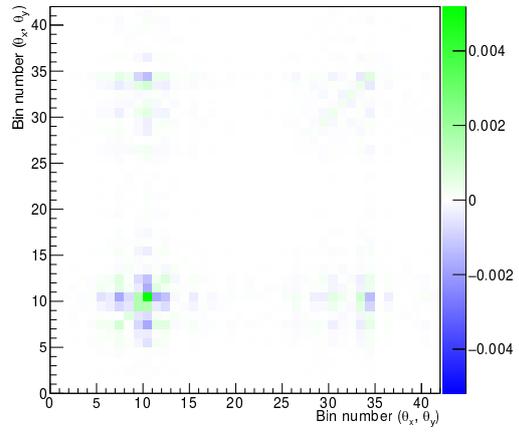


(d) Focus coils magnetic field uncertainty

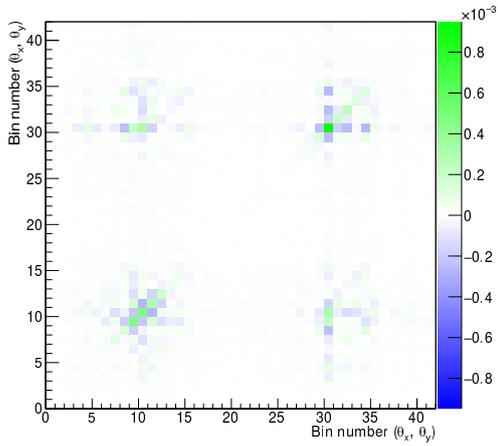
Figure A.28: Momentum selection and field uncertainty covariance matrices for the 180 MeV/c convoluted dataset



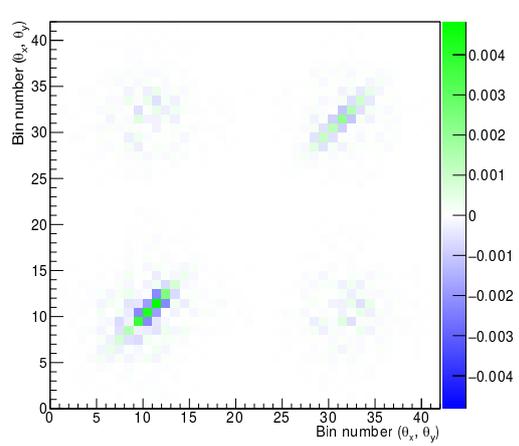
(a) Statistical errors



(b) Alignment shift of ± 1 mm

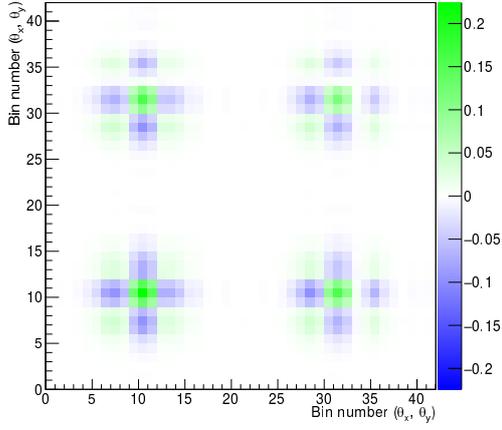


(c) Alignment rotation of ± 1 mrad

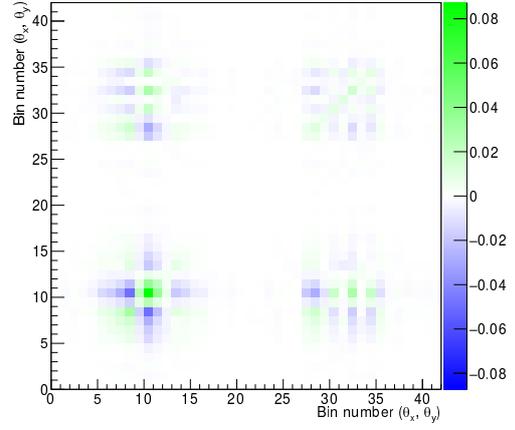


(d) Momentum-dependent shift correction

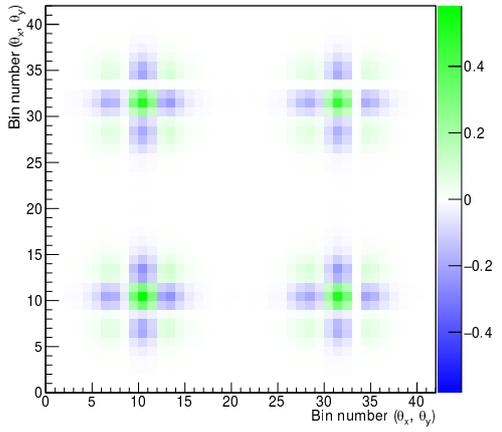
Figure A.29: Statistical and alignment covariance matrices for the 190 MeV/c LiH dataset



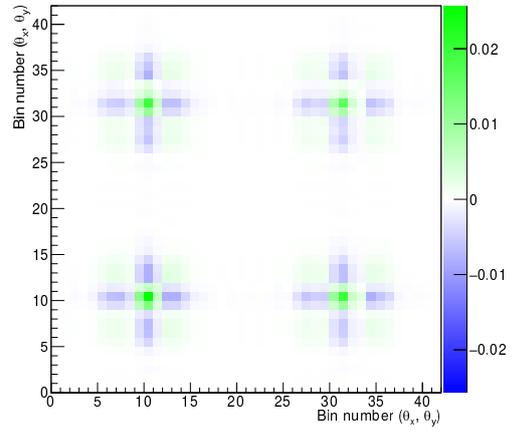
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

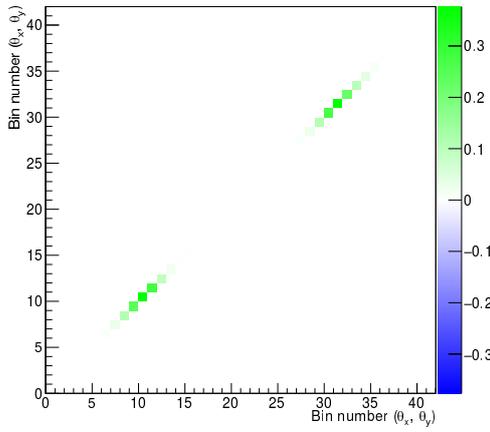


(c) Downstream magnetic field uncertainty

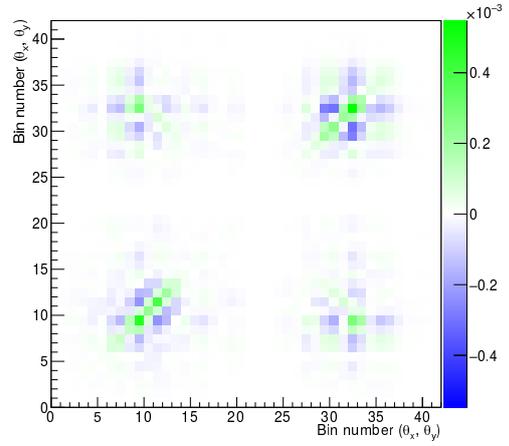


(d) Focus coils magnetic field uncertainty

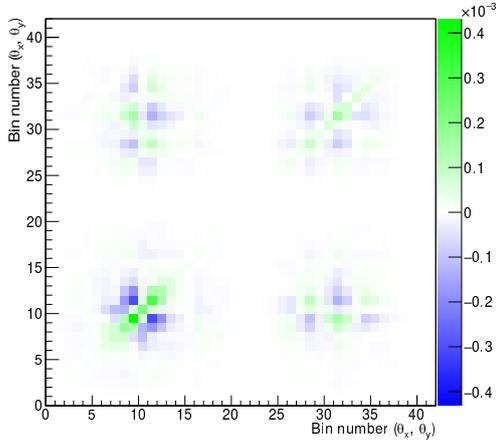
Figure A.30: Momentum selection and field uncertainty covariance matrices for the 190 MeV/c LiH dataset



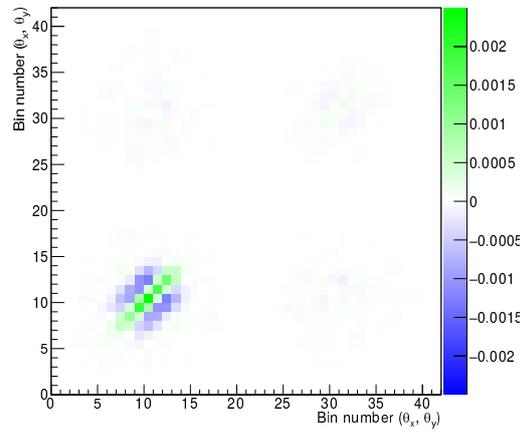
(a) Statistical errors



(b) Alignment shift of ± 1 mm

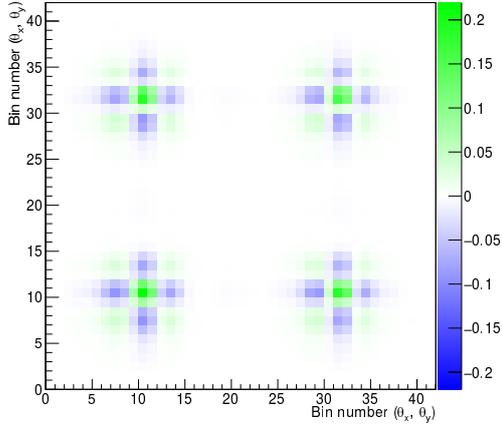


(c) Alignment rotation of ± 1 mrad

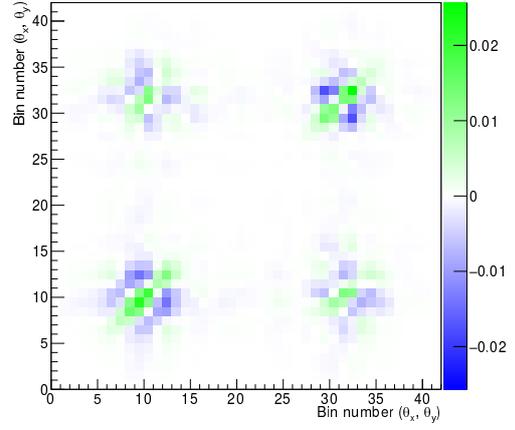


(d) Momentum-dependent shift correction

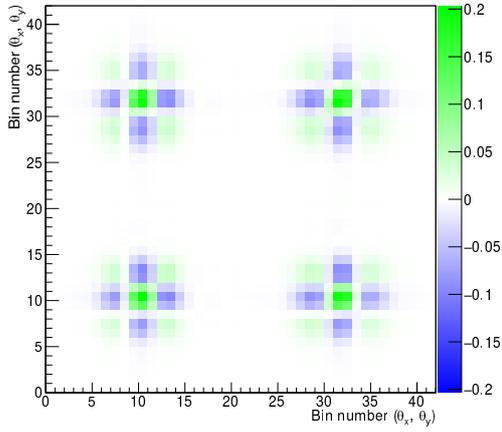
Figure A.31: Statistical and alignment covariance matrices for the 190 MeV/c convoluted dataset



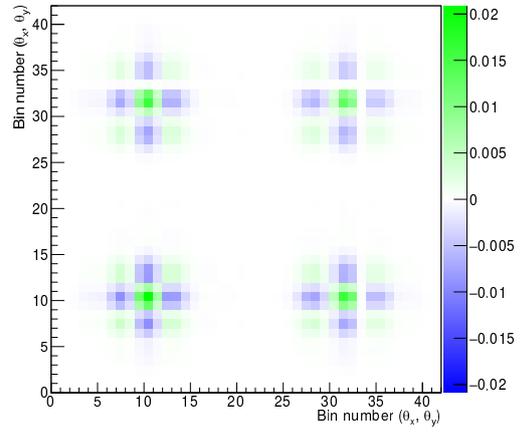
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

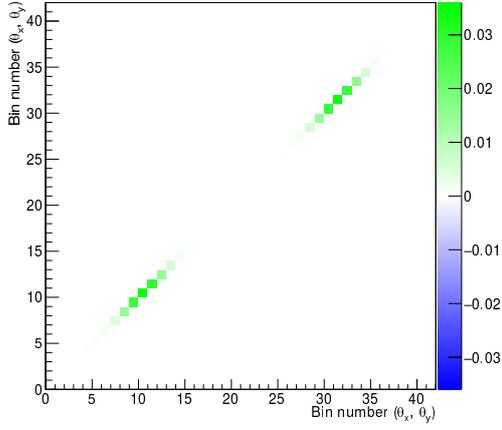


(c) Downstream magnetic field uncertainty

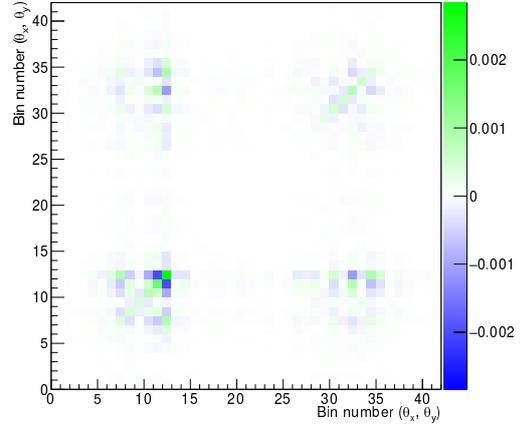


(d) Focus coils magnetic field uncertainty

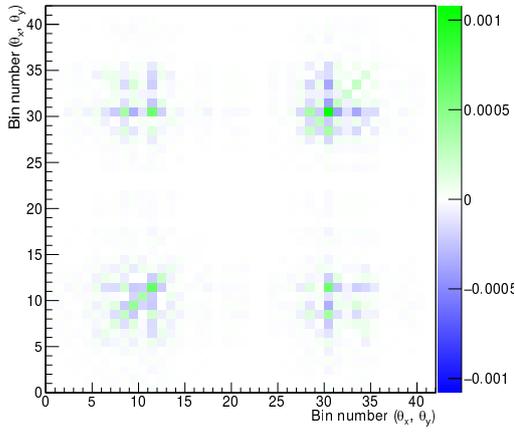
Figure A.32: Momentum selection and field uncertainty covariance matrices for the 190 MeV/c convoluted dataset



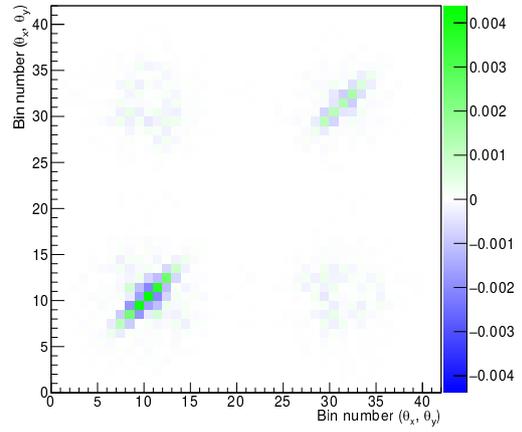
(a) Statistical errors



(b) Alignment shift of ± 1 mm

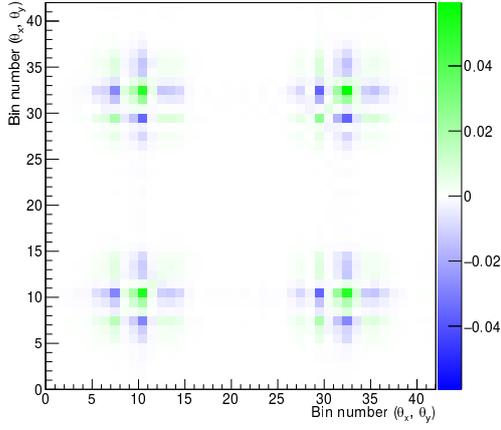


(c) Alignment rotation of ± 1 mrad

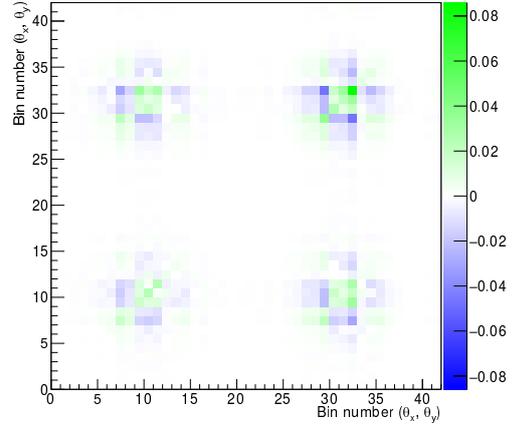


(d) Momentum-dependent shift correction

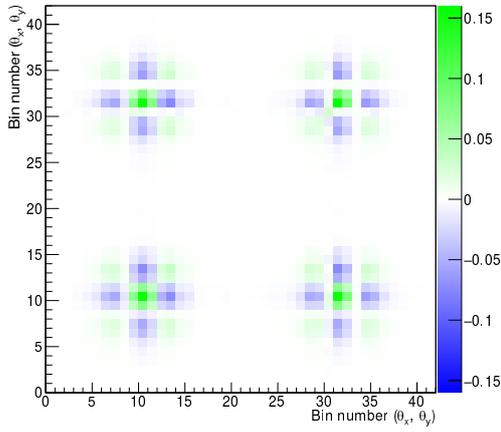
Figure A.33: Statistical and alignment covariance matrices for the 200 MeV/c LiH dataset



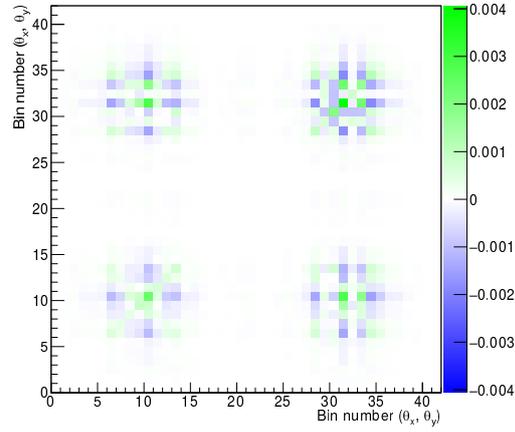
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

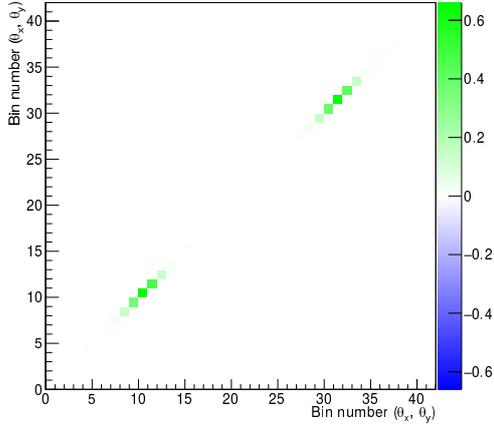


(c) Downstream magnetic field uncertainty

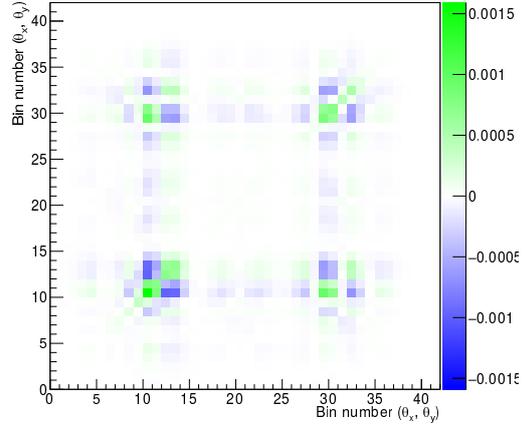


(d) Focus coils magnetic field uncertainty

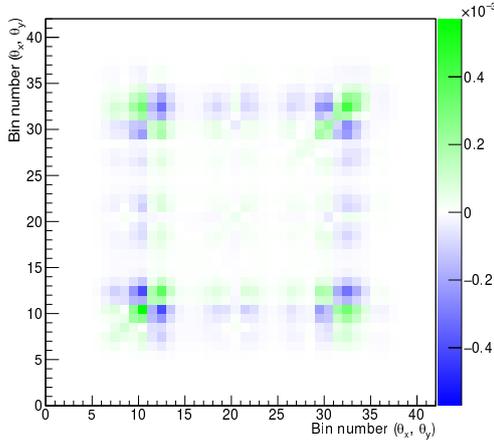
Figure A.34: Momentum selection and field uncertainty covariance matrices for the 200 MeV/c LiH dataset



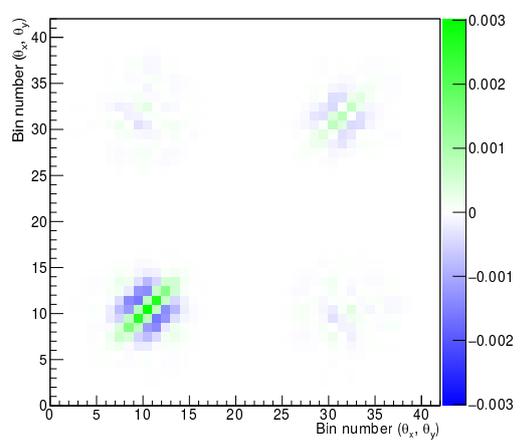
(a) Statistical errors



(b) Alignment shift of ± 1 mm

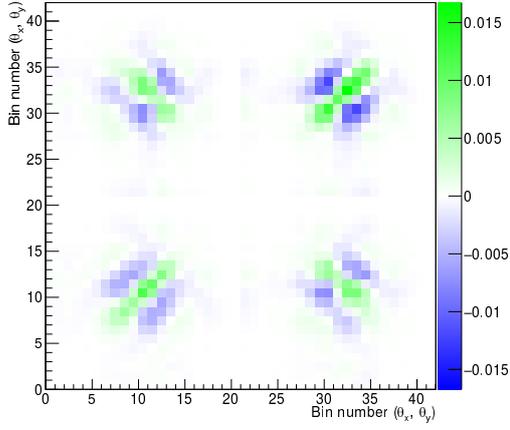


(c) Alignment rotation of ± 1 mrad

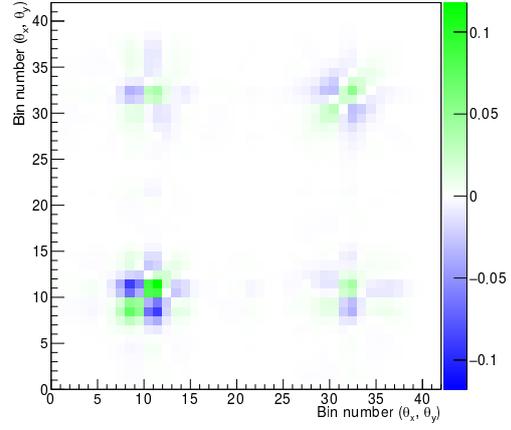


(d) Momentum-dependent shift correction

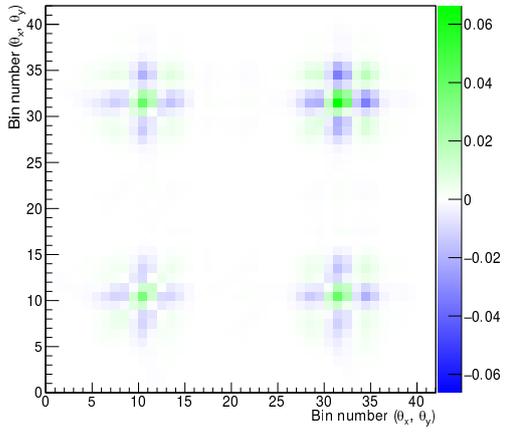
Figure A.35: Statistical and alignment covariance matrices for the 200 MeV/c convoluted dataset



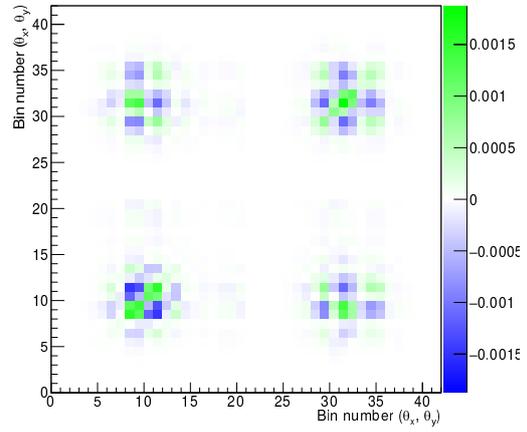
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

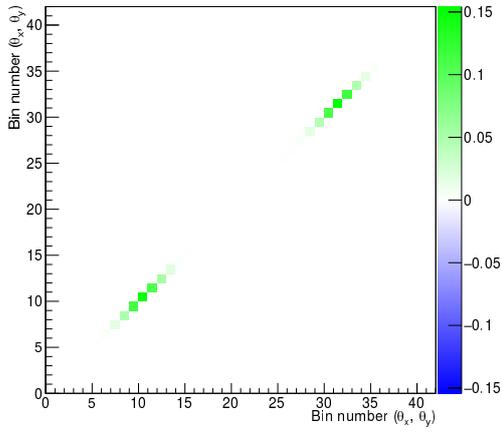


(c) Downstream magnetic field uncertainty

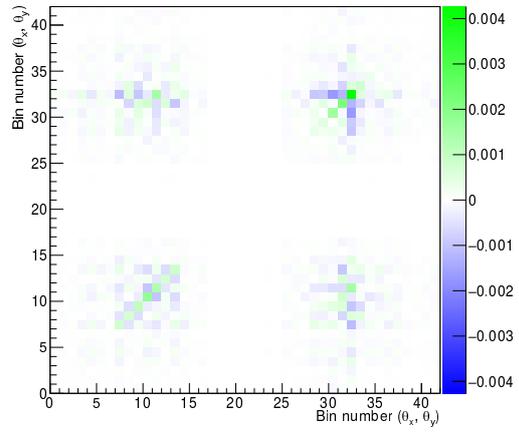


(d) Focus coils magnetic field uncertainty

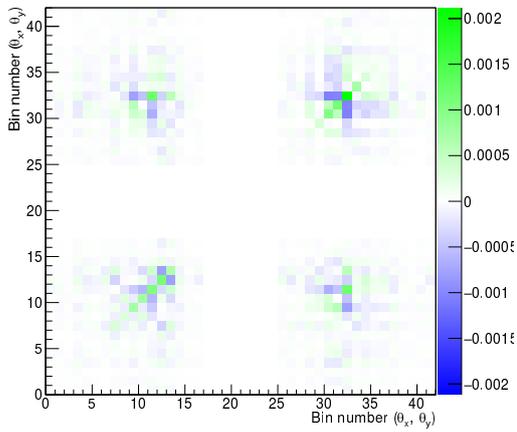
Figure A.36: Momentum selection and field uncertainty covariance matrices for the 200 MeV/c convoluted dataset



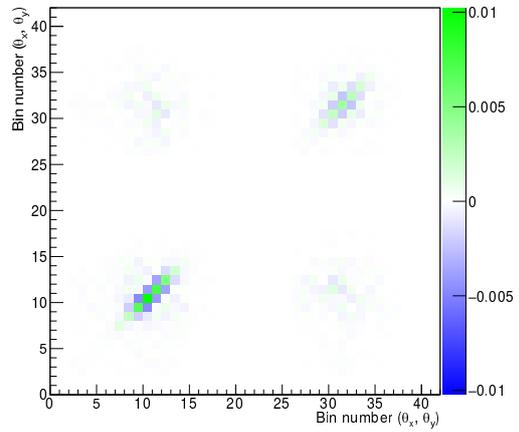
(a) Statistical errors



(b) Alignment shift of ± 1 mm

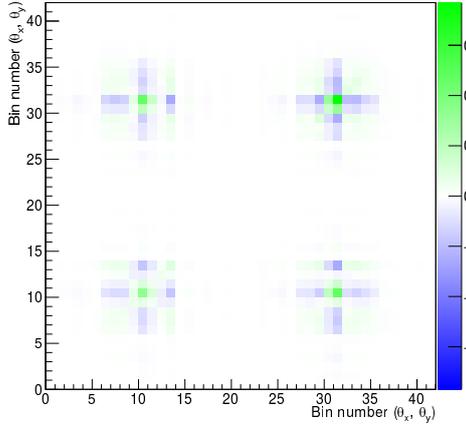


(c) Alignment rotation of ± 1 mrad

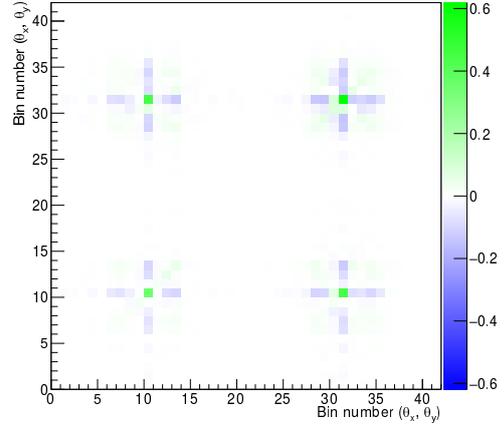


(d) Momentum-dependent shift correction

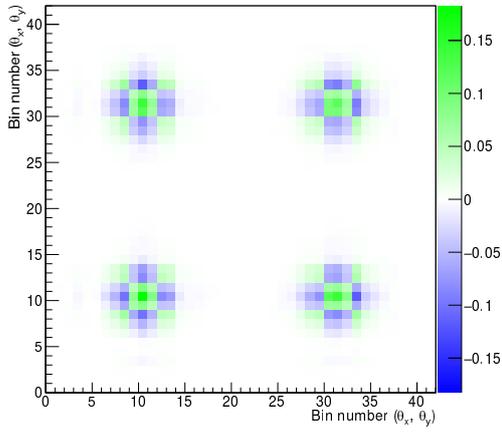
Figure A.37: Statistical and alignment covariance matrices for the 220 MeV/c LiH dataset



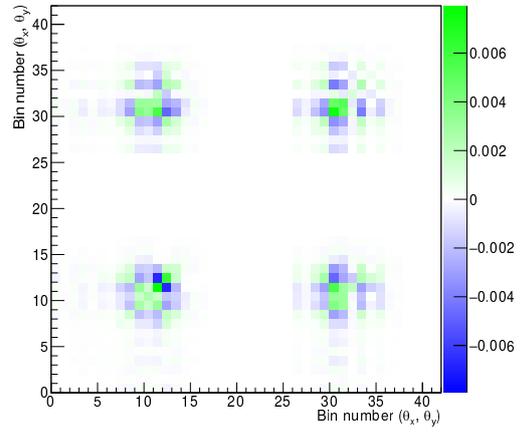
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

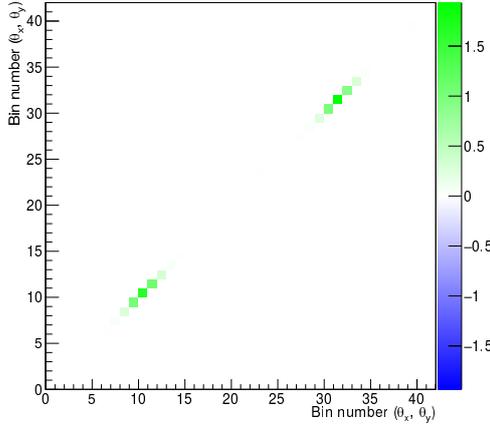


(c) Downstream magnetic field uncertainty

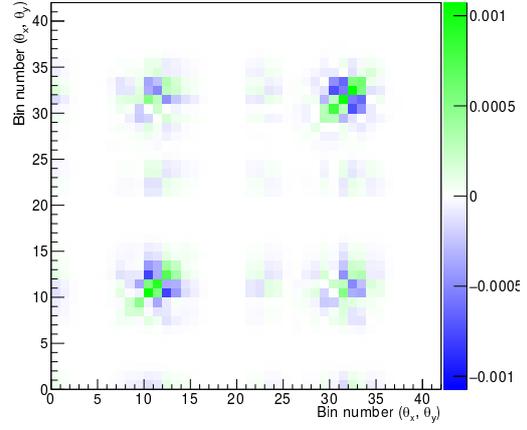


(d) Focus coils magnetic field uncertainty

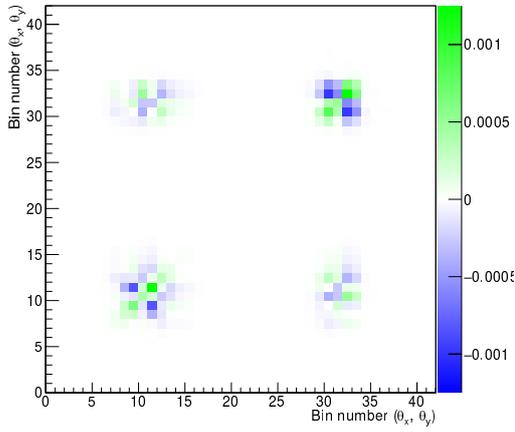
Figure A.38: Momentum selection and field uncertainty covariance matrices for the 220 MeV/c LiH dataset



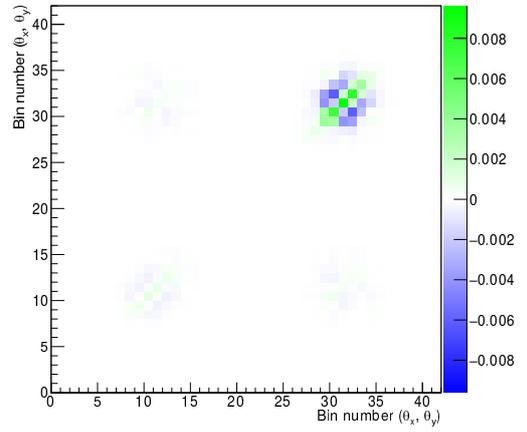
(a) Statistical errors



(b) Alignment shift of ± 1 mm

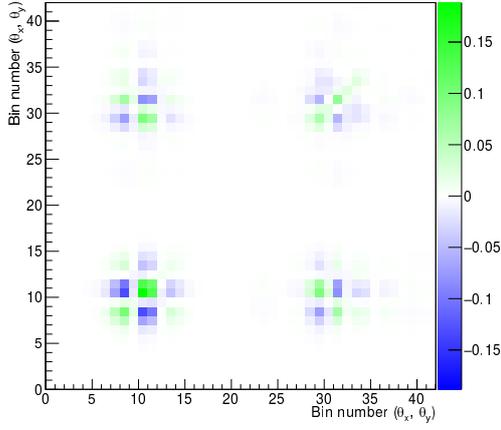


(c) Alignment rotation of ± 1 mrad

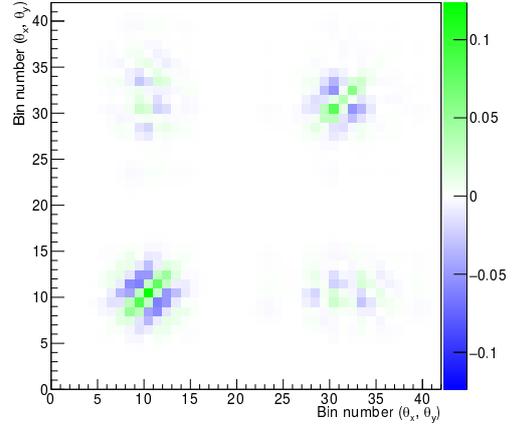


(d) Momentum-dependent shift correction

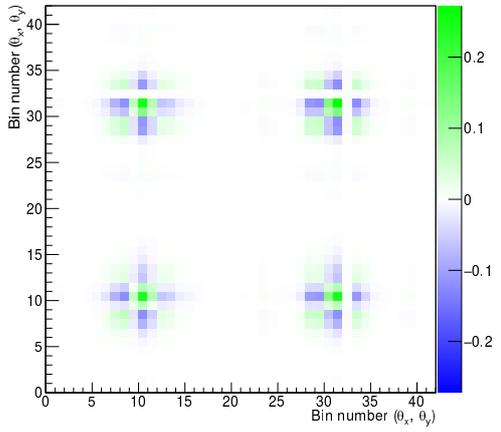
Figure A.39: Statistical and alignment covariance matrices for the 220 MeV/c convoluted dataset



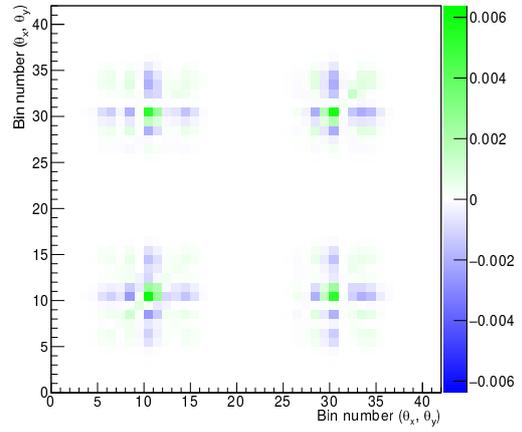
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

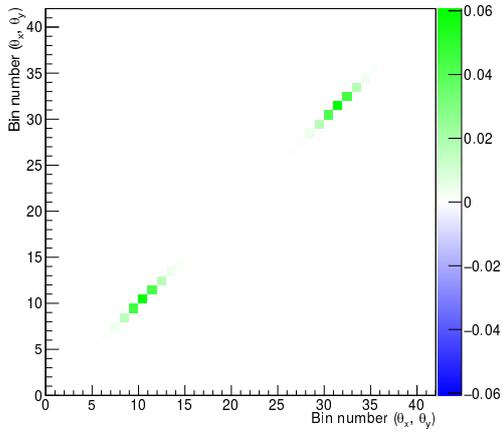


(c) Downstream magnetic field uncertainty

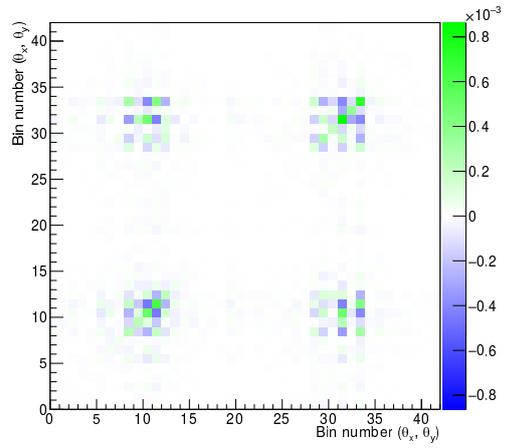


(d) Focus coils magnetic field uncertainty

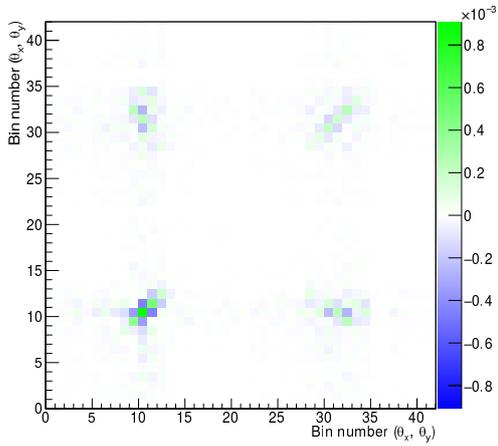
Figure A.40: Momentum selection and field uncertainty covariance matrices for the 220 MeV/c convoluted dataset



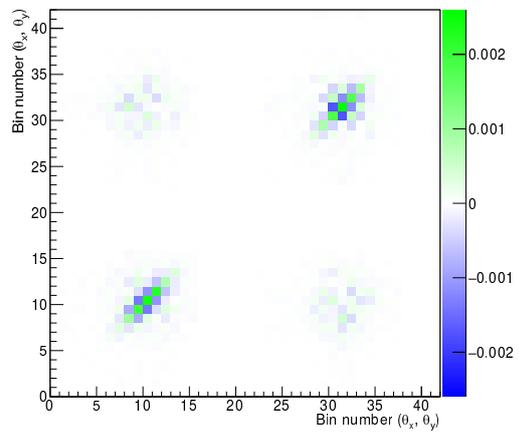
(a) Statistical errors



(b) Alignment shift of ± 1 mm

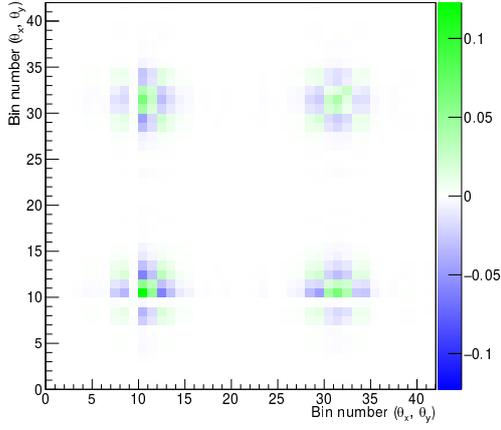


(c) Alignment rotation of ± 1 mrad

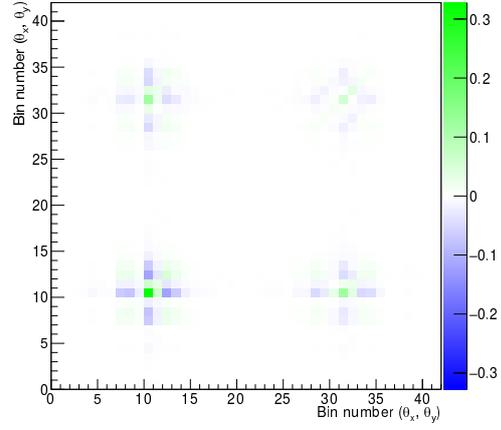


(d) Momentum-dependent shift correction

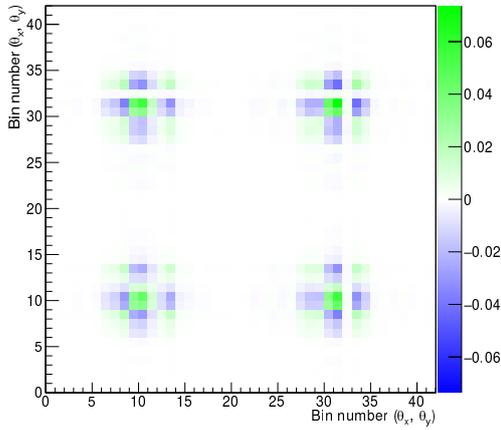
Figure A.41: Statistical and alignment covariance matrices for the 230 MeV/c LiH dataset



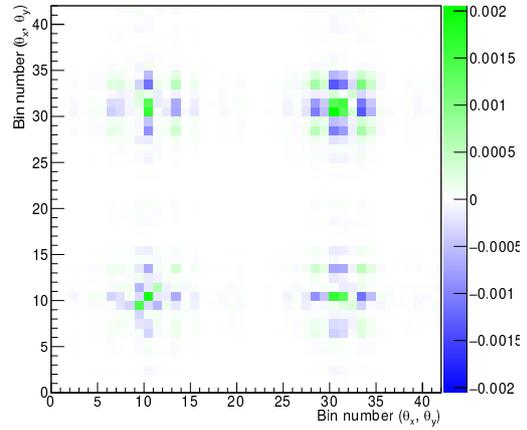
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

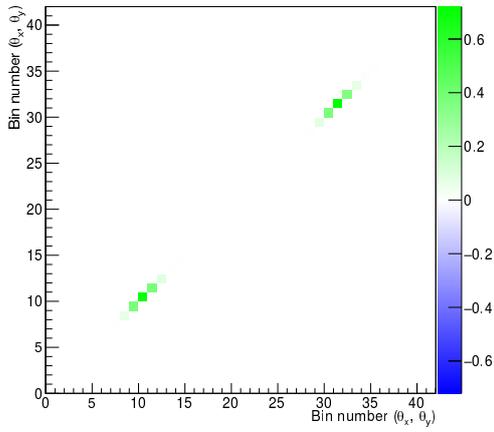


(c) Downstream magnetic field uncertainty

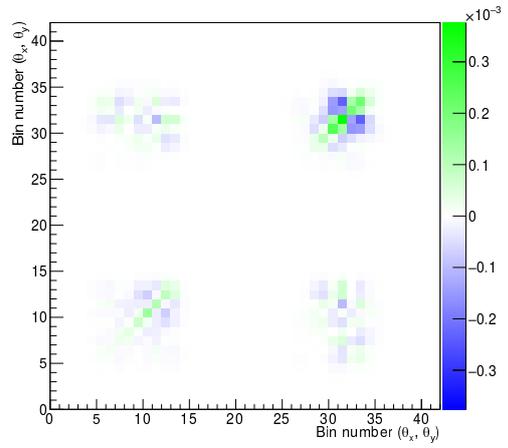


(d) Focus coils magnetic field uncertainty

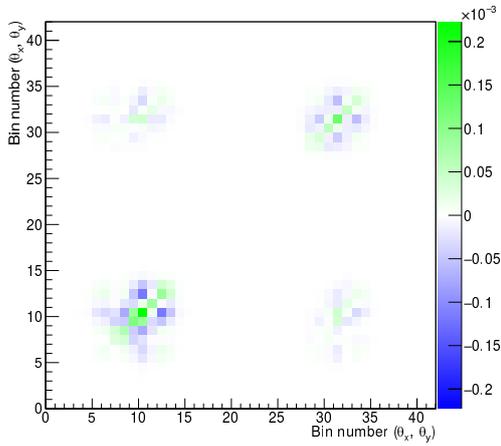
Figure A.42: Momentum selection and field uncertainty covariance matrices for the 230 MeV/c LiH dataset



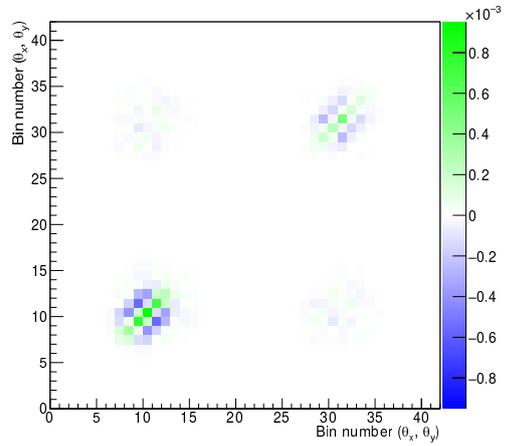
(a) Statistical errors



(b) Alignment shift of ± 1 mm

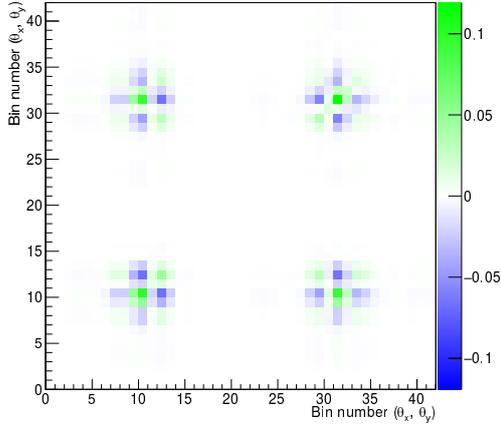


(c) Alignment rotation of ± 1 mrad

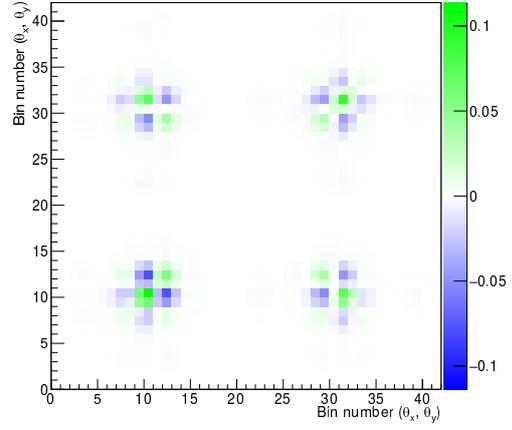


(d) Momentum-dependent shift correction

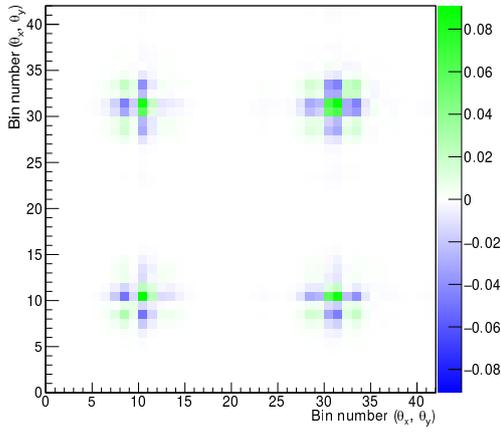
Figure A.43: Statistical and alignment covariance matrices for the 230 MeV/c convoluted dataset



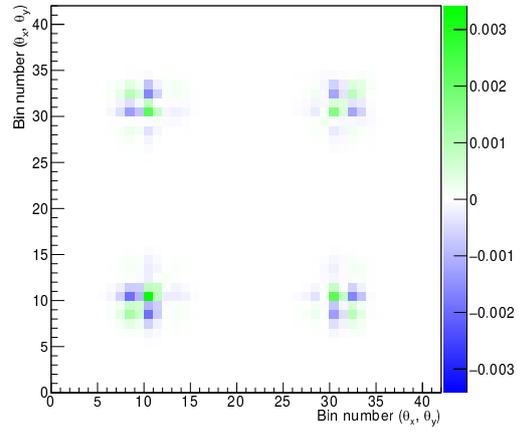
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

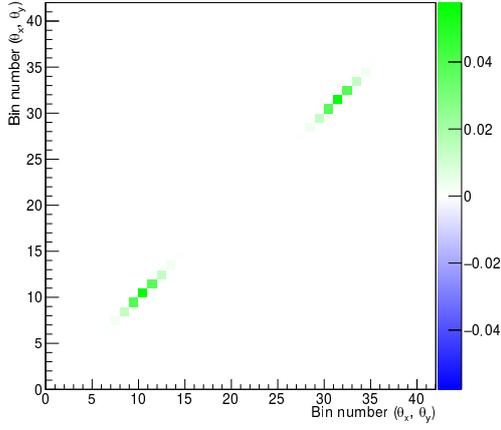


(c) Downstream magnetic field uncertainty

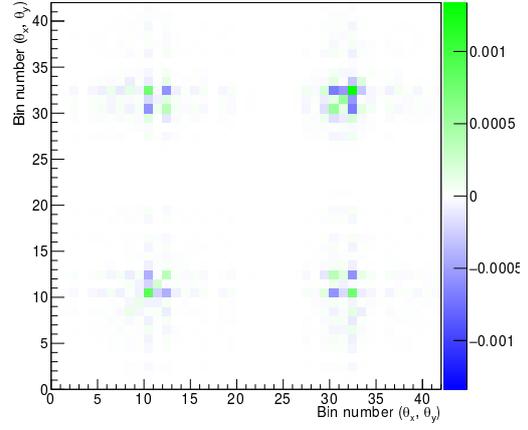


(d) Focus coils magnetic field uncertainty

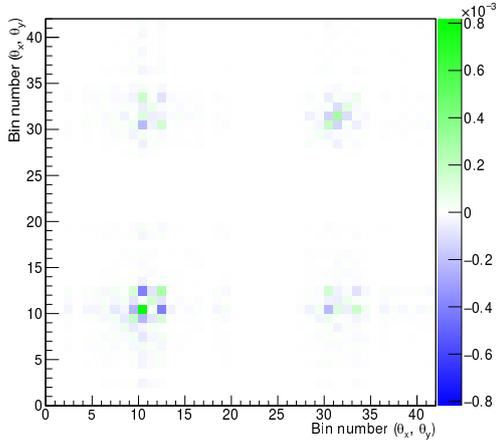
Figure A.44: Momentum selection and field uncertainty covariance matrices for the 230 MeV/c convoluted dataset



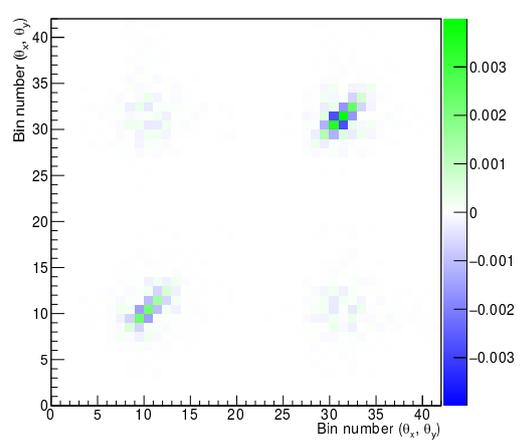
(a) Statistical errors



(b) Alignment shift of ± 1 mm

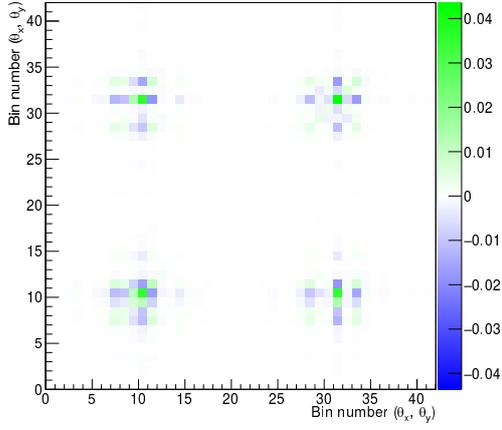


(c) Alignment rotation of ± 1 mrad

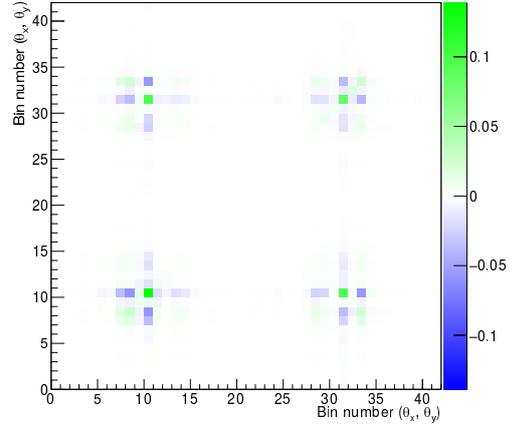


(d) Momentum-dependent shift correction

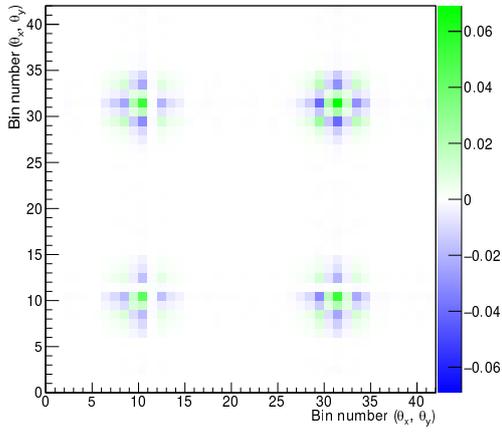
Figure A.45: Statistical and alignment covariance matrices for the 240 MeV/c LiH dataset



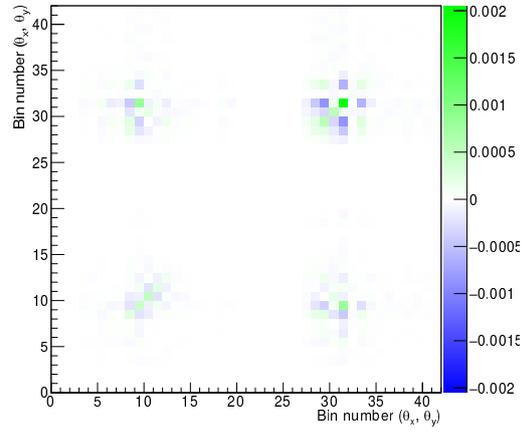
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty

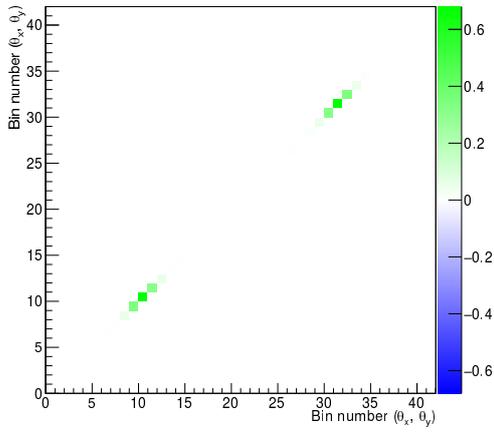


(c) Downstream magnetic field uncertainty

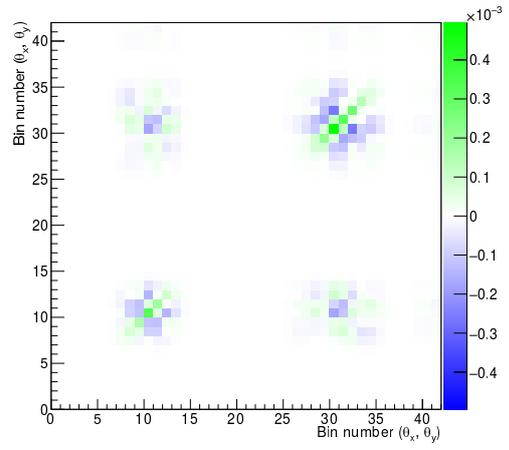


(d) Focus coils magnetic field uncertainty

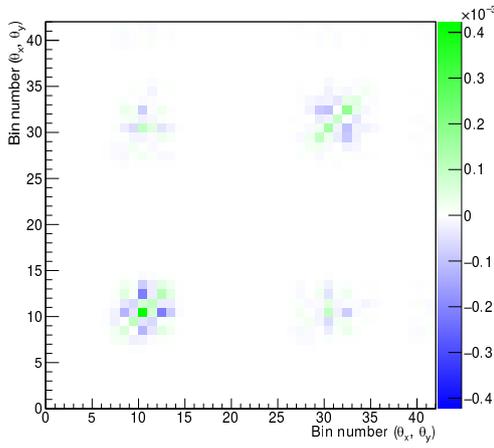
Figure A.46: Momentum selection and field uncertainty covariance matrices for the 240 MeV/c LiH dataset



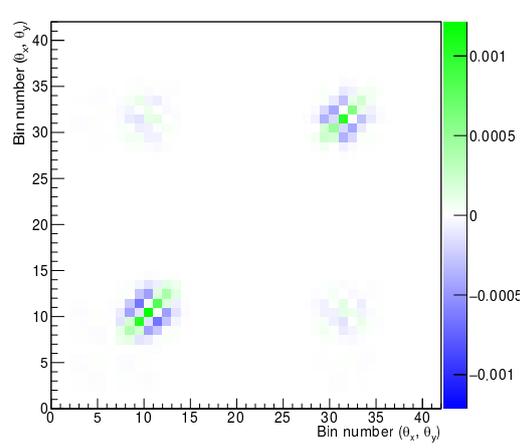
(a) Statistical errors



(b) Alignment shift of ± 1 mm

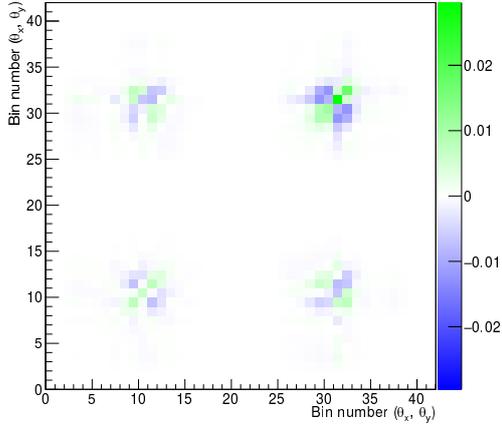


(c) Alignment rotation of ± 1 mrad

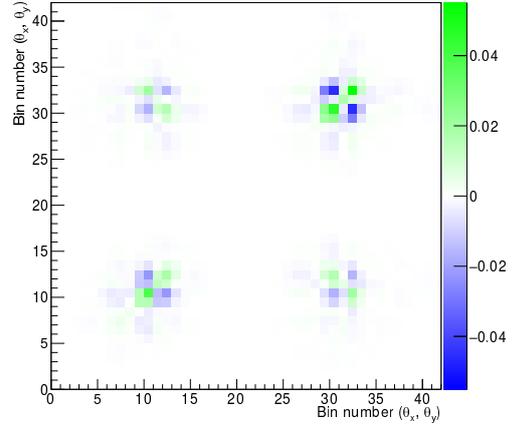


(d) Momentum-dependent shift correction

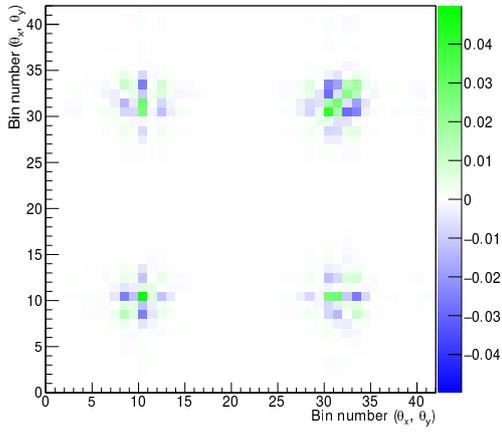
Figure A.47: Statistical and alignment covariance matrices for the 240 MeV/c convoluted dataset



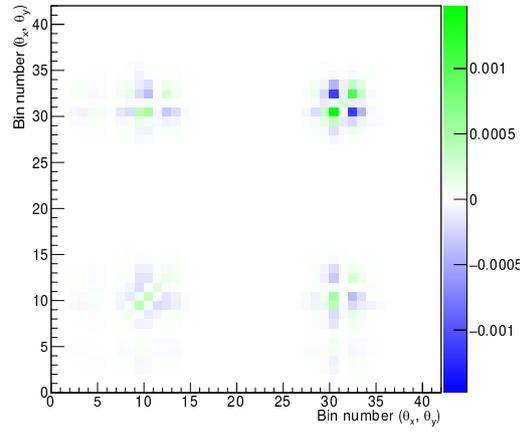
(a) Momentum selection uncertainty



(b) Upstream magnetic field uncertainty



(c) Downstream magnetic field uncertainty



(d) Focus coils magnetic field uncertainty

Figure A.48: Momentum selection and field uncertainty covariance matrices for the 240 MeV/c convoluted dataset

A.2 Combined Covariance Matrices

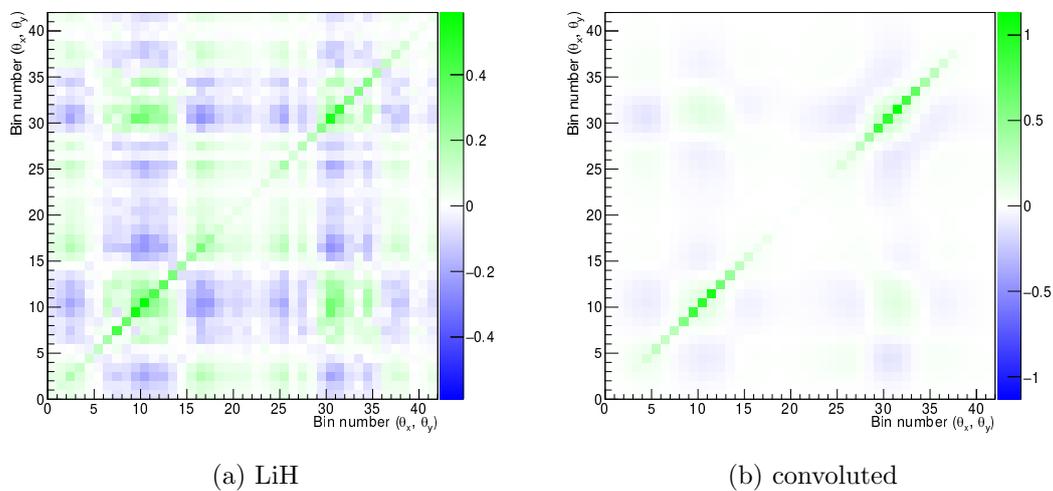


Figure A.49: Combined (incl. statistical errors) covariance matrices 120 MeV/c

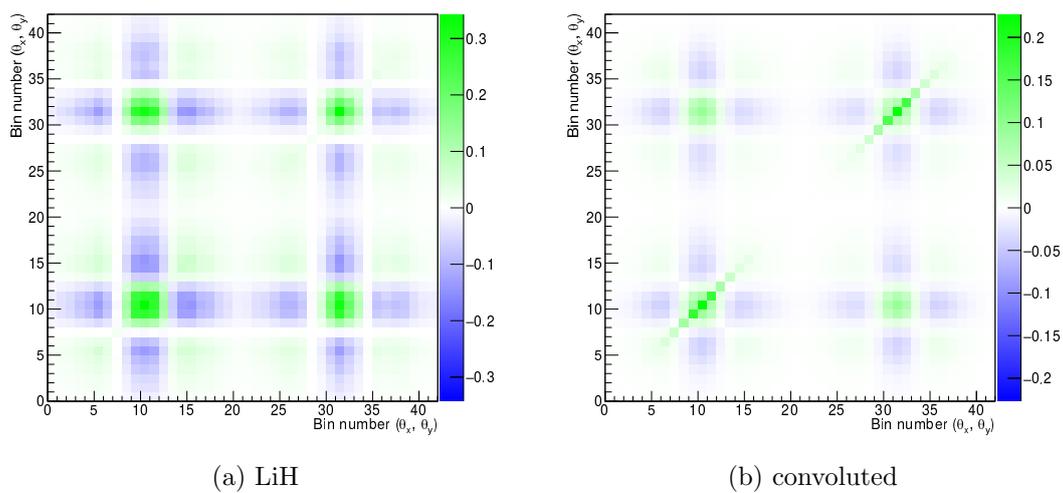
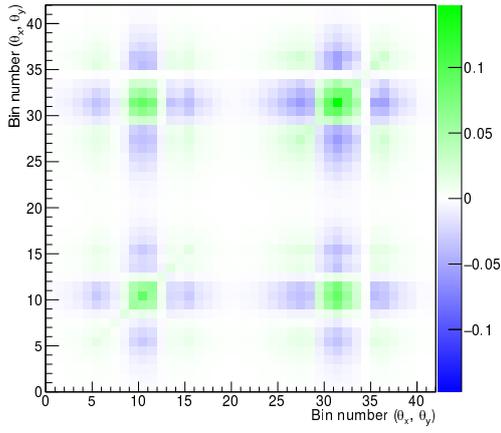
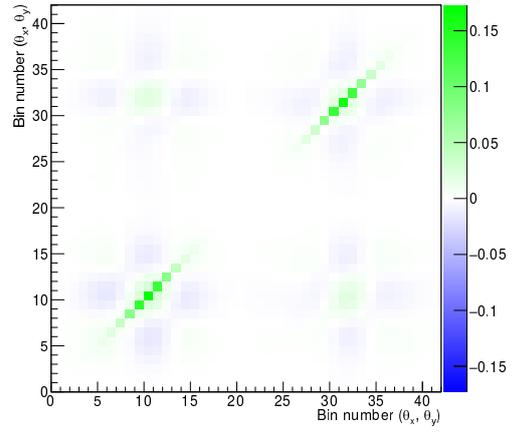


Figure A.50: Combined (incl. statistical errors) covariance matrices 130 MeV/c

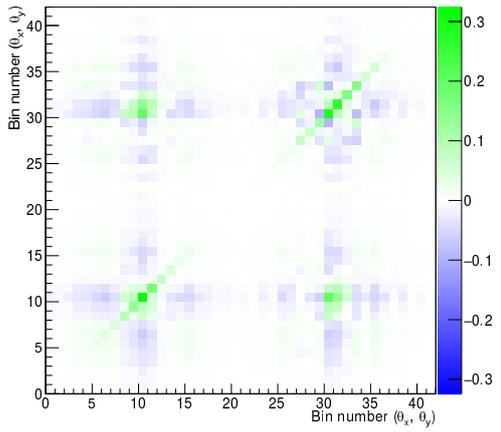


(a) LiH

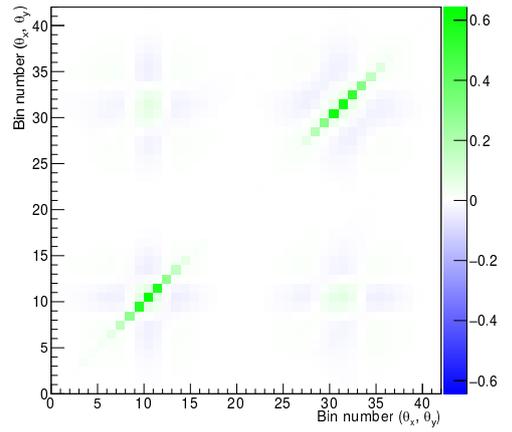


(b) convoluted

Figure A.51: Combined (incl. statistical errors) covariance matrices 140 MeV/c

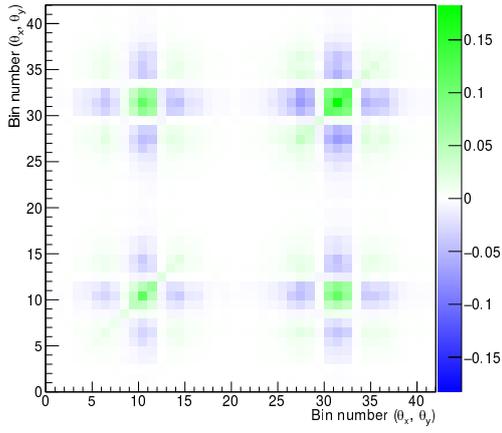


(a) LiH

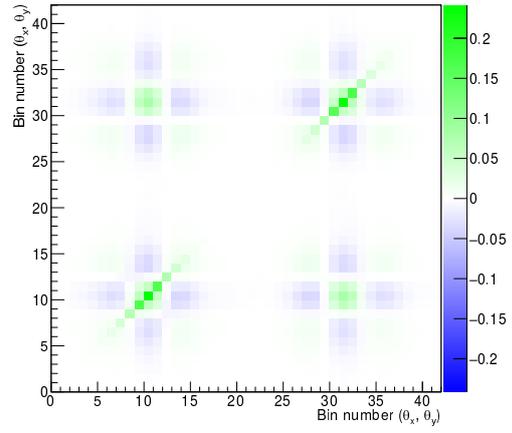


(b) convoluted

Figure A.52: Combined (incl. statistical errors) covariance matrices 150 MeV/c

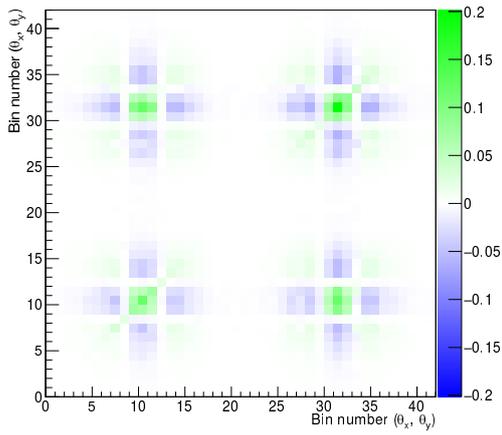


(a) LiH

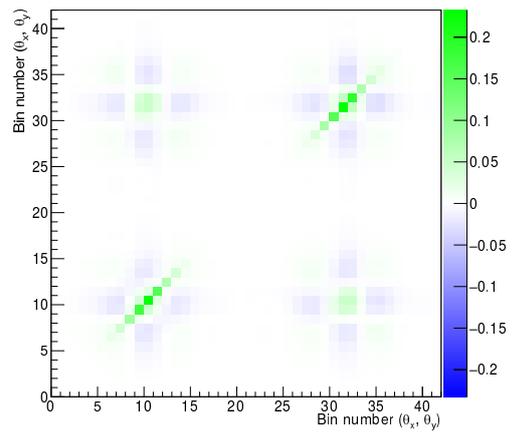


(b) convoluted

Figure A.53: Combined (incl. statistical errors) covariance matrices 160 MeV/c

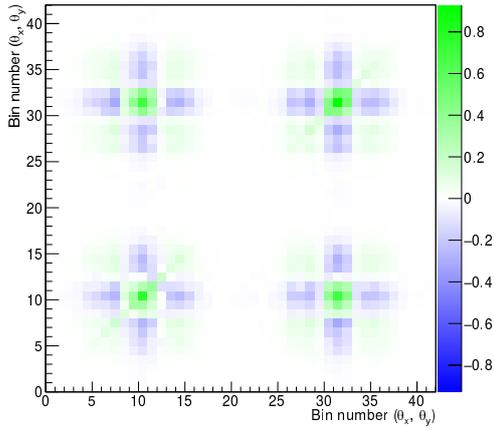


(a) LiH

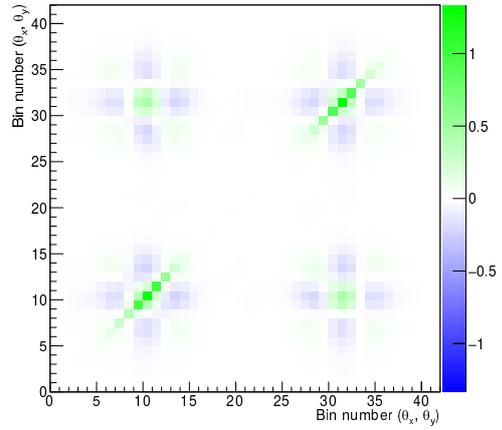


(b) convoluted

Figure A.54: Combined (incl. statistical errors) covariance matrices 170 MeV/c

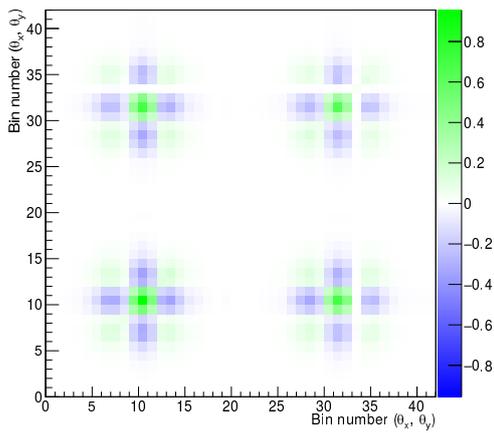


(a) LiH

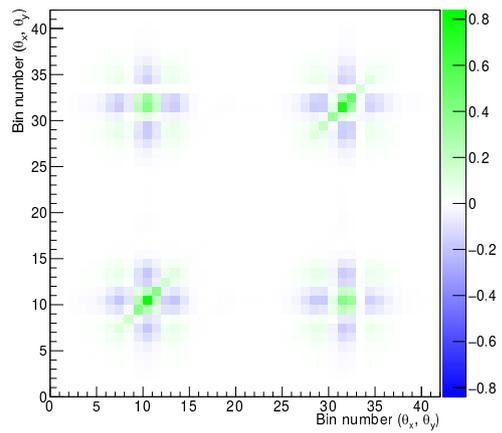


(b) convoluted

Figure A.55: Combined (incl. statistical errors) covariance matrices 180 MeV/c

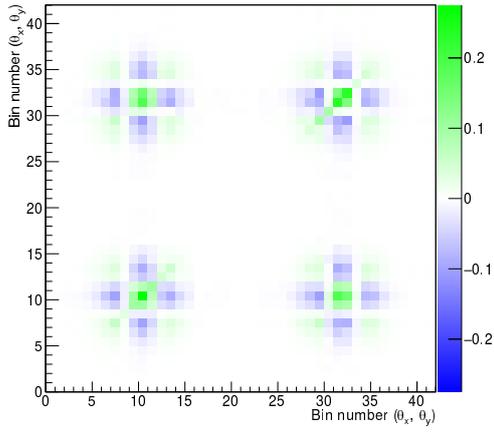


(a) LiH

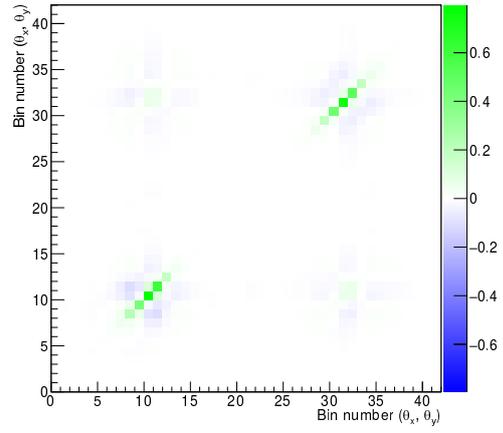


(b) convoluted

Figure A.56: Combined (incl. statistical errors) covariance matrices 190 MeV/c

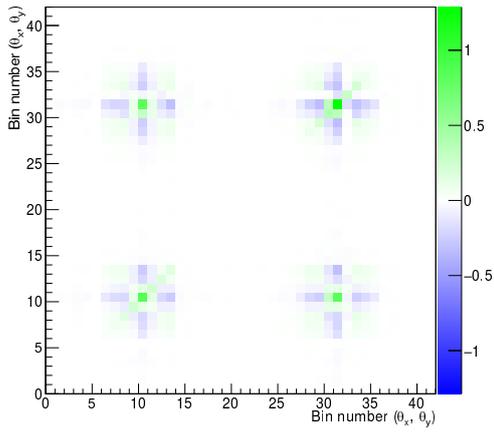


(a) LiH

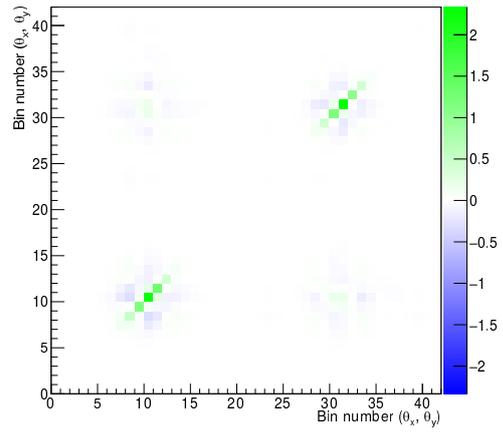


(b) convoluted

Figure A.57: Combined (incl. statistical errors) covariance matrices 200 MeV/c

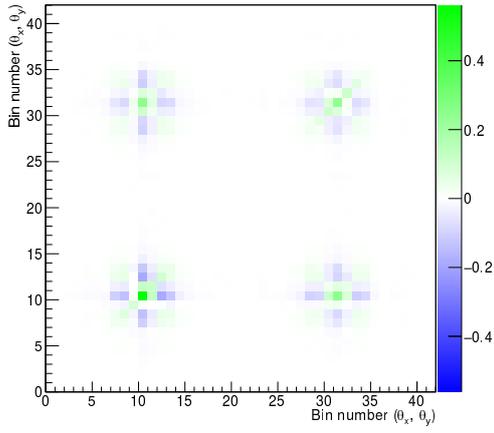


(a) LiH

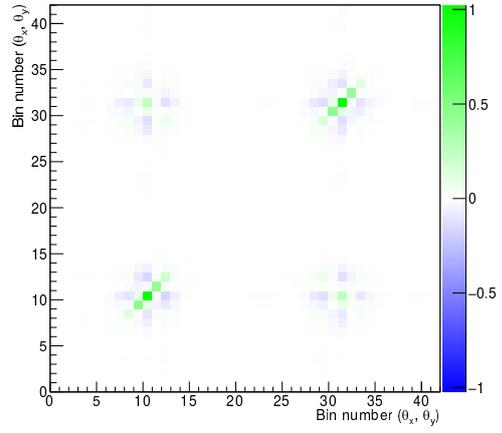


(b) convoluted

Figure A.58: Combined (incl. statistical errors) covariance matrices 220 MeV/c

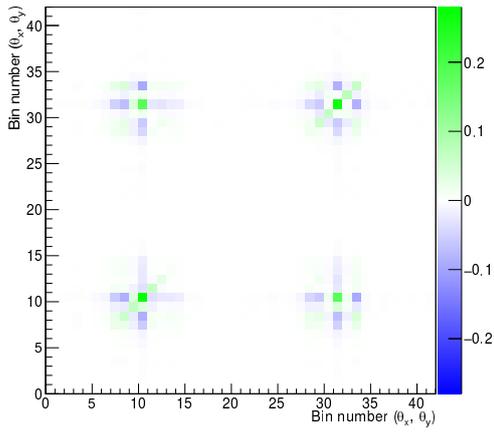


(a) LiH

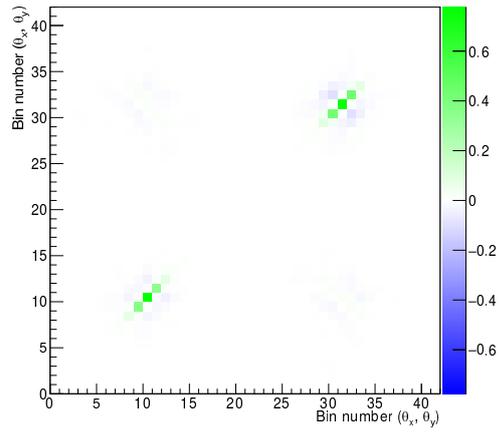


(b) convoluted

Figure A.59: Combined (incl. statistical errors) covariance matrices 230 MeV/c



(a) LiH



(b) convoluted

Figure A.60: Combined (incl. statistical errors) covariance matrices 240 MeV/c

Bibliography

- [1] W. Pauli. *Fünf Arbeiten zum Ausschliessungsprinzip und zum Neutrino*. Texte zur Forschung. Wissenschaftliche Buchgesellschaft, [Abt. Verlag], 1977. ISBN 9783534067336.
- [2] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956. doi: 10.1126/science.124.3212.103.
- [3] Ivan V. Anicin. *The Neutrino: Its past, present and future*. 2005.
- [4] K. Kodama et al. Observation of tau neutrino interactions. *Phys. Lett.*, B504: 218–224, 2001. doi: 10.1016/S0370-2693(01)00307-0.
- [5] K. A. Olive. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016. doi: 10.1088/1674-1137/40/10/100001.
- [6] Raymond Davis, Jr., Don S. Harmer, and Kenneth C. Hoffman. Search for neutrinos from the sun. *Phys. Rev. Lett.*, 20:1205–1209, 1968. doi: 10.1103/PhysRevLett.20.1205.
- [7] K. S. Hirata et al. Observation of B-8 Solar Neutrinos in the Kamiokande-II Detector. *Phys. Rev. Lett.*, 63:16, 1989. doi: 10.1103/PhysRevLett.63.16.
- [8] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. *Phys. Rev.*, C60:055801, 1999. doi: 10.1103/PhysRevC.60.055801.
- [9] D. Vignaud. The GALLEX solar neutrino experiment. *Nucl. Phys. Proc. Suppl.*, 60B:20–29, 1998. doi: 10.1016/S0920-5632(97)00498-2.
- [10] S. Fukuda et al. Solar B-8 and hep neutrino measurements from 1258 days of Super-Kamiokande data. *Phys. Rev. Lett.*, 86:5651–5655, 2001. doi: 10.1103/PhysRevLett.86.5651.

- [11] Q. R. Ahmad et al. Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ${}^8\text{B}$ solar neutrinos at the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 87:071301, 2001. doi: 10.1103/PhysRevLett.87.071301.
- [12] Y. Fukuda et al. Study of the atmospheric neutrino flux in the multi-GeV energy range. *Phys. Lett.*, B436:33–41, 1998. doi: 10.1016/S0370-2693(98)00876-4.
- [13] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002. doi: 10.1103/PhysRevLett.89.011301.
- [14] K. Zuber. *Neutrino Physics, Second Edition*. Series in High Energy Physics, Cosmology and Gravitation. CRC Press, 2011. ISBN 9781420064728.
- [15] F. Capozzi, E. Lisi, A. Marrone, D. Montanino, and A. Palazzo. Neutrino masses and mixings: Status of known and unknown 3ν parameters. *Nucl. Phys.*, B908:218–234, 2016. doi: 10.1016/j.nuclphysb.2016.02.016.
- [16] Francesco Capozzi, Eleonora Di Valentino, Eligio Lisi, Antonio Marrone, Alessandro Melchiorri, and Antonio Palazzo. Global constraints on absolute neutrino masses and their ordering. *Phys. Rev.*, D95(9):096014, 2017. doi: 10.1103/PhysRevD.95.096014.
- [17] A. Gando et al. Constraints on θ_{13} from A Three-Flavor Oscillation Analysis of Reactor Antineutrinos at KamLAND. *Phys. Rev.*, D83:052002, 2011. doi: 10.1103/PhysRevD.83.052002.
- [18] K. Abe et al. Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target. *Phys. Rev.*, D91(7):072010, 2015. doi: 10.1103/PhysRevD.91.072010.
- [19] M. Ravonel Salzgeber. Anti-neutrino oscillations with T2K. 2015.
- [20] P. Adamson et al. First measurement of electron neutrino appearance in NOvA. *Phys. Rev. Lett.*, 116(15):151806, 2016. doi: 10.1103/PhysRevLett.116.151806.
- [21] P. Adamson et al. First measurement of muon-neutrino disappearance in NOvA. *Phys. Rev.*, D93(5):051104, 2016. doi: 10.1103/PhysRevD.93.051104.
- [22] K. Abe et al. The T2K Experiment. *Nucl. Instrum. Meth.*, A659:106–135, 2011. doi: 10.1016/j.nima.2011.06.067.

- [23] J. H. Choi et al. Observation of Energy and Baseline Dependent Reactor Antineutrino Disappearance in the RENO Experiment. *Phys. Rev. Lett.*, 116(21):211801, 2016. doi: 10.1103/PhysRevLett.116.211801.
- [24] F. P. An et al. New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay. *Phys. Rev. Lett.*, 115(11):111802, 2015. doi: 10.1103/PhysRevLett.115.111802.
- [25] Roger Wendell. Atmospheric Neutrino Oscillations at Super-Kamiokande. *PoS, ICRC2015:1062*, 2016.
- [26] M. G. Aartsen et al. Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data. *Phys. Rev.*, D91(7):072004, 2015. doi: 10.1103/PhysRevD.91.072004.
- [27] R. Acciarri et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE). 2016.
- [28] James Strait et al. Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE). 2016.
- [29] K. Abe et al. Physics Potentials with the Second Hyper-Kamiokande Detector in Korea. 2016.
- [30] K. Abe et al. Physics Potentials with the Second Hyper-Kamiokande Detector in Korea. 2016.
- [31] Yu-Feng Li. Overview of the Jiangmen Underground Neutrino Observatory (JUNO). *Int. J. Mod. Phys. Conf. Ser.*, 31:1460300, 2014. doi: 10.1142/S2010194514603007.
- [32] M. G. Aartsen et al. PINGU: A Vision for Neutrino and Particle Physics at the South Pole. 2016.
- [33] J. Scott Berg, A. Blondel, A. Bross, J. Morfin, K. Long, Jurgen Pozimski, P. Soler, R. Tsenov, and M. Zisman. The International Design Study for the Neutrino Factory. *ICFA Beam Dyn. Newslett.*, 55:54–72, 2011.
- [34] S. Choubey et al. International Design Study for the Neutrino Factory, Interim Design Report. 2011.

- [35] X. Ding, J. S. Berg, D. Cline, and H. G. Kirk. Optimization of a mercury jet target for a neutrino factory or a muon collider. *Phys. Rev. ST Accel. Beams*, 14:111002, 2011. doi: 10.1103/PhysRevSTAB.14.111002.
- [36] F. J. P. Soler. nuSTORM: Neutrinos from Stored Muons. In *Proceedings, Topical Research Meeting on Prospects in Neutrino Physics (NuPhys2014): London, UK, December 15-17, 2014*, 2015.
- [37] Daniel M. Kaplan. Muon Colliders and Neutrino Factories. *EPJ Web Conf.*, 95:03019, 2015. doi: 10.1051/epjconf/20159503019.
- [38] H. Wiedemann. *Particle Accelerator Physics*. Graduate Texts in Physics. Springer International Publishing, 2015. ISBN 9783319183176.
- [39] E. Wilson and E.J.N. Wilson. *An Introduction to Particle Accelerators*. Oxford University Press, 2001. ISBN 9780198508298.
- [40] David Neuffer. Principles and Applications of Muon Cooling. *Conf. Proc.*, C830811:481–484, 1983.
- [41] MICE: An International Muon Ionisation Cooling Experiment, Technical Reference Document. 2005.
- [42] Isis.stfc.ac.uk. How isis works - in depth, 2015. URL <http://www.isis.stfc.ac.uk/about/how-isis-works---in-depth4371.html>.
- [43] M. Bogomilov et al. The MICE Muon Beam on ISIS and the beam-line instrumentation of the Muon Ionization Cooling Experiment. *JINST*, 7:P05009, 2012. doi: 10.1088/1748-0221/7/05/P05009.
- [44] C.N. Booth, P. Hodgson, L. Howlett, R. Nicholson, E. Overton, et al. The design, construction and performance of the MICE target. *JINST*, 8:P03006, 2013. doi: 10.1088/1748-0221/8/03/P03006.
- [45] Holger Witte, Stephen Plate, J. Scott Berg, Jason Tarrant, and Alan Bross. Partial Return Yoke for MICE Step IV and Final Step. In *Proceedings, 6th International Particle Accelerator Conference (IPAC 2015): Richmond, Virginia, USA, May 3-8, 2015*, page WEPJE027, 2015.
- [46] S. Ishimoto, S. Suzuki, M. Yoshida, M. A. Green, Y. Kuno, and W. Lau. Liquid Hydrogen Absorber for MICE. *Conf. Proc.*, C100523:421–423, 2010.

- [47] Jean-Baptiste Lagrange, Christopher Hunt, Vittorio Palladino, and Jaroslaw Pasternak. The MICE Demonstration of Muon Ionization Cooling. In *Proceedings, 7th International Particle Accelerator Conference (IPAC 2016): Busan, Korea, May 8-13, 2016*, page TUPMY004, 2016. doi: 10.18429/JACoW-IPAC2016-TUPMY004.
- [48] A. Dobbs, D. Forrest, and Soler. F. J. P. The MICE luminosity monitor. *J.Phys.Conf.Ser.*, 408:012084, 2013. doi: 10.1088/1742-6596/408/1/012084.
- [49] R. Bertoni, M. Bonesini, A. de Bari, G. Cecchet, and S. Terzo. The Construction and Laboratory Tests for the MICE TOF0/1 Detectors. *Internal Note*, MICE-NOTE-DET-241, 2009.
- [50] R. Bertoni, M. Bonesini, A. de Bari, G. Cecchet, Y. Karadzhov, and R. Mazza. The Construction of the MICE TOF2 Detector. *Internal Note*, MICE-NOTE-DET-286, 2009.
- [51] R. Bertoni, M. Bogomilov, M. Bonesini, A. de Bari, G. Cecchet, et al. Analysis of PID detectors (TOF and KL) performances in the MICE 2010 Run. *Internal Note*, MICE-NOTE-DET-337, 2011.
- [52] A. Dobbs, K. Long, E. Santos, D. Adey, P. Hanlet, and C. Heidt. The Reconstruction Software for the Muon Ionization Cooling Experiment Trackers. *J. Phys. Conf. Ser.*, 513:022008, 2014. doi: 10.1088/1742-6596/513/2/022008.
- [53] M. Ellis et al. The Design, construction and performance of the MICE scintillating fibre trackers. *Nucl. Instrum. Meth.*, A659:136–153, 2011. doi: 10.1016/j.nima.2011.04.041.
- [54] F. Ambrosino et al. Measurement of the charged kaon lifetime with the KLOE detector. *JHEP*, 01:073, 2008. doi: 10.1088/1126-6708/2008/01/073.
- [55] R. Sandström. Comparison of two calorimeter designs. *Internal Note*, MICE-NOTE-DET-146, 2006.
- [56] Ruslan Asfandiyarov. *Totally Active Scintillator Tracker-Calorimeter for the Muon Ionization Cooling Experiment*. PhD thesis, Geneva U., 2014.
- [57] D. Adams et al. Electron-Muon Ranger: performance in the MICE Muon Beam. *JINST*, 10(12):P12012, 2015. doi: 10.1088/1748-0221/10/12/P12012.

- [58] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, January 2008. ISSN 0001-0782. doi: 10.1145/1327452.1327492.
- [59] L. Cremaldi, M. Drews, D. Kaplan, D. Rajaram, D. Sanders, D. Summers, and M Winter. Progress on Cherenkov Reconstruction in MICE. *Internal Note*, MICE-NOTE-DET-473, 2015.
- [60] A. Dobbs, C. Hunt, K. Long, E. Santos, M. A. Uchida, P. Kyberd, C. Heidt, S. Blot, and E. Overton. The reconstruction software for the MICE scintillating fibre trackers. 2016.
- [61] URL http://heplnv152.pp.rl.ac.uk/maus/MAUS_latest_version/doc/index.html.
- [62] C. Runge. *Graphical Methods: A Course of Lectures Delivered in Columbia University, New York, October, 1909, to January, 1910.* ... Graphical Methods. Columbia University Press, 1912.
- [63] Celeste Pidcott. *Multiple Scattering and Particle Identification in the Muon Ionization Cooling Experiment.* PhD thesis, University of Warwick, 2016.
- [64] E. Rutherford. The scattering of α and β particles by matter and the structure of the atom. *Phil. Mag. Ser.6*, 21:669–688, 1911. doi: 10.1080/14786440508637080.
- [65] G. Wentzel. Zwei Bemerkungen über die Zerstreung korpuskularer Strahlen als Beugungserscheinung. *Zeitschrift für Physik*, 40(8):590–593, Aug 1926. ISSN 0044-3328. doi: 10.1007/BF01390457.
- [66] Timothy Carlisle. *Step IV of the Muon Ionization Cooling Experiment (MICE) and the multiple scattering of muons.* PhD thesis, Oxford U., 2013.
- [67] G. Molière. Theorie der Streuung schneller geladener Teilchen II. Mehrfach- und Vielfachstreuung. *Zeitschrift Naturforschung Teil A*, 3:78–97, February 1948. doi: 10.1515/zna-1948-0203.
- [68] H. A. Bethe. Molière’s theory of multiple scattering. *Phys. Rev.*, 89:1256–1266, Mar 1953. doi: 10.1103/PhysRev.89.1256.
- [69] U. Fano. Inelastic collisions and the molière theory of multiple scattering. *Phys. Rev.*, 93:117–120, Jan 1954. doi: 10.1103/PhysRev.93.117.

- [70] H. W. Lewis. Multiple scattering in an infinite medium. *Phys. Rev.*, 78:526–529, Jun 1950. doi: 10.1103/PhysRev.78.526.
- [71] GEANT Collaboration. GEANT4 : Physics Reference Manual. Nov 2012. URL <https://cds.cern.ch/record/2237422>. CERN Program Library Long Writeups.
- [72] S. Goudsmit and J. L. Saunderson. Multiple scattering of electrons. *Phys. Rev.*, 57: 24–29, Jan 1940. doi: 10.1103/PhysRev.57.24.
- [73] V N Ivanchenko, O Kadri, M Maire, and L Urban. Geant4 models for simulation of multiple scattering. *Journal of Physics: Conference Series*, 219(3):032045, 2010.
- [74] D. Attwood et al. The scattering of muons in low Z materials. *Nucl. Instrum. Meth.*, B251:41–55, 2006. doi: 10.1016/j.nimb.2006.05.006.
- [75] W W M Allison. Calculations of energy loss and multiple scattering (elms) in molecular hydrogen. *Journal of Physics G: Nuclear and Particle Physics*, 29(8): 1701, 2003.
- [76] Mikael Kuusela. Unfolding: A statistician’s perspective. Conference Slides, Phystat- ν , 9 2016.
- [77] N. D. Gagunashvili. Pearson’s chi-square test modifications for comparison of un-weighted and weighted histograms and two weighted histograms. *PoS*, ACAT:060, 2007.
- [78] John Columba Nugent. *Multiple coulomb scattering in the MICE experiment*. PhD thesis, Glasgow U., 2017.