

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

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/177206>

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



Beam Studies in Muon and Neutrino Physics

by

Thomas William Lord

Thesis

Submitted to the University of Warwick

in partial fulfilment of the requirements

for admission to the degree of

Doctor of Philosophy

Department of Physics

December 2021

Contents

List of Tables	v
List of Figures	viii
Acknowledgments	xxiv
Declarations	xxv
Abstract	xxvi
Introduction	1
I MICE	
The Muon Ionization Cooling Experiment	3
Chapter 1 Background	4
1.1 Beam Optics	6
1.1.1 Phase-Space and Trace-Space	6
1.1.2 Solenoid Optics	9
1.1.3 Liouville's Theorem	10
1.1.4 Single Particle Amplitude	10
1.2 Beam Matching	11
1.3 Cooling	13
1.3.1 Radiative Cooling	13
1.3.2 Stochastic Cooling	13
1.3.3 Ionisation Cooling	14
1.3.4 Frictional Cooling	15
1.3.5 Conclusion	16
Chapter 2 The Muon Ionisation Cooling Experiment (MICE)	17
2.1 Experimental Goals	17
2.2 MICE Beamline	18
2.2.1 MICE Coordinate System	19

2.2.2	ISIS Proton Synchrotron	19
2.2.3	MICE Target	19
2.2.4	Beam Magnets	22
2.2.5	Proton Absorber	23
2.2.6	Diffuser	24
2.2.7	Partial Return Yoke (PRY)	24
2.3	Cooling Channel	24
2.3.1	Steps I to VI	24
2.3.2	Absorber and Focus Coil (AFC) Module	26
2.3.3	'Solenoid' and 'Flip' Modes	27
2.4	Detectors	28
2.4.1	Luminosity Monitor	28
2.4.2	Time-of-Flight (TOF) Detectors	29
2.4.3	Cherenkov Detectors	30
2.4.4	Scintillating Fibre Trackers	31
2.4.5	KLOE-Light Calorimeter	31
2.4.6	Electron Muon Ranger	32
2.5	MICE Software Framework	33
2.5.1	MICE Analysis User Software (MAUS)	33
2.5.2	Simulation	33
2.5.3	Reconstruction	34
2.5.4	Configuration Database	35
2.5.5	Controls, Monitoring and Data Acquisition	36
Chapter 3 Tracker		37
3.1	Tracker Construction	37
3.2	Track Reconstruction	38
3.2.1	Digitisation	39
3.2.2	Clustering	39
3.2.3	Spacepoint Formation	40
3.2.4	Pattern Recognition	41
3.2.5	Kalman Filter	41
3.2.6	Track Reconstruction Performance	43
3.3	Global Reconstruction	49
3.3.1	Track Matching	49
3.3.2	Particle Identification	50
3.4	Conclusion	51
Chapter 4 Emittance and Amplitude Analysis		52
4.1	Beam Configurations	52
4.2	Data Selection	53

4.2.1	Upstream Sample	54
4.2.2	Downstream Sample	55
4.2.3	Simulation	60
4.2.4	Real versus Simulated Sample Selection	65
4.2.5	Hybrid Simulation	74
4.3	Calculation of Amplitudes	79
4.4	Effects of Detector Reconstruction on Amplitude	82
4.4.1	Track Reconstruction Inefficiency	82
4.4.2	Beam Impurity	82
4.4.3	Reconstruction Resolution	83
4.4.4	Corrections	91
4.5	Calculation of Uncertainties	96
4.5.1	Statistical Uncertainties	96
4.5.2	Systematic Uncertainties	97
4.6	Results	114
4.6.1	Amplitude PDFs and CDFs	114
4.6.2	Corrected vs Uncorrected Distributions	119
4.6.3	PDF and CDF Ratios	121
Chapter 5 Canonical Angular Momentum		125
5.1	Canonical Angular Momentum Build-up	125
5.2	Analysis	127
5.2.1	Single Particle Distributions	127
5.2.2	Tracker Evolution	133
5.3	Summary	138
Chapter 6 Conclusions and Future Work		139
6.1	Conclusion	139
6.2	Future Work	140
II DUNE		
The Deep Underground Neutrino Experiment		141
Chapter 7 Background		142
7.1	Neutrino Discovery	142
7.2	Neutrinos in the Standard Model	144
7.3	Neutrinos Beyond The Standard Model	145
7.3.1	Flavour Oscillations	145
7.4	Current Neutrino Sector Parameters	150

Chapter 8 The DUNE Experiment	152
8.1 DUNE Physics Goals	153
8.2 The DUNE Beamline	153
8.2.1 The Target and Horns	154
8.3 Near Detector Complex	156
8.3.1 ND-LAr	156
8.3.2 ND-GAr	157
8.3.3 SAND	158
8.4 Far Detector Complex	158
Chapter 9 DUNE-PRISM	162
9.1 Off-Axis Neutrino Beams	162
9.1.1 Off-Axis Flux at DUNE	163
9.2 Flux Matching	166
9.2.1 Tikhonov Regularisation	168
9.3 Summary	173
Chapter 10 Beam-PRISM	174
10.1 Beam Options	174
10.1.1 Non-Nominal Run-Time	175
10.2 Method	183
10.2.1 Possible Approaches	183
10.2.2 Compressed Sensing	184
10.3 Flux-Matching Objectives	188
10.4 Results	189
10.4.1 PRISM On-Axis	189
10.4.2 Expanding DUNE-PRISM	193
10.4.3 Restricted Off-Axis DUNE-PRISM	199
10.5 Summary	206
Chapter 11 Conclusions and Future Work	207
11.1 Summary	207
11.2 Future Work	208
Appendix A Beam-PRISM Solutions	210
A.1 Restricted Off-Axis Range Solutions	210
A.1.1 25.5m	210

List of Tables

4.1	Table of MICE beam configurations.	52
4.2	Upstream sample TOF01 and TKU Momentum vs TOF01 selection criteria for each beam configuration.	55
4.3	Upstream sample selection for number of reconstructed events, shown to 3 significant figures. Shown for 2017-02-6 3-140 and 4-140 samples.	56
4.4	Upstream sample selection for number of reconstructed events, shown to 3 significant figures. Shown for 2017-02-6 6-140 and 10-140 samples.	57
4.5	Downstream sample selection for number of reconstructed events, shown to 3 significant figures. Shown for 2017-02-6 3-140 and 4-140 samples.	58
4.6	Downstream sample selection for number of reconstructed events, shown to 3 significant figures. Shown for 2017-02-6 6-140 and 10-140 samples.	59
4.7	Upstream sample selection for number of reconstructed MC events, shown to 3 significant figures. Shown for 2017-02-6 3-140 and 4-140 samples.	61
4.8	Upstream sample selection for number of reconstructed MC events, shown to 3 significant figures. Shown for 2017-02-6 6-140 and 10-140 samples.	62
4.9	Downstream sample selection for number of reconstructed MC events, shown to 3 significant figures. Shown for 2017-02-6 3-140 and 4-140 samples.	63
4.10	Downstream sample selection for number of reconstructed MC events, shown to 3 significant figures. Shown for 2017-02-6 6-140 and 10-140 samples.	64
4.11	Upstream sample selection for number of reconstructed hybrid simulation events, shown to 3 significant figures. Shown for 2017-02-6 3, 6, and 10-140 LH ₂ and 4-140 no absorber samples.	75

4.12	Downstream sample selection for number of reconstructed hybrid simulation events, shown to 3 significant figures. Shown for 2017-02-6 3, 6, and 10-140 LH ₂ and 4-140 no absorber samples. . . .	76
4.13	Upstream MC truth sample selection for number of reconstructed hybrid simulation events, shown to 3 significant figures. Shown for 2017-02-6 3, 6, and 10-140 LH ₂ and 4-140 no absorber samples. . . .	77
4.14	Downstream MC truth sample selection for number of reconstructed hybrid simulation events, shown to 3 significant figures. Shown for 2017-02-6 3, 6, and 10-140 LH ₂ and 4-140 no absorber samples.	78
7.1	Summary of neutrino oscillation parameters determined from a recent global analysis.	151
9.1	A sample DUNE-PRISM run plan, outlining the yearly distribution of time spent at various on- and off-axis positions. . . .	164
10.1	Horn configurations, their operating currents and the central off-axis positions of the detector utilised in the example iterative flux-matching solution with coefficients shown in figure 10.9. The solution allows for HC fluxes at currents up to 330 kA. . .	187
10.2	Horn configurations and their operating currents utilised in the on-axis solution for flux-matching, allowing HC fluxes at or below the nominal current of 293 kA.	189
10.3	Horn configurations and their operating currents utilised in the on-axis solution for flux-matching, allowing for HC fluxes at currents up to 350 kA.	191
10.4	Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 27.5 m.	200
10.5	Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 20.5 m.	201
10.6	Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 14.5 m.	204
11.1	Sources of systematic flux uncertainty considered for the nominal-flux DUNE-PRISM analysis.	209
A.1	Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 25.5 m.	210

A.2	Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 23.5 m.	212
A.3	Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 18.5 m.	214

List of Figures

1.1	Block diagrams of neutrino factory and muon collider scenarios, showing the same shared complex up to the green dashed line. [1]	5
1.2	The extrema, intercepts and correlation of a trace-space ellipse described in terms of the Twiss parameters.	7
1.3	Trace-space evolution of beam filamentation from beam mismatching, producing non-linear emittance growth. The ideal matched ellipse is shown in red, with the actual beam ellipse in black.	12
1.4	Illustration of steering error, showing larger emittance ellipse achieved after filamentation.	12
1.5	Ionisation cooling principle shown in 3 steps.	14
2.1	The MICE beamline, Step IV.	18
2.2	Schematic of the ISIS proton synchrotron.	20
2.3	Schematic of the MICE target.	21
2.4	The Random Acronym for Target System (RATS) target control interface.	22
2.5	Kinematic limits for muon production from pion decays at MICE and simulated pion and muon spectra after the decay solenoid.	23
2.6	Cross-section schematic of the MICE diffuser and photo of the diffuser during testing.	24
2.7	The MICE Partial Return Yoke.	25
2.8	The 6 proposed steps for study of muon ionization cooling at MICE.	26
2.9	Solenoid field model through the cooling channel and the corresponding linear optics evolution of RMS beam width for 3, 4, 6, and 10 mm nominal emittance beams.	27
2.10	Cross-section of the luminosity monitor, viewed side-on relative to the MICE target particle spill	28
2.11	Blown-out aerogel Cherenkov counter, showing the entrance window, mirror, aerogel mosaic, acetate window, GORE reflector panel, exit window, and 8 inch PMT in iron shield.	30

2.12	Cross-section of the scintillating fibres embedded in extruded lead within the KL.	31
2.13	Blown-out image of the construction of the EMR.	32
3.1	Schematic of the carbon-fibre tracker frame and photograph prior to installation.	38
3.2	Arrangement and cross-sectional layout of the doublet layers U, V, and W in each scintillating fibre station.	39
3.3	Illustration of Kuno's conjecture along planes U, V, and W. Summing the signed distances from the center of each fibre to the plane centre, α_U, α_V and α_W , will equal zero for any three simultaneously coincident fibres.	40
3.4	Reconstruction residuals (measured - true) for TKU S1 track-point x positions in mm. Each beam has an RMS of 0.4 mm.	44
3.5	Reconstruction residuals (measured - true) for TKU S1 track-point y positions in mm. Each beam has an RMS of 0.4 mm.	44
3.6	Reconstruction residuals (measured - true) for TKD S1 track-point x positions in mm. Each beam has an RMS of 0.4 mm.	45
3.7	Reconstruction residuals (measured - true) for TKD S1 track-point y positions in mm. Each beam has an RMS of 0.4 mm.	45
3.8	Reconstruction residuals (measured - true) for TKU S1 track-point p_x momentum in MeV/c. Each beam has an approximate RMS of 1.3 MeV/c.	46
3.9	Reconstruction residuals (measured - true) for TKU S1 track-point p_y momentum in MeV/c. Each beam has an approximate RMS of 1.3 MeV/c.	46
3.10	Reconstruction residuals (measured - true) for TKD S1 track-point p_x momentum in MeV/c. Each beam has an approximate RMS of 1.5 MeV/c.	47
3.11	Reconstruction residuals (measured - true) for TKD S1 track-point p_y momentum in MeV/c. Each beam has an approximate RMS of 1.5 MeV/c.	47
3.12	Reconstruction residuals (measured - true) for TKU S1 track-point p_z momentum in MeV/c. Each beam has an approximate RMS of 2.3 MeV/c.	48

3.13	Reconstruction residuals (measured - true) for TKD S1 track-point p_z momentum in MeV/c. Each beam has an approximate RMS of 3.2 MeV/c.	48
4.1	Reconstructed momentum distributions at TKU Station 1 for an upstream muon sample with all selection requirements other than the momentum selection requirement applied. Plots momentum distributions (left) before and (right) after dipole tuning, with real data events in black and MC events in yellow.	60
4.2	Reconstructed TKU momentum versus TOF01 time-of-flight for all events in (top) data and (bottom) MC which reconstruct a single track in TKU and a single spacepoint in TOFs 1 and 2. The upper and lower bounds of the TKU momentum vs TOF01 selection is shown by the blue curves in each distribution. The TKU momentum and TOF01 requirements are shown in red.	66
4.3	Upstream sample selection requiring 1 reconstructed spacepoint in TOF1, showing real data events in black and MC events in yellow.	67
4.4	Upstream sample selection requiring 1 reconstructed spacepoint in TOF0, showing real data events in black and MC events in yellow.	67
4.5	Upstream sample selection requiring 1 reconstructed track in TKU, showing real data events in black and MC events in yellow.	68
4.6	Upstream sample selection requiring a $\chi^2 > 8$ for reconstructed TKU tracks, showing real data events in black and MC events in yellow.	68
4.7	Upstream sample selection requiring max fiducial radius less than 150 mm for reconstructed TKU tracks, showing real data events in black and MC events in yellow.	69
4.8	Upstream sample selection requiring a TOF01 time-of-flight consistent with the muon hypothesis, showing real data events in black and MC events in yellow. Dotted lines are included to guide the eye.	69
4.9	Upstream sample selection requiring a reconstructed momentum of 140 ± 5 for tracks in TKU, showing real data events in black and MC events in yellow.	70
4.10	Upstream sample selection requiring a global track propagated upstream which passes through the US virtual diffuser face at radius < 90 mm, showing real data events in black and MC events in yellow.	70

4.11	Upstream sample selection requiring a global track propagated upstream which passes through the DS virtual diffuser face at radius < 90 mm, showing real data events in black and MC events in yellow.	71
4.12	Downstream sample selection requiring 1 reconstructed track in TKD, showing real data events in black and MC events in yellow.	71
4.13	Downstream sample selection requiring a $\chi^2 > 8$ for reconstructed TKD tracks, showing real data events in black and MC events in yellow.	72
4.14	Downstream sample selection requiring max fiducial radius less than 150 mm for reconstructed TKD tracks, showing real data events in black and MC events in yellow.	72
4.15	Downstream sample selection requiring a reconstructed momentum of $140 + 30$ MeV/c, -50 MeV/c for tracks in TKD, showing real data events in black and MC events in yellow.	73
4.16	Example RMS ellipses resulting from iterative recalculation of the covariance matrix. The two half-samples are coloured in shades of red and green, with darker shades as the amplitude cut approaches the distribution centre.	80
4.17	Example MC truth amplitude PDF and CDF distributions for a simulated 4-140 beam passed through a lithium hydride absorber. Normalised bin contents of the upstream and downstream distributions are denoted by red and green markers respectively, with an arbitrary smooth curve added to represent the possible underlying data distribution. PDFs are normalised against the largest US bin, CDFs are normalised against the largest DS bin.	81
4.18	Example MC truth amplitude PDF and CDF distributions for a simulated 6-140 beam passed through a lithium hydride absorber. Normalised bin contents of the upstream and downstream distributions are denoted by red and green markers respectively, with an arbitrary smooth curve added to represent the possible underlying data distribution. PDFs are normalised against the largest US bin, CDFs are normalised against the largest DS bin.	81
4.19	Reconstruction inefficiencies in the tracker as a function of different 2D phase-space variables for a simulated 10-140 no absorber beam.	83
4.20	Reconstruction residuals of x and y in TKU, calculated as reconstructed minus true phase-space variables.	84
4.21	Reconstruction residuals of p_x and p_y in TKU, calculated as reconstructed minus true phase-space variables.	85

4.22	Reconstruction residuals of x and y in TKD, calculated as reconstructed minus true phase-space variables.	86
4.23	Reconstruction residuals of p_x and p_y in TKD, calculated as reconstructed minus true phase-space variables.	87
4.24	Comparison of reconstructed and true x , y , p_x , and p_y phase-space variables in TKU station 1. Black points show reconstructed variables for MC events, yellow histogram shows true variables for MC events.	88
4.25	Comparison of reconstructed and true x , y , p_x , and p_y phase-space variables in TKD station 1. Black points show reconstructed variables for MC events, yellow histogram shows true variables for MC events.	89
4.26	Hall probe readings over time from several hall probes positioned along the MICE cooling channel.	90
4.27	US migration matrices for 3, 4, 6, and 10-140 beams describing the probability of TKU reconstruction in each amplitude bin compared to an event's true amplitude bin.	92
4.28	DS migration matrices for 3, 4, 6, and 10-140 beams describing the probability of TKD reconstruction in each amplitude bin compared to an event's true amplitude bin.	93
4.29	Inefficiency correction factors for each DS MC truth amplitude bin for (top to bottom) 3, 4, 6, and 10-140 beams.	94
4.30	Hybrid MC upstream amplitude PDF resolutions for (top left) 3-140, (top right) 4-140, (bottom left) 6-140, and (bottom right) 10-140 beams before and after applying the efficiency and migration correction.	95
4.31	Hybrid MC downstream amplitude PDF resolutions for (top left) 3-140, (top right) 4-140, (bottom left) 6-140, and (bottom right) 10-140 beams before and after applying the efficiency and migration correction.	95
4.32	Comparison of reconstructed TKU x-position distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	100
4.33	Comparison of reconstructed TKU y-position distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	101

4.34	Comparison of reconstructed TKU p_x momentum distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	102
4.35	Comparison of reconstructed TKU p_y momentum distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	103
4.36	Comparison of reconstructed TKU χ^2 distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	104
4.37	Comparison of reconstructed TKD x-position distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	105
4.38	Comparison of reconstructed TKD y-position distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	106
4.39	Comparison of reconstructed TKD p_x momentum distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	107
4.40	Comparison of reconstructed TKD p_y momentum distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	108
4.41	Comparison of reconstructed TKD χ^2 distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	109
4.42	Comparison of TKD maximum radial excursion at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.	110

4.43	Change in the downstream efficiency correction as a result of each systematic variation of the simulation geometry, shown for 3-140 (top left), 4-140 (top right), 6-140 (bottom left), and 10-140 (bottom right) beams.	111
4.44	Change of the diagonal terms in the upstream migration matrix as a result of each systematic variation of the simulation geometry, shown for 3-140 (top left), 4-140 (top right), 6-140 (bottom left), and 10-140 (bottom right) beams. Errors are zero by definition currently as there are no performance systematics on the upstream distributions.	112
4.45	Change of the diagonal terms in the downstream migration matrix as a result of each systematic variation of the simulation geometry, shown for 3-140 (top left), 4-140 (top right), 6-140 (bottom left), and 10-140 (bottom right) beams.	113
4.46	Normalised amplitude distributions for reconstructed data events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Blue dashed lines indicate the approximate aperture scraping limit.	115
4.47	Normalised amplitude distributions for reconstructed data events traversing the LH ₂ empty and LH ₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Blue dashed lines show the approximate aperture scraping limit. Blue dashed lines indicate the approximate aperture scraping limit.	115
4.48	Normalised cumulative amplitude distributions for reconstructed data events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Blue dashed lines indicate the approximate aperture scraping limit.	116
4.49	Normalised cumulative amplitude distributions for reconstructed data events traversing the LH ₂ empty and LH ₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Blue dashed lines indicate the approximate aperture scraping limit.	116
4.50	Normalised amplitude distributions for reconstructed MC events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions.	117

4.51	Normalised amplitude distributions for reconstructed MC events traversing the LH ₂ empty and LH ₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions.	117
4.52	Normalised cumulative amplitude distributions for reconstructed MC events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. . . .	118
4.53	Normalised cumulative amplitude distributions for reconstructed MC events traversing the LH ₂ empty and LH ₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions.	118
4.54	(Hollow) Uncorrected and (filled) corrected amplitude distributions for reconstructed data events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Where hollow markers are not visible, they coincide with the filled markers.	120
4.55	(Hollow) Uncorrected and (filled) corrected amplitude distributions for reconstructed data events traversing the LH ₂ empty and LH ₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Where hollow markers are not visible, they coincide with the filled markers.	120
4.56	Ratios of the amplitude PDF distributions for reconstructed data and MC events traversing the no absorber and LiH absorber states. The y-axis for the 3-140 beam is extended to accommodate the larger variation in PDF ratio. Coloured bands show combined statistical and systematic errors for (blue) measured data and (red) simulated MC. Blue dashed lines indicate the approximate aperture scraping limit.	122
4.57	Ratios of the amplitude PDF distributions for reconstructed data and MC events traversing the LH ₂ empty and LH ₂ full absorber states. Coloured bands show combined statistical and systematic errors for (blue) measured data and (red) simulated MC. Blue dashed lines indicate the approximate aperture scraping limit.	122
4.58	Ratios of the amplitude CDF distributions for reconstructed data and MC events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (blue) measured data and (red) simulated MC. Blue dashed lines indicate the approximate aperture scraping limit.	123

4.59	Ratios of the amplitude CDF distributions for reconstructed data and MC events traversing the LH ₂ empty and LH ₂ full absorber states. Coloured bands show combined statistical and systematic errors for (blue) measured data and (red) simulated MC. Blue dashed lines indicate the approximate aperture scraping limit. .	123
5.1	Normalised single particle distributions of reconstructed kinetic angular momentum components, L_{kin} , of the DS muon sample at TKU station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.	129
5.2	Normalised single particle distributions of reconstructed field angular momentum components, L_{field} , of the DS muon sample at TKU station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.	129
5.3	Normalised single particle distributions of reconstructed kinetic angular momentum components, L_{kin} , of the DS muon sample at TKD station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.	130
5.4	Normalised single particle distributions of reconstructed field angular momentum components, L_{field} , of the DS muon sample at TKD station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.	130
5.5	Normalised single particle distributions of reconstructed canonical angular momentum, L_{canon} , of the DS muon sample at TKU station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.	131
5.6	Normalised single particle distributions of reconstructed canonical angular momentum, L_{canon} , of the DS muon sample at TKD station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.	131
5.7	Normalised single particle distributions of the change in canonical angular momentum of the DS muon sample across the absorber for 3, 4, 6, and 10-140 beams. Reconstructed data events are shown in black, reconstructed MC events in yellow. .	132
5.8	Evolution of the mean kinetic angular momentum, L_{kin} , of the beam at each tracker station for beams 3,4,6, and 10-140. Black points show reconstructed data events, red points show reconstructed MC events, and purple shows MC truth variables for events in the MC sample.	134

5.9	Evolution of the mean field angular momentum, L_{field} , of the beam at each tracker station for beams 3,4,6, and 10-140. Black points show reconstructed data events, red points show reconstructed MC events, and purple shows MC truth variables for events in the MC sample.	135
5.10	Evolution of the mean canonical angular momentum of the beam at each tracker station for beams 3,4,6, and 10-140. Black points show reconstructed data events, red points show reconstructed MC events, and purple shows MC truth variables for events in the MC sample.	136
5.11	Evolution of \mathcal{L} through the cooling channel, calculated at each tracker station for beams 3,4,6, and 10-140. Black points show reconstructed data events, red points show reconstructed MC events, and purple shows MC truth variables for events in the MC sample.	137
7.1	Measurements of the hadron production cross-section around the Z resonance, with curves indicating the predicted hadronic production cross-sections for two, three, and four neutrino species.	143
7.2	Projection of spin, \bar{s} , on the direction of particle motion, \bar{p} , for a particle, q , showing right- and left-handed helicities.	144
7.3	The most recent constraints on the PMNS oscillation parameters δ_{CP} , $\sin^2 \theta_{13}$, and $\sin^2 \theta_{23}$ from the T2K collaboration.	147
7.4	The effect of different δ_{CP} values on $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ appearance probabilities as a function of neutrino energy for $\delta_{CP} = -\pi/2, 0, \text{ and } \pi/2$	148
7.5	A visual representation of the two possible neutrino mass orderings: (left) the normal hierarchy (NH) and (right) inverted hierarchy (IH). The relative proportions of red (ν_e), blue (ν_μ), and green (ν_τ) segments in each mass eigenstate box corresponds to the relative probability of finding each flavour eigenstate in that mass eigenstate.	150
7.6	Two-dimensional confidence intervals from T2K, NO ν A, and MINOS at the 90% and 99% confidence level for $\sin^2 \theta_{23}$ versus Δm_{31}^2 for normal ordering and inverted ordering.	150
8.1	The planned facilities and baseline of the DUNE and LBNF projects, depicting the near site at Fermilab and the far site at SURF.	152
8.2	The Neutrino beamline and DUNE near detector hall at Fermilab, Il.	154

8.3	Planned design and layout of the magnetic focusing horns in the target hall.	155
8.4	Illustrated layout of the subdetectors within the near detector hall. The beam points right to left, passing in order through ND-LAr, ND-GAr, then SAND.	157
8.5	Layout of the underground caverns housing the DUNE FD and cryogenics at SURF, South Dakota. The arrow shows the direction of the neutrino beam. The long vertical structure on the right shows the refurbished Ross Shaft, providing surface access.	159
8.6	Operating principle of a single-phase LArTPC module.	160
9.1	Neutrino energy spectrum as a function of parent pion energy at different off-axis angles away from the pion momentum, θ	162
9.2	FHC muon neutrino flux through each off-axis near detector 50 cm flux window. The colour axis is normalised to neutrinos per GeV per cm ² per POT.	165
9.3	FHC muon neutrino flux through each 4 m wide fiducial volume detector stop.	165
9.4	(Top) Flux matching between an oscillated FD flux (red) and the linear combination of ND fluxes (green) obtained by the linear least squares method shown in eqn. 9.2 and (bottom) their corresponding coefficients. The coefficients correspond to the weighting of each 50 cm flux window, with flux windows at larger distance off-axis having larger coefficient number.	167
9.5	(Top) Flux matching between an oscillated FD flux (red) and the linear combination of ND fluxes (green) obtained using Tikhonov regularisation with $\lambda = 10^{-8}$, and (bottom) the corresponding regularised coefficients.	169
9.6	(Top) Flux matching between an oscillated FD flux (red) and the linear combination of ND fluxes (green) obtained using Tikhonov regularisation with $\lambda = 10^{-9}$, and (bottom) the corresponding regularised coefficients.	170
9.7	An L-curve (top left) constructed from the evolution of the residual norm and solution norm terms, ρ and η , and its curvature (top right), with the resulting value of $\lambda = 3.192 \times 10^{-9}$ chosen producing the flux-matching (middle) and regularised coefficients (bottom).	172

10.1	On-axis ν_μ flux distributions for a 3 m fiducial width ND 575m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper row first, for FHC horn combinations A, AB, AC, ABC and BC respectively. . . .	176
10.2	Normalised on-axis ν_μ flux distributions for a 3 m fiducial width ND 575m downstream, generated for 5×10^8 protons on target (POT) each. All fluxes are normalised with respect to the nominal operating flux. Fluxes are shown, left to right, upper row first, for horn combinations A, AB, AC, ABC and BC respectively.	177
10.3	Off-axis ν_μ flux spectra produced by varying horn A current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA respectively.	178
10.4	Off-axis ν_μ flux spectra produced by varying horn A and B current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA respectively. . . .	179
10.5	Off-axis ν_μ flux spectra produced by varying horn A and C current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA respectively. . . .	180
10.6	Off-axis ν_μ flux spectra produced by varying horn A, B and C current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 0, 25, 50, 75, 100, 125, 150, 200, 250, 260, 270, 280, 293, 300, 310, 330 and 350 kA respectively.	181
10.7	Off-axis ν_μ flux spectra produced by varying horn B and C current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA respectively. . . .	182

10.8	Evolution of iteratively reweighted coefficients over 12 iterations for an oscillated FD flux-matching solution. Coefficients 270 to 318, shown between two black lines, show the variation of nominal PRISM fluxes with constant regularisation in response to the reweighting of contributing HC fluxes. Coefficient ranges for different horn configurations are shown in figure 10.9.	186
10.9	Coefficients resulting from the filtered solution of Fig. 10.8. Green lines separate coefficients corresponding to flux-windows from differing current powerings, red lines separate differing horn combinations.	186
10.10	Oscillated ν_μ FD fluxes showing the impact of different predicted values for (left) Δm_{atm}^2 and (right) $\sin^2 \theta_{23}$	188
10.11	Energy-weighted iterative flux-matching over a targeted energy range of 0-10 GeV utilising eight on-axis nominal fluxes spanning -1.5 m to 2.5 m off-axis position, combined with 4 m wide on-axis HC fluxes produced from utilising up to 293 kA horn current. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.	190
10.12	Energy-weighted iterative flux-matching over a targeted energy range of 0-10 GeV utilising eight on-axis nominal fluxes spanning -1.5 m to 2.5 m off-axis position, combined with 4 m wide on-axis HC fluxes produced from utilising up to 350 kA horn current. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.	192

10.13	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 32.5 m off-axis position, combined with 4 m wide HC fluxes up to 30.5 m off-axis produced from utilising horn currents up to 350 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.	195
10.14	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 32.5 m off-axis position, combined with 4 m wide HC fluxes up to 30.5 m off-axis produced from utilising horn currents up to 293 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.	196
10.15	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 32.5 m off-axis position, combined with 4 m wide HC fluxes up to 30.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.	197
10.16	Energy-weighted flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 32.5 m off-axis position, combined with a single 4 m wide on-axis 280 kA flux-window. The two panels show (top) The resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (bottom) solution coefficient values.	198

10.17	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 27.5 m off-axis position, combined with 4 m wide HC fluxes up to 26.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.	202
10.18	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 20.5 m off-axis position, combined with 4 m wide HC fluxes up to 18.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.	203
10.19	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 14.5 m off-axis position, combined with 4 m wide HC fluxes up to 12.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.	205
A.1	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 25.5 m off-axis position, combined with 4 m wide HC fluxes up to 22.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.	211

A.2	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 23.5 m off-axis position, combined with 4 m wide HC fluxes up to 22.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.	213
A.3	Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 18.5 m off-axis position, combined with 4 m wide HC fluxes up to 16.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.	215

Acknowledgments

I owe a huge debt of thanks to the friends, colleagues and loved ones who have guided, reassured and supported me throughout the last 4 years.

First, I have to thank my mum for supporting me through my life and encouraging me to always put in the work required to succeed.

For my supervisor, Steve, your attitude to the business of our research has always been extraordinarily encouraging, even as you presented an enduringly cynical view on each day. While lockdowns and office closures obfuscated the real world, your support and encouragement has been unwavering.

A huge thank you to my colleagues on the MICE collaboration who provided an atmosphere that was both warm and friendly and highly productive. Arriving at RAL for MICE shifts in my first week as a new student was hugely daunting and inspiring; were it not for the colleagues there, I'm sure those long nights would have been much harder to endure. There are many of you and I am grateful to you all. A particular debt of thanks to Paolo who made a huge effort in the early days to answer my endless emails, and Chris, for your guidance and leadership as head of the analysis team.

A great thanks to Mary, for readily inviting me to work closely with her at BNL, where her guidance and enthusiasm for research was ever encouraging.

For the friends I have shared an office with, whose camaraderie and guidance I will miss.

For Yash, Chris, and Sam, who made the long months of writing much more bearable.

And, of course, for Pub Wednesday, and all who sail in her.

Declarations

The material presented in this thesis has not been submitted for a degree at another university. The results presented here are the author's own work unless otherwise stated.

Chapter 1 is drawn from available literature and publications. Chapters 2 and 3 are derived from MICE publications, internal notes, and some literature; the characterisation of track reconstruction performance was produced by the author. The amplitude analysis in chapter 4 was first developed by Chris Rogers and the MICE Collaboration, with the validation and characterisation of 'Solenoid' mode data conducted by the author. The analysis in Chapter 5 was developed and conducted by the author.

Chapters 7 and 8 are drawn from the relevant literature, with technical descriptions of the DUNE experiment derived from information available in the collaboration's design reports. The PRISM technique was first developed by Mark Scott, with the application method described in chapter 9 developed by the DUNE-PRISM group. The analysis in chapter 10 was developed and conducted by the author, with input from the DUNE-PRISM analysis group.

Abstract

The Muon Ionization Cooling Experiment (MICE) was designed to provide the first demonstration of muon ionization cooling by measuring the reduction in phase-space volume induced in a muon beam passed through a low- Z absorber material within the MICE cooling channel. This technique has been validated by the MICE collaboration using a cooling channel in ‘flip’ mode, where the on-axis magnetic field reverses direction across the absorber. In this thesis, an analysis of the cooling effect in ‘solenoid’ mode, where the on-axis magnetic field points in the same direction throughout the channel, has been implemented, with the change in phase-space volume characterised by calculation of single-particle amplitudes. The ratio of amplitudes of muons before and after traversing the absorber material show a pronounced cooling effect for beams which observe a filled liquid hydrogen absorber or lithium hydride disk, with no equivalent cooling seen by beams traversing an empty vacuum or unfilled hydrogen vessel.

The Deep Underground Neutrino Experiment (DUNE) is a future long-baseline neutrino oscillation experiment, with a planned near detector site at Fermilab, Il., USA, and a far detector site under construction in South Dakota, USA. The DUNE-PRISM program aims to reduce sensitivity of the extracted oscillation parameters to interaction model uncertainties and detector effects by replicating the oscillated far detector neutrino energy spectrum using linear combinations of near detector measurements, utilising a series of near detector measurements at increasing angles off of the neutrino beam’s axis. To advance this technique, an analysis examining to what degree additional flux contributions obtained from varying the magnetic focusing horn currents can improve the agreement between linear combinations of near detector measurements and the far detector spectra has been conducted. Immediate improvement is obtained at energies above 3.5 GeV with the addition of a single additional flux contribution. Reductions in the maximum allowed off-axis distance can be compensated for with the addition of a small number of flux contributions.

Introduction

Accelerated particle beams composed of electrons, protons, or ions have been instrumental in developing our understanding of particle physics and have been utilised outside of the field to develop a greater understanding of many scientific disciplines. The equivalent utilisation of beams of muons or neutrinos has yet to be seen.

Present day muon beams [2–4], produced through the decay of pions, themselves produced from proton-target interactions, provide much lower brightness than their electron, proton or ion counterparts, limiting their usage in accelerator facilities. The production of high-brightness muon beams would enable their utility across a number of applications, including the realisation of a muon collider, a lepton-antilepton collider capable of an effective energy reach which greatly exceeds that of equivalent centre of mass energy hadron colliders [5, 6]. In addition, the decay of muons within a well-controlled beam produces muon and electron neutrinos with well-known flux and energy composition, and can act as the source for a Neutrino Factory [7]. Such a facility would allow for a huge suite of measurements to characterise neutrino parameters such as their interaction cross-sections and oscillation parameters with high precision, dramatically reducing the measurement uncertainties involved. Both of these facilities require novel techniques to provide a sufficiently high muon beam luminosity, as conventional methods to cool a muon beam (reduce the phase-space volume) are ineffective due to the particle’s short lifetime.

Neutrinos are similarly produced as a tertiary beam at accelerator sources, and coupled with the neutrino’s extremely low interaction cross-section, are particles whose properties have been poorly understood throughout most of the 20th century, with recent experiments only now beginning to provide measurements with the power to make sense of neutrino properties. The development of next-generation high-precision neutrino experiments is currently underway, but the measurements these experiments take will be limited by the systematic uncertainties they must control, with large contributions expected from absolute neutrino flux and interaction cross-sections. In lieu of the well-defined beams that would be produced by a Neutrino Factory facility, any analysis from these next-generation experiments must explore ways to combat

other uncertainties, such as those arising from neutrino interaction modelling.

The study and understanding of these beams and the development of techniques tailored to them is necessary to advance their usage in the next-generation of particle accelerator facilities.

Part I discusses the Muon Ionization Cooling Experiment's role in demonstrating a technique which can effectively cool a muon beam over the short timescales required and the results of this technique when applied to an ionization cooling channel with constant magnetic field polarity throughout.

Part II explores one such next-generation neutrino experiment, the Deep Underground Neutrino Experiment, and how a PRISM-style technique can be adapted to provide the most effective reduction in measurement bias arising from neutrino interaction mis-modelling.

Part I

MICE

The Muon Ionization Cooling
Experiment

Chapter 1

Background

Muon acceleration presents an appealing technology for future high energy physics experiments. Recent proposals have been made to construct a muon collider within the CERN LHC tunnel, aimed at exploiting the particular strengths of this facility [6]. A muon collider would offer a unique combination of point-like particle collisions, resulting in precise centre-of-mass energies, with a high energy reach due to the significant reduction in energy losses from synchrotron radiation¹ compared with an equivalent electron-positron collider. Such a facility could serve as a Higgs factory to provide precision measurements of the standard model Higgs resonance. Similar technology could be used to construct a neutrino factory such as NuMax, a facility designed to provide an intense and well-defined flux of neutrinos from muon decays [8]. By producing a well-defined beam in this way, the uncertainties associated with flux systematics would be significantly reduced, with an unparalleled neutrino beam intensity.

The most essential component in the development of next-generation muon facilities which remains untested is the cooling stage. To enable facilities such as a neutrino factory or muon collider, novel techniques must be developed which can collimate and focus a highly divergent muon beam within sufficiently short timescales. Figure 1.1 shows the potential layout of these facilities, with differing amounts of cooling required in each case.

Ionisation cooling is the only feasible process which can sufficiently cool a muon beam without experiencing significant losses due to muon decay. Whilst ionisation cooling has been considered as an effective method to cool particle beams since the 1960s, the method has never been practically tested. For any future facility which relies on this method to be built, the underlying physics principles which enable it must be validated and the expected design performance shown experimentally.

Our best understanding of the ionisation cooling process is contained within

¹Synchrotron radiation energy loss is proportional to mass^{-4} , meaning a $\times 200^{-4}$ reduction in losses for muons compared with electrons.

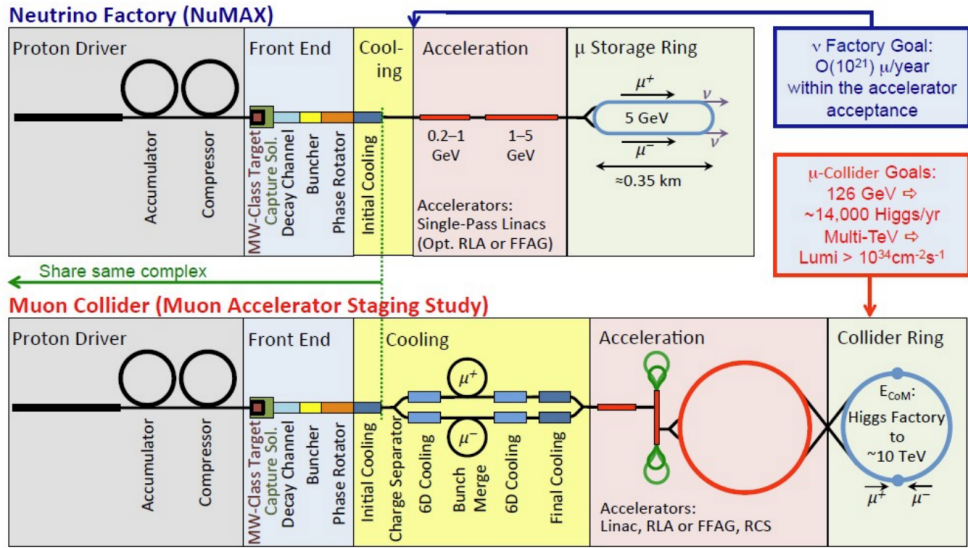


Figure 1.1: Block diagrams of neutrino factory and muon collider scenarios, showing the same shared complex up to the green dashed line. [1]

the Monte Carlo (MC) simulation toolkits used to model particle propagation and interactions, such as GEANT4. These same simulation toolkits are used to validate design choices for projects such as ν STORM (Neutrinos from Stored Muons) [9] and the Muon Accelerator Program [10]. As such, the validation of these simulation tools is an essential secondary objective of the Muon Ionisation Cooling Experiment (MICE), in addition to providing the first demonstration of the reduction of the transverse phase-space of a muon beam with ionisation cooling. In particular, multiple coulomb scattering of muons is not currently understood to high precision. The verification of muon multiple coulomb scattering algorithms will be crucial for accurate modelling of future muon ionisation channels. In the analyses to follow, comparison of simulated MC distributions and real data distributions will provide a means of assessing the validity of these models.

1.1 Beam Optics

To preface our understanding of muon cooling and the motivations of MICE, some discussion of beam optics is required. This will provide the means of characterising and quantifying any observed cooling effect.

1.1.1 Phase-Space and Trace-Space

Each particle in a beam can be completely described by the evolution of its 4D position and momentum vectors, (x, y, z, t) and (p_x, p_y, p_z, E) , where x and y are the transverse (perpendicular to the beam) spatial coordinates, z the longitudinal (parallel to the beam) spatial coordinate, p_x , p_y , and p_z are the corresponding momenta, t the laboratory reference frame time, and E the energy. The distribution of an ensemble of particles in a beamline can then be characterised by their coordinates in six-dimensional phase-space, parameterised as (x, p_x, y, p_y, z, E) , or in trace-space using their geometric coordinates (x, x', y, y', z, E) where $x' = p_x/p_z$ and $y' = p_y/p_z$ (other alternatives are often used). The particle density at any point in phase space is given by the distribution function of the beam. A simple example is a Gaussian distribution in each phase-space variable, with a corresponding multivariate-Gaussian representation at higher order. The distribution function of a 2D phase-space plane in its simplest form may be described by a bi-Gaussian with

$$f(x, x') = \frac{1}{2\pi\sigma_x\sigma_{x'}} \exp\left(-\frac{(x - \langle x \rangle)^2}{2\sigma_x^2}\right) \exp\left(-\frac{(x' - \langle x' \rangle)^2}{2\sigma_{x'}^2}\right), \quad (1.1)$$

symmetric around the means of x and x' , however this simple representation has no coordinate-momentum correlation. Such a distribution is described entirely by the mean and variance in each phase-space dimension. For a correlated n -dimensional multivariate-Gaussian distribution, the mean and variance generalise to a vector of means for each phase-space variable and an $n \times n$ dimensional covariance matrix.

In linear beam dynamics, an alternative representation of such a distribution is given by the Twiss parameters α_t , β_t , and γ_t , along with the geometric Root Mean Squared (RMS) emittance ε_x^{rms} [11]. Each parameter describes a collective property of the beam distribution related to the RMS beam ellipse at a given point, as highlighted in figure 1.2, capturing the correlation between x and x' in α_t , the physical beam size in β_t , or the spread of angular x' values in γ_t . In linear systems, the motion of particles in each plane is often largely decoupled and can be treated as three independent 2D phase-space planes of horizontal, vertical, and longitudinal coordinates. By considering each 2D phase-space plane separately, we can obtain the Twiss parameters for that plane, defined in Fig 1.2. This provides a useful form to describe the properties

of such an ensemble of particles. The Twiss parameters are also known as the Courant-Snyder functions and are connected by the relation

$$\gamma_t = \frac{1 + \alpha_t^2}{\beta_t}, \quad (1.2)$$

hence only two of these three parameters are independent [11].




Figure 1.2: The extrema, intercepts and correlation of a trace-space ellipse described in terms of the Twiss parameters.

The Twiss parameters can be obtained from the solution to Hill's equation within a periodic transverse focusing system [12]. Hill's equation:

$$x''(s) - k(s)x(s) = 0, \quad (1.3)$$

with $k(s)$ the focusing strength coefficient, has a periodic solution given by

$$x(s) = \sqrt{\epsilon} \sqrt{\beta(s)} \cos(\phi(s)), \quad (1.4)$$

where the phase $\phi(s) = \int ds/\beta(s) + \phi_0$, and $f(s)$ denotes f as a function of s , with s path length. Differentiating gives

$$x'(s) = \sqrt{\epsilon} \left[\frac{\beta'(s)}{2\beta^{1/2}(s)} \cos(\phi(s)) - \beta^{1/2}(s) \phi'(s) \sin(\phi(s)) \right], \quad (1.5)$$

which together with equation 1.4 and the periodic β -function belonging to the focusing system can be plotted over multiple periods to define a trace-space ellipse corresponding to the preferred trace-space of the periodic focusing system at that point. Where a beam's actual trace-space ellipse has the same orientation and aspect ratio as the focusing system ellipse, the beam is termed matched (see section 1.2) and the β and ϵ in the solution to Hill's equation

are the same as the Twiss parameters β_t and ε_x^{rms} .

The Twiss parameters appear organically elsewhere. For a 2D covariance matrix, $\Sigma_{\mathbf{x}}$, which describes the horizontal $x - x'$ 2D trace-space,

$$\Sigma_{\mathbf{x}} = \begin{pmatrix} \sigma_{xx} & \sigma_{xx'} \\ \sigma_{x'x} & \sigma_{x'x'} \end{pmatrix} = \begin{pmatrix} \text{Var}(x, x) & \text{Cov}(x, x') \\ \text{Cov}(x', x) & \text{Var}(x'x') \end{pmatrix} = \varepsilon_x^{rms} \begin{pmatrix} \beta_t & -\alpha_t \\ -\alpha_t & \gamma_t \end{pmatrix}. \quad (1.6)$$

The RMS ellipse which corresponds to a particular contour of constant probability density in the trace-space is then parameterised as

$$\gamma_t x^2 + 2\alpha_t x x' + \beta_t x'^2 = \varepsilon_x^{rms} \quad (1.7)$$

and contains the central 68% of the particle distribution. Furthermore, the RMS emittance of the beam describes the area of the ellipse divided by π and can be directly calculated from the determinant of the covariance matrix,

$$\varepsilon_x^{rms} = \sqrt{|\Sigma_{\mathbf{x}}|}. \quad (1.8)$$

In higher dimensional trace-space characterised by an $n \times n$ covariance matrix, Σ , the associated n -dimensional emittance, ε_{nD}^{rms} , is calculated as

$$\varepsilon_{nD}^{rms} = \sqrt[n]{|\Sigma|}. \quad (1.9)$$

For beams with de-coupled horizontal and vertical phase-space, the 4-dimensional emittance is simply related to the pair of 2D emittances ε_x and ε_y by

$$\varepsilon_{4D} = \sqrt{\varepsilon_x \varepsilon_y}. \quad (1.10)$$

In the case of a solenoidal magnetic channel such as the one utilised in MICE, the uniform longitudinal magnetic field imparts a perpendicular Lorentz force on beam particles with transverse components of motion, inducing a Larmor rotation around their gyrocentres. As a result, x-y correlations do exist in the ensemble, coupled by the non-zero mechanical angular momentum, L_{kin} , where

$$\langle L_{kin} \rangle = \langle xy' - yx' \rangle, \quad (1.11)$$

providing additional contributions to the 4-dimensional determinant. The importance of angular momentum in MICE is discussed in greater detail in chapter 5.

It is useful to consider the normalized transverse RMS emittance of the beam, ε^n , as it provides a statistical description of the extension of the beam in trace-space and is invariant under the combined effects of linear transverse forces and longitudinal acceleration. This is not the case for ε^{rms} since the mean

energy of the beam changes from acceleration or traversing some medium, and so the components x' and y' which are defined with respect to the momentum are no longer conserved. To deal with this, we define the normalised transverse emittance as

$$\varepsilon_{\perp}^n = \beta\gamma\sqrt[4]{|\Sigma|}, \quad (1.12)$$

with β and γ the usual relativistic quantities which compensate for the distortion of the trace-space ellipse under acceleration or energy loss. Both the normalised transverse emittance and the volume of the RMS ellipsoid are constants of linear motion in vacuum.

1.1.2 Solenoid Optics

Solenoid channel optics like those used in the MICE beamline operate with an inherent x-y coupling in the transverse plane arising from the Larmor rotation of particles as they propagate. Hence proper treatment of the beam optics requires covariance matrix characterisation in at least 4 dimensions. We can parameterise the covariance matrix Σ_{\perp} which describes the 4D transverse trace-space (x, x', y, y') in terms of the Twiss parameters as [13]

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

where α_{\perp} , β_{\perp} and γ_{\perp} are the 4D Twiss parameters, ε_{\perp} is the geometric 4D RMS emittance, $\kappa = qB_z/2p_z$, and \mathcal{L} is a dimensionless parameter related to the mechanical angular momentum L_{kin}

$$\langle L_{kin} \rangle = \langle xp_y - yp_x \rangle = -2p_z\varepsilon_{\perp}(\beta_{\perp}\kappa - \mathcal{L}). \quad (1.14)$$

The canonical angular momentum, L_{canon} , is more directly related to \mathcal{L} by

$$\langle L_{canon} \rangle \simeq \langle xp_y - yp_x + \kappa r^2 p_z \rangle = 2mc\varepsilon^n \mathcal{L}. \quad (1.15)$$

This canonical angular momentum is generally a conserved quantity in solenoidal optics and is given by the mean of the single particle canonical angular momenta,

$$\langle L_{canon} \rangle = \langle xp_y^c \rangle - \langle yp_x^c \rangle, \quad (1.16)$$

where p_x^c and p_y^c are the canonical particle momenta.

1.1.3 Liouville's Theorem

For any beam only subject to conservative forces such as solenoid or dipole fields, Liouville's Theorem states that the phase-space volume is conserved [14]. Where the phase-space distribution undergoes only linear transformations, this extends to the conservation of beam emittance (though non-linearities in beam optics can cause a measured emittance growth as highlighted in section 1.2). In this case, variations in the Twiss parameters correspond to an induced skew or rotation of the RMS ellipse while the occupied area remains constant.

Furthermore, by considering a beam undergoing acceleration, it is evident that the geometric emittance cannot be conserved. This arises from the non-canonical nature of the parameters $x' = dx/dz$ and $y' = dy/dz$, such that any increase in p_z results in an apparent shrinking of x' and y' . Instead, the normalised emittance (defined in equation 1.12), which is explicitly conserved, is more commonly discussed.

Liouville's Theorem demands the application of non-conservative forces to provide a viable beam cooling method.

1.1.4 Single Particle Amplitude

Amplitude characterises the distance of a point in normalised trace-space from a beam centroid. In a stable magnetic lattice², the trajectories of particles can be decomposed into a set of decoupled amplitudes, with one for each decoupled trace-space. These amplitudes are conserved to first order, with the mean of all single-particle amplitudes for a beam defining its emittance. The single-particle amplitude, A_i , in a decoupled 2D trace-space defines an ellipse contour of equal amplitude along which particles may migrate, similar to the RMS ellipse. The area of this ellipse is entirely determined by the amplitude A_i of the particle, with the single-particle emittance defined as

$$\varepsilon_i = \pi A_i^2 \quad (1.17)$$

in units of mm or $mm \text{ rad}$. This is often given with the π absorbed into the units as $\varepsilon_i = A_i^2$ in units of $\pi \text{ mm mrad}$. The RMS emittance of the beam is then proportional to the mean of the amplitudes.

For a coupled 4D transverse phase-space such as that generated in a solenoidal magnetic field, the single-particle transverse amplitude of a particle at a point $p = (x, p_x, y, p_y)$ within a beam ensemble is defined as

$$A_{\perp} = \varepsilon_{\perp}^n (p - \bar{p})^T \Sigma^{-1} (p - \bar{p}), \quad (1.18)$$

²A magnetic lattice is a composition of magnets along the path of an enclosed particle beam.

with \bar{p} the mean position of the beam in phase-space. Each transverse particle amplitude corresponds to a locus of points which define a hyper-ellipsoid in phase-space passing through the position of the particle.

With a matched beam³, A_{\perp} is a conserved quantity on average up to some dynamic aperture, with conservation even in the case of a non-matched beam in the linear approximation. This is because for a matched beam, particle migration from low to high amplitudes along phase-space curves is balanced by equal migration from high to low amplitudes. The influence of non-linear magnetic fields at large transverse amplitudes leads to non-linear particle motion which can result in eventual particle loss [12].

1.2 Beam Matching

Beam matching involves the modelling of a suitable magnetic lattice to transport and focus a beam such that the optical properties of the beam match the required phase-space ellipse of the beamline at the point of machine transfer. This is done by individual adjustment of each magnet so as to vary the focusing strength, modelled through a combination of linear optics and full particle-in-field simulations. Where beam matching occurs in the context of circular colliders, beams take multiple passes through the same magnetic lattice, and so matching traditionally involves ensuring the magnetic lattice and beam (Twiss) parameters are periodic so as to minimise beam losses over successive passes.

In the case of a linear beamline such as in MICE, it is still desirable to have a good match between the Twiss parameters exiting the transport section of the beamline and the cooling channel phase-space ellipse. Poorly matched beams are susceptible to ‘blow-up’ and suffer greater difficulties achieving the required focusing at the absorber, with particle losses through the cooling channel more likely.

When perfect matching with the machine ellipse does not occur, the emittance ellipse of the incoming beam will have a difference in shape and/or orientation. As the beam propagates, particles will rotate in transverse phase-space along the phase-space ellipse defined by the machine. For a strictly linear machine, this rotation occurs at the same angular frequency for all particles, but small non-linearities in the magnetic channel cause the frequency to depend on the amplitude of the particle. The phase-space distribution is deformed by this process, first into an S-shaped area, eventually effectively occupying a new larger ellipse, as in Fig. 1.3. After this process, known as filamentation, the beam will occupy a larger phase-space area with an increased RMS emittance.

Where a beam observes large non-linear behaviour at its extremities in

³A matched beam is one whose phase-space ellipse has the same shape, position and orientation as the magnetic lattice ellipse.




Figure 1.3: Trace-space evolution of beam filamentation from beam mismatching, producing non-linear emittance growth. The ideal matched ellipse is shown in red, with the actual beam ellipse in black.

the tails of the distribution, this effect is more pronounced and filamentation occurs more quickly. MICE is particularly susceptible to these effects due to the tertiary nature of the muon beam. As the phase-space volume of the MICE beam occupies a much larger volume than traditional particle beams, more particles occupy non-linear areas of the channel. These non-linearities can be further amplified by the additional complexity of the magnetic lattice employed.

Steering Errors

For a beam injected into a machine with an error in position and angle with respect to the design machine phase-space, the emittance ellipses of the beam are no longer concentric with the admittance ellipses of the machine. Instead, the emittance and admittance ellipses are radially displaced by an amount dependent on the steering error. Such steering errors lead to double-humped distributions in phase-space, as particles disperse along their respective phase-space ellipses.




Figure 1.4: Illustration of steering error, showing larger emittance ellipse achieved after filamentation.

1.3 Cooling

The production of muon beams commonly relies on a tertiary decay process, initiated with proton-target collisions at a proton accelerator facility. Proton-nucleus collisions produce secondary beams composed of pions, kaons, muons, and other particle species in smaller quantities. Meson decay in these beams produce significant numbers of muons diffusely populating large areas of phase-space. The tertiary muon beam is then itself at risk from muon decay, imposing a limited manipulation time before unacceptable losses are accrued. Typical accelerator facilities which operate electron or proton beams can produce beams at low initial emittances, avoiding substantial requirements for emittance reduction techniques, and when emittance reduction is necessary operate with much longer particle lifetimes. For effective muon storage and collider facilities, tighter constraints on the beam envelope's phase-space with many orders higher brightness are required [6], necessitating improved emittance reduction techniques.

Beam cooling is the process of reducing the phase-space volume occupied by an ensemble of beam particles without particle loss. As highlighted in section 1.1.3, the phase-space volume of a beam is invariant under the application of conservative forces. To circumvent Liouville's theorem, a number of solutions have been proposed and pioneered, each with its own limitations.

1.3.1 Radiative Cooling

Radiative damping was first suggested as a cooling mechanism in the 1950s [15], taking advantage of the synchrotron radiation produced by radially accelerated charged particles. The emission of synchrotron light by particles orbiting within a magnetic field is preferentially emitted in the direction of motion, creating an oppositional damping force and reducing both longitudinal and transverse momenta. With longitudinal momentum lost in the process restored by RF cavities, a net reduction in the transverse emittance of the beam is achieved. Due to the dependence of emitted radiative power on mass as m^{-4} , this technique only facilitates beam cooling of low-mass particles, at the scale of the electron mass. Due to this mass dependence, muon beams cannot effectively utilise radiative cooling, but are equally less susceptible to synchrotron losses in circular accelerators, enabling the higher collision energy reach expected at future muon colliders.

1.3.2 Stochastic Cooling

Stochastic cooling, invented by Simon van der Meer in the early 1970s, acts to reduce the energy spread and angular divergence of a charged particle beam

through application of successive corrective steering pulses to a circulating beam bunch [16]. To provide appropriately oriented/timed pulses, the beam profile is first measured using its electric charge distribution at an earlier position in the beamline, with a corresponding corrective pulse provided by a kicker magnet downstream. The measurement and corrective pulse normally occur at 90° separation in betatron phase such that a position displacement at the former becomes an angular displacement at the latter, where betatron phase is the phase, ϕ , of the periodic solution to Hill's equation, shown in Eq. 1.4.

This method is referred to as "stochastic" due to the statistical basis which underlies it. Each application corrects only the average of a small subsample of the beam and over many repetitions applied to successive samples the variance of beam phase-space properties are reduced, resulting in a reduced emittance.

The technique is applied at particle accelerator rings at facilities worldwide, including the Super Proton Synchrotron at CERN, with cooling times spanning a few seconds to minutes. Whilst the stability of protons allows hadron accelerators to provide progressive cooling over long timescales without experiencing significant beam losses, due to the short muon lifetime (only $2.2 \mu\text{s}$ at rest) cooling times of this length are insufficient to provide the level of emittance reduction required at future muon facilities as a significant number of beam muons would decay during cooling.

1.3.3 Ionisation Cooling

Ionisation cooling has undergone a half-century of development, first proposed by G.I. Budker and A.N. Strinsky in the 1960s [17] and later developed as an approach to muon cooling by D. Neuffer in the 1990's [18, 19]. This technique operates analogously to radiative cooling, with energy loss inducing a reduction in both transverse and longitudinal momenta and RF cavity re-acceleration restoring the longitudinal component in the direction of the beam only, effectively reducing the transverse geometric emittance of the beam. In ionization cooling, energy is lost as the beam particles traverse an absorber and ionise electrons in the material. The mean energy loss per distance travelled




Figure 1.5: Ionisation cooling principle shown in 3 steps.

follows the Bethe-Bloch formula [20, 21], and to first order obeys

$$\left\langle \frac{dE}{ds} \right\rangle \propto \frac{Z}{A} \frac{1}{\beta^2}, \quad (1.19)$$

with $\frac{Z}{A}$ the ratio of atomic number (Z) to atomic mass (A), and β the standard relativistic velocity. This relation informs the choice of absorber material, with the optimal energy loss resulting from a high value of $\frac{Z}{A}$ and a low β .

The cooling effect is counteracted by heating from Multiple Coulomb scattering (MCS) in the absorber increasing the beam's angular spread. Since the effects of MCS increase with Z , the absorber material choice favours low- Z materials such as liquid hydrogen, liquid helium and lithium hydride. These two contributions to the change in beam emittance over a path length s are described by [18]

$$\frac{d\varepsilon_{\perp}^n}{ds} = -\frac{1}{\beta^2} \left| \frac{dE_{\mu}}{ds} \right| \frac{\varepsilon_{\perp}}{E_{\mu}} + \frac{\beta_{\perp} (13.6 \text{ MeV})^2}{2\beta^3 E_{\mu} m_{\mu} c^2 X_0}, \quad (1.20)$$

with the first term on the RHS describing the emittance reduction due to cooling, and the second describing the emittance increase due to heating. Here X_0 is the radiation length of the absorber material. The equilibrium emittance, $\varepsilon_{\perp eq}^n$, is the point at which these effects are equal and opposite, and is given by

$$\varepsilon_{\perp eq}^n = -\frac{\beta_{\perp} (13.6 \text{ MeV})^2}{2\beta m_{\mu} X_0 \left\langle \frac{dE}{ds} \right\rangle}, \quad (1.21)$$

where 13.6 MeV is a fitted value from Highland's formula [22]. Note that $\left\langle \frac{dE}{ds} \right\rangle$ is always negative and so the equilibrium emittance is positive. Furthermore, a beam with small β_{\perp} will undergo less heating and so have a lower equilibrium emittance. This effect arises as the increase in angular spread from MCS is smaller relative to the angular spread of the beam.

While the principles of the method appear straightforward, the implementation of high accelerating gradient RF cavities immersed in multi-Tesla magnetic fields is challenging due to difficulties in RF breakdown suppression, with field lines extending well beyond the solenoids themselves [23]. However, ionisation cooling is the only cooling method which can operate on short enough timescales and satisfy the demands of a muon collider.

1.3.4 Frictional Cooling

Frictional cooling [24, 25] operates similarly to ionization cooling, utilising the momentum loss which occurs when passing muons through an absorbing material, but instead requires much lower energies, in the range of a few keV. At such low energies, the muon no longer ionizes atoms within the absorber,

instead losing energy through the excitation of absorber atoms, with μ^+ and μ^- differences arising from elastic scattering on nuclei and charge exchange reactions. Below roughly 10 keV, muons observe an increased stopping power with increasing energy. By application of a longitudinal electric field, an equilibrium energy is defined where muons experience equal energy loss as the energy gained from the electrostatic field. Recent designs have utilised gaseous absorbing media, with differing optimal materials for μ^+ and μ^- beams, but have been limited by low efficiencies.

1.3.5 Conclusion

Ionisation Cooling has never undergone real-world testing. The MICE experiment was built to study the physical effects of equation 1.20 and will be discussed in the next chapter.

Chapter 2

The Muon Ionisation Cooling Experiment (MICE)

The realisation of a muon collider or neutrino factory relies on the development of accelerator technology capable of sufficiently reducing the phase-space volume (emittance) occupied by the muon beam prior to injection into the succeeding storage ring or accelerator. The Muon Ionization Cooling Experiment (MICE) project was established to provide the first demonstration of emittance reduction via ionization cooling, to validate our current understanding of ionization cooling by measurement and comparison with current models, and to provide the groundwork for development of these future facilities.

2.1 Experimental Goals

To provide a demonstration of ionization cooling, MICE requires a beamline which can provide a well-controlled muon beam, a cooling channel with one or more absorbers, and a detector suite capable of reconstructing the change in particle properties through the cooling channel. In order to properly investigate Eq. 1.20, the beam delivered by the MICE beamline must be variable between different average incoming momenta and emittances. The beam configurations used are discussed in Sec. 4.1.

There are a number of beam quantities which would be present in an ideal beam [26]. The aim throughout MICE data-taking for each beam configuration was for the following criteria (or the best compromise between them) to be fulfilled simultaneously:

1. A matched beam in the upstream spectrometer solenoid, downstream of the diffuser. A matched beam is one where the curves in phase space are uniformly filled, such that if a muon moves to a different area in phase-space, another muon will replace it. The matching conditions are:
$$\beta_x = \beta_y = \beta = 2p_z[GeV/c]/(0.3B[T]), \alpha_x = \alpha_y = \alpha = 0$$

2. Beam size at the (upstream) diffuser is such that the rms width in x and y satisfies $\sigma_x^2 = \sigma_y^2 = (\varepsilon_N \beta) / \gamma$, with $\gamma = E / m_\mu$
3. Beam angular divergence at upstream diffuser surface ($\sigma_{x'}$ and $\sigma_{y'}$), increased by multiple Coulomb scattering in the diffuser ($\sigma_{\theta, MCS}^2$), is equal to the desired angular divergence in SSU, i.e, $\sigma_{x'}^2 + \sigma_{\theta, MCS}^2 = \sigma_{y'}^2 + \sigma_{\theta, MCS}^2 = \varepsilon^n / (\beta \gamma)$
4. Beam central momentum matches nominal momentum p_z once corrected for the propagated energy loss in materials encountered before reaching the central absorber.

Due to the difference in emittance between the horizontal and vertical planes, it is not possible to fulfil all matching conditions simultaneously, with an average match obtained from the 4D covariance instead. The success of this matching condition was further compromised by the failure of one of the downstream matching coils, M1 (matching coils are discussed in section 2.3.1).

2.2 MICE Beamline




Figure 2.1: The MICE beamline, Step IV. Quadrupoles are numbered Q1-9, Dipoles D1/2, Time-of-Flight detectors as TOF0-2. Ckov, KL and EMR refer to the Cherenkov, KLOE-Light Calorimeter, and Electron Muon Ranger detectors respectively. Trackers are differentiated as being upstream or downstream with respect to the absorber [27]. Layout (a) shows aerial view, (b) side-on view.

The MICE beam was produced at the ISIS proton synchrotron facility, where the MICE target was fed parasitically by dipping into the proton beam and extracting resulting pions into the MICE beamline. These pions were

transported towards the cooling channel, with the resulting tertiary muon beam produced by their in-flight decays. The MICE target's intersection with the ISIS synchrotron and the subsequent beamline are shown in figure 2.1.

2.2.1 MICE Coordinate System

MICE uses a set of local and global coordinate systems, with one local coordinate system for each detector system, each oriented equivalently with respect to the muon beam. The beam direction is aligned along the positive z -axis in each case, with x and y forming a right-handed coordinate system such that positive y corresponds to increasing height above floor level. The origin of each local system is centred on the beam axis in x and y , with z -position defined from the mean z -position of the detector. The z -axis may equivalently be referred to as longitudinal, and the x and y axes as transverse, both with respect to the beam. The global MICE coordinate system uses the same axis orientation, with origin defined from the central z -position of Dipole 2.

2.2.2 ISIS Proton Synchrotron

ISIS is an 800 MeV proton synchrotron based at the Rutherford Appleton Laboratory, producing beams of neutrons and muons. H^+ ions are passed through a RF quadrupole where they are accelerated into 665 KeV bunches 5ns long. A linac increases proton energy to 70 MeV, after which protons are injected into the main ISIS synchrotron ring. After injection, protons are accelerated to 800 MeV over a 10 ms acceleration cycle, comprised of $\sim 10,000$ orbits, before being emptied by a kicker magnet directed towards the two ISIS target stations, shown in figure 2.2. Collision with either of the tungsten targets at Target Station 1 or 2 (TS1 or TS2) is used to produce neutron beams, whilst a carbon target upstream of TS1 produces muons. During the 10 ms following each acceleration cycle, the focusing and bending magnet currents are reduced to their initial values in preparation for the next proton injection. The ISIS synchrotron repeats this proton acceleration process every 20 ms, providing a beam spill repetition rate of 50 Hz. During each acceleration cycle the proton beam initially fills most of the beampipe, with an initial vertical radius of ~ 67 mm, and shrinks to a radius of about 48 mm during acceleration.

2.2.3 MICE Target

The MICE target, shown in figure 2.3, was comprised of a magnetically operated titanium-alloy hollow cylinder of outer radius 2.975 mm, inner radius 2.275 mm, and length 528 mm. The target was contained in a vacuum attachment to the ISIS synchrotron beamline and operated by dipping parasitically into the ISIS proton beam during 1 in every 50 acceleration cycles, providing the MICE




Figure 2.2: Schematic of the ISIS proton synchrotron, showing the central acceleration ring, the MICE target, and target stations 1 and 2 (TS1 and TS2) [28].

particle spill at a rate of 1 Hz. This was later reduced to 0.8 Hz as every 5th MICE spill was skipped due to beam loss requirements.

Production of the desired muon flux required a penetration depth into the beam halo of at least 5 mm over a short (~ 2 ms) interval and hence a minimum travel distance of about 24 mm and an acceleration of around 900 ms^{-2} . Additionally, the target removal time was required to be under 10 ms to avoid the succeeding injection cycle. These requirements ensured no target obstruction of the wide beam envelope during early stages of each acceleration cycle and that proton sampling only occurred close to the maximum energy. To achieve this, the target entered the beam halo during the last 1-2 ms of each 10 ms acceleration cycle, sampling protons for less than half of the target dip parabola.

The driving force to generate this acceleration was provided by a linear actuator consisting of sintered neodymium-iron-bore magnets on the target shaft placed inside a twenty-four coil stator. Target operation under this accel-




Figure 2.3: Schematic of the MICE target [26].

eration can cause significant wear, frictional heating, and dust-production, with earlier iterations of the MICE target suffering increasingly shortened lifetimes due to dust-induced failure from bearing degradation [29]. Minimisation of shaft-to-bearing friction was then crucial to preserve the required performance over long lifetimes, where the surrounding vacuum environment meant standard lubrication was not possible/prohibited. As a result, radiation-hardened polyimide Vespel SCP-5000 was chosen as bearing material and contact surfaces on the target shaft were coated in Diamond-like carbon (DLC).

Precise target positioning was provided by an optical vane measurement system using readout from three (fibre-coupled solid state) lasers. The vane provided $150\ \mu\text{m}$ measurement resolution via 157 slots along its length, each of width $0.3\ \text{mm}$, affording continuous reliable dip-depth monitoring. Consistent dip-depth ensured each target actuation would result in a consistent energy profile and comparable number of particles per beam spill, usually approaching 100 particle triggers per spill. Due to long-term variations in the precise position of the beam edge and intensity of the beam halo, the target system was also required to be easily adaptable to these varying beam parameters. To accommodate these variations, a flexible user interface, RATS (Random Acronym for Target System), Fig. 2.4, was created to allow target operation by non-expert MICE operators who could immediately respond to proton beam drift.

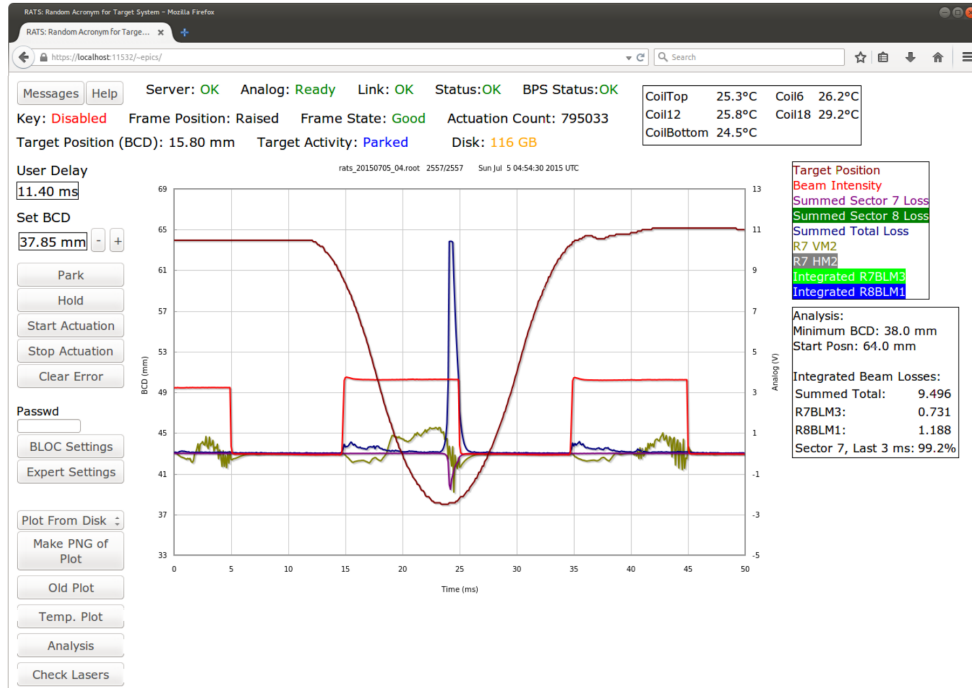


Figure 2.4: The Random Acronym for Target System (RATS) target control interface [30].

2.2.4 Beam Magnets

A series of focusing quadrupole magnets, labelled Q1 to Q3 in figure 2.1, provided particle capture for pions produced in the target. Following these, particles passed through two dipole magnets (D1 and D2) which directed the beam towards the MICE hall and provided momentum selection either side of the 5 Tesla superconducting decay solenoid (DS). Two more sets of quadrupole triplets Q4-6 and Q7-9 followed, separated by the first detectors of the MICE detector suite.

Quadrupoles

Quadrupoles 1-3 were Type-IV quadrupoles, featuring an effective circular aperture of 101.5 mm radius and a 200 A, 30 V power supply. As pion production at the MICE target occurred isotropically, Q1-3 were positioned at 3 m downstream to capture pions at the highest possible rate. Further downstream, Q4-6 and Q7-9 were instead Type-QC quadrupoles with effective circular aperture 176 mm radius and 400 A, 70 V power supplies.

Dipoles

The bending magnets, dipoles 1 and 2, were rectangular Type-1 dipoles, each with horizontal and vertical apertures 508 mm and 152 mm respectively. Both




Figure 2.5: Left: Kinematic limits for muon production from pion decays at MICE, showing (red) upper and (blue) lower momentum limits. Tuning D2 to the backwards going muon peak results in an extremely pure muon sample, with $< 1\%$ pion contamination. Right: Simulation showing (yellow) pion and (white) muon spectra after the decay solenoid. Dipole momentum selection when tuned to the backward-going muon peak is shown by the green band [26].

dipoles acted to bend the beam in the direction of the MICE hall and with careful tuning of D1 and D2, momentum selection of the beam was performed prior to any particle reconstruction. D1, positioned upstream of the DS, was used for momentum selection of a mostly pion beam entering the DS, whilst D2 performed selection of a mostly muon beam leaving the DS. To produce a high purity muon beam, backwards decaying muons could be selected using D2, reducing pion contamination in the beam to less than the 1% level, as in Fig. 2.5. For a higher particle rate, equal momentum selections in D1 and D2 were chosen, producing beams containing electrons, muons and pions.

Decay Solenoid

The decay solenoid (DS) was a long 5 Tesla superconducting solenoid magnet, so-named as its purpose was to increase the path length of particles traversing the MICE beamline and increase the proportion of muons in the beam. [26]

2.2.5 Proton Absorber

In the case of positively selected beam polarity, potentially large proton flux could traverse a significant fraction of beamline. To reduce proton backgrounds in the detector, variable thickness sheets of borated-polyethylene could be inserted into the the beam, with thicknesses tailored to the beam momentum. This would have had limited effect on pions and muons whilst removing a majority of the proton background before reaching/saturating the MICE detector suite/systems, due to the difference in proton and pion energy loss per unit length.




Figure 2.6: Left: Cross-section schematic of the MICE diffuser. Right: Photo of diffuser and the containing stainless steel drum being tested with iris petals in the closed position [26].

2.2.6 Diffuser

The diffuser was composed of a series of individually controlled variable thickness tungsten and brass irises, shown in figure 2.6, deployable in the upstream solenoid to diffuse the beam through multiple Coulomb scattering [26]. By preselecting a subset of these irises, different input beam emittances could be produced prior to the upstream tracker, enabling study of cooling performance over these beam emittances. As the diffuser was required to operate in fields of up to 4 Tesla, a conventional electromagnetic motor system was unsuitable for removing or inserting irises into the beam line; an air compressor system was instead used to close or retract the diffuser.

2.2.7 Partial Return Yoke (PRY)

As the MICE solenoids were discovered to produce substantial stray magnetic field, problematic for MICE hall equipment, the development of the Partial Return Yoke (PRY) was commissioned. The PRY, shown in figure 2.7, partially enclosed the MICE cooling channel (section 2.3) with a low carbon steel structure, reducing the stray magnetic field in the MICE hall by about 1 Gauss (0.1 mT) [31].

2.3 Cooling Channel

2.3.1 Steps I to VI

The cooling channel was the central component of MICE. Over the lifetime of the experiment, MICE has operated several different cooling channel configurations. The outlined evolution of the experiment, known as steps I to VI (fig 2.8), describe the planned construction and growth of the experiment in complexity, with each step facilitating a specific scientific goal. In step I, the MICE




Figure 2.7: The MICE Partial Return Yoke [31].

beamline was commissioned, supplying muons to calibrate a minimal set of PID detectors. The following step II introduced the first of the spectrometer solenoids, allowing for the optimisation of match quality between the beamline and cooling channel. Step III added the second spectrometer solenoid and enabled a cross-calibration of the two by comparison of measurements from particles traversing both spectrometers. The inclusion of ionization cooling absorbers would only arrive in step IV, enclosed within an Absorber and Focus Coil (AFC) module. Plans for steps V and VI included first one and then two sets of RF and Coupling Coil (RFCC) modules respectively, with each accompanied by an additional AFC module.

The data presented in this thesis was taken during Step IV, in what would come to be the channel's final configuration after the experiment was re-baselined. This re-baselining of the project was largely due to the cancellation of the US Muon Program in 2014 and the subsequent US Department of Energy's decision to cease US involvement in MICE in 2017, prompting other agencies to fail to extend MICE funding also. The change in scope removed the RF accelerating cavities, operating with a cooling channel composed of a pair of instrumented superconducting solenoid spectrometers (SSU upstream and SSD downstream) enclosing the beamline and trackers, and a single central absorber station with its own focus coil, the Absorber Focus Coil (AFC) module.

To provide the magnetic fields required to contain and control the beam, the channel operated 12 superconducting solenoid coils. The spectrometers utilised 5 superconducting coils each, with the AFC containing its own pair. The coils in each superconducting solenoid were designated as End2 (E2), Centre (C), End1 (E1), Match2 (M2), and Match1 (M1), in order of appearance as particles moved through the upstream cooling channel, with the downstream




Figure 2.8: The 6 proposed steps for study of muon ionization cooling at MICE, culminating in a cooling channel with 3 absorber modules, each separated by an RF module containing 4 RF accelerating cavities [26].

coils appearing in reverse order. The failure of the M1 SSU magnet due to an unplanned magnet quench in 2017 lead to the compromised operation of the channel under non-optimal non-matching configurations during the latter stage of MICE data-taking.

The solenoid magnets were capable of producing fields of 4 T, uniform to within 5% throughout the tracker and consistent at under 0.2% variation radially at each detector interface. The channel could be run in either ‘solenoid’ mode, where each magnet has the same polarity, or ‘flip’ mode, where the polarity alternates across the absorber. This configuration provided the first measurements of normalised transverse emittance reduction without re-acceleration due to ionization cooling.

2.3.2 Absorber and Focus Coil (AFC) Module

To improve cooling performance, beam optics through the cooling channel were designed to strongly focus the beam at the centre of the absorber, minimising the beta function there. The heating term within the cooling equation shown in eqn. 1.20 scales with β_{\perp} , such that a low β_{\perp} corresponds to a lower equilibrium emittance, as seen in eqn. 1.21. To achieve this the AFC module, which consisted of two superconducting coils in a housing assembly surrounding an absorber, was designed and operated. The coils are referred to as the upstream

and downstream Focus Coils (FCs). The assembly allowed for an absorber to be removed and replaced between runs, enabling comparison of performance across materials. Data taking took place with a 21-litre aluminium vessel of liquid hydrogen in both vessel-empty and vessel-full states, a 65 mm lithium hydride disk, a polyethylene wedge for 6D cooling, and an empty vacuum. The AFC was also required for running in ‘flip’ or ‘solenoid’ modes with appropriate field across the absorber.

2.3.3 ‘Solenoid’ and ‘Flip’ Modes

The MICE cooling channel took data in both flip and solenoid modes. In an extended cooling channel or an ionization-cooled muon beamline facility, ‘flip’ mode allows for the repetition of many cooling channel cells until a state of equilibrium beam emittance is reached, whereas ‘solenoid’ mode causes an increase in canonical angular momentum over successive passes, mismatching the beam to the lattice and reducing cooling performance [13]. In return for this mismatching, lower currents are required to operate at a comparable field strength. In a facility such as a neutrino factory or muon collider, a combination of field modes could be operated to compromise between engineering/power requirements and performance.




Figure 2.9: (Top) Solenoid field model B_z through the cooling channel, contrasting (solid) solenoid mode and (dashed) flip mode. Black lines show the field along the beam axis, green lines show the field at a displacement of 16 cm. (Bottom) Linear optics evolution of the RMS beam width through the cooling channel for nominal beams of 3, 4, 6 and 10 mm emittance. Blue dashed lines show the absorber and tracker station positions.

2.4 Detectors

The MICE beamline was fitted with a suite of detectors for beam monitoring, particle ID and particle-by-particle momentum-position reconstruction. The first of these when traversing the beamline was the luminosity monitor located in the ISIS synchrotron hall. TOF0, positioned much further downstream, after the bending magnets D1 and D2 and quads Q4-6, began the core MICE detector suite and along with TOF1 and TOF2 captured particle time-of-flight through the beamline. The remaining detectors were interspersed along the beamline after TOF0 following the layout shown in figure 2.1.

2.4.1 Luminosity Monitor

The Luminosity Monitor (LM) is the closest detector to the target station, situated within the ISIS synchrotron hall. The LM is positioned 10 m from the target interaction point at equal height to the ISIS ring, and is oriented at equal and opposite angle to the MICE beamline, $\pm 25^\circ$ relative to the ISIS ring. The monitor is constructed from two pairs of scintillators, each with its own readout PMT, with the larger scintillator pair positioned at greater distance from the target to compensate for angular divergence (Fig. 2.10). The two scintillator pairs are separated by a 150 mm thick polyethylene sheet which shields the latter PMT pair from low-energy protons and pions. By counting coincident triggers of each scintillation pair, as well as the four scintillator system, a comparison of simulated and real particle rate can be made.




Figure 2.10: Cross-section of the luminosity monitor, viewed side on relative to the MICE target particle spill [32].

2.4.2 Time-of-Flight (TOF) Detectors

Time-of-flight detectors (TOFs) produced time-of-flight measurements for particles triggering two or more TOFs, providing PID discrimination by measuring particle speed through sections of the MICE beamline. The first of these detectors, TOF0, was located immediately following quads Q4-6, with TOF1 positioned just upstream and TOF2 immediately downstream of the cooling channel. This layout is shown in figure 2.1. Each detector was composed of two planes of 25 mm thick scintillator slabs oriented orthogonally to one-another, with readout from scintillation light directed into photomultiplier tubes (PMTs) at both edges of each slab.

The slab width was 4 cm in TOF0 and TOF1 and 6 cm in TOF2, forming a total active area of $40 \times 40 \text{ cm}^2$, $42 \times 42 \text{ cm}^2$, and $60 \times 60 \text{ cm}^2$ for TOF0, TOF1, and TOF2 respectively. The TOFs were measured to have a timing resolution of 50 ps at TOF0 and TOF2, and 60 ps at TOF1, making use of leading edge timing discrimination above a given threshold to achieve such fine resolution. After refurbishing of TOFs 1 and 2, this was improved to a measured intrinsic time resolution of 55 ps for TOF0, 53 ps for TOF1 and 50 ps for TOF2 [33].

By combining the measurement from both slab planes, particle position when crossing the scintillator could be constrained in both x and y . This was improved by considering the additional signal arrival time required for light propagation from the particle crossing-point to the PMT, given by [26]:

$$t_i = t_0 + \frac{L/2 \pm x}{v_{eff}} + \delta_i; \quad (2.1)$$

where t_0 is the particle arrival time, x its distance from the counter centre, L the scintillator length, v_{eff} the effective velocity of light in the scintillator slab and δ_i the combined time delay from cables, PMT transmit time, and other causes. The transverse impact position, u , of a particle on a TOF station was reconstructed from the difference between time measurements of the two PMTs i, j at the ends of a given counter as:

$$u = \frac{v_{eff}}{2} \times ((t_i - \delta_i) - (t_j - \delta_j)). \quad (2.2)$$

Translation into MICE coordinate system is simple, with $x = u$ for vertical slabs and $y = u$ for horizontal slabs, and this procedure gives spatial resolution for the crossing point of a particle of approximately 1 cm.¹

Coincident slab hits are reconstructed as a TOF spacepoint, and the corrected time-difference between two TOF stations, e.g. $t_{TOF1} - t_{TOF0}$,

¹Spatial resolutions in x and y are given by $\sigma_x = \sigma_y = \sqrt{2} \times v_{eff} \times \sigma_t \approx 1 \text{ cm}$, with σ_t the time resolution for the TOF station in question.




Figure 2.11: Blown-out aerogel Cherenkov counter, showing a) entrance window, b) mirror, c) aerogel mosaic, d) acetate window, e) GORE reflector panel, f) exit window, and g) 8 inch PMT in iron shield [26].

henceforth referred to as TOF01 (and similarly for TOF02), gives the particle speed. This is combined with the tracker reconstruction (see section 2.4.4) and the approximate momentum as set by dipoles D1 and D2, providing an estimate of particle mass and discrimination of pion or electron events. The resulting pion rejection is in excess of 99% [26].

2.4.3 Cherenkov Detectors

The pair of Cherenkov detectors, CkovA and CkovB, follow TOF0 as shown in figure 2.1. A blown-out schematic of the detector design is shown in figure 2.11. The two Cherenkov detectors were designed to provide $\mu - \pi$ separation at higher momenta where TOF separation struggles to provide particle discrimination. The detectors were constructed with differing aerogel densities of 0.225 gcm^{-3} and 0.370 gcm^{-3} for CkovA and CkovB respectively, which therefore have refractive indices of 1.07 and 1.12. The resulting difference in muon thresholds of $p_{\mu,a}^{th} = 278 \text{ MeV}/c$ and $p_{\mu,b}^{th} = 210 \text{ MeV}/c$ and in pion thresholds of $p_{\pi,a}^{th} = 367 \text{ MeV}/c$ and $p_{\pi,b}^{th} = 277 \text{ MeV}/c$ in Ckovs A and B respectively provided pion discrimination using these detectors. For the 240 MeV/c beams, pions are above threshold only in CkovB, whereas muons are above in both Ckovs A and B. For 200 MeV/c beams, pions are below threshold in both cherenkov detectors, whilst the higher momenta muons remain above threshold for CkovB. Note that particles traversing the Cherenkov detectors have higher momenta than the nominal beam momentum within the cooling channel. Overall for muons in the momentum range 210 MeV/c to 365 MeV/c,

the two Cherenkov detectors guarantee better than 99.7% purity [26]. In the case of muons below 210 MeV/c threshold where there is no $\mu - \pi$ separation, the TOFs provide good separation instead.

2.4.4 Scintillating Fibre Trackers

MICE operates two scintillating fibre (SciFi) trackers, each contained within the bore of the superconducting solenoid which provides a uniform magnetic field. The upstream and downstream trackers are referred to as TKU and TKD respectively. Each tracker is composed of 5 stations of scintillating fibre structures, with each station constructed from three distinctly oriented SciFi planes. Particles propagating through the solenoid induce scintillation light in multiple fibres in each station, providing (x, y, z) particle positioning for each tracker station. Combination of each station measurement provides a particle track through the tracker.

For particle events recorded in the presence of the uniform solenoid field, spatial positioning at each station allows reconstruction of a helical track and hence particle momentum. For events recorded without an active solenoid field, spatial positioning reconstructs straight tracks (excluding scattering effects) and hence contains no momentum information. The design and construction of the MICE tracker is covered in greater detail in chapter 3.

2.4.5 KLOE-Light Calorimeter

The KLOE-Light Calorimeter (KL) was a sampling electromagnetic calorimeter (ECal) based on the design of the ECal used by the KLOE experiment at Frascati National Laboratory (LNF) in Italy. The KL was composed of 21 horizontal cells of sandwiched scintillating fibres within extruded lead foils (see Fig. 2.12), with an active volume of $93 \times 93 \times 4 \text{ cm}^3$. Scintillation




Figure 2.12: Cross-section of the (blue) scintillating fibres embedded in extruded lead within the KL [34].

light was read out by Hamamatsu PMTs positioned on both sides of a cell. The KL performed pre-showering and provided energy deposition and timing information, permitting discrimination between muons and decay electrons for higher energy beams. As the lead foils used in the KL were significantly thinner than those in the KLOE detector, the MICE detector was dubbed KLOE-Light.

2.4.6 Electron Muon Ranger

The Electron Muon Ranger (EMR) was a fully active scintillating tracking calorimeter, positioned at the downstream end of the beamline. The detector was composed of 48 planes of triangular scintillator bars with each plane arranged from 59 bars as shown in figure 2.13. Adjacent planes were rotated by 90° , alternating in orientation between parallels along the x-axis and y-axis. Scintillation light was collected by fibres which transport the signal to Hamamatsu photo-multiplier tube readout.

The function of the EMR was to provide PID through electron and muon discrimination, distinguished from one another based on their track topologies. Muons leave clear, straight tracks in the EMR whilst electrons make tracks which exhibit strong scattering in the scintillating materials. Within a mixed particle test beam of momentum range spanning 100 to 400 MeV/c, operation of the EMR alone provided a 98.6% efficiency in electron identification, resulting in a muon purity of greater than 99.8% [36]. Additionally, low energy muon calorimetry is provided from the range of stopped muons in the detector, identified from a clear Bragg peak exhibited in each case.




Figure 2.13: Blown-out image of the construction of the EMR [35].

2.5 MICE Software Framework

2.5.1 MICE Analysis User Software (MAUS)

The MICE Analysis User Software (MAUS) is the collaboration's particle tracking, detector reconstruction and physics analysis software framework, succeeding the G4MICE software package [37] used in the early stages of the experiment. The package is designed for Monte Carlo (MC) simulation and digitisation, shared event reconstruction for both real data and MC, and simple analysis. Any higher level analysis presented here is built on the MAUS software's event reconstruction.

The core functionality of MAUS is based on the GEANT4 toolkit's physical modelling and particle tracking processes. Detector electronics response for interacting particles within each active detector volume is also modelled in MAUS, providing the raw detector output when simulating events. The philosophy of the MICE reconstruction is such that data events and simulated events are not distinguished between by MAUS, and hence once digitised, MC events are reconstructed in the same manner as raw true data, with both built into higher level objects via the same algorithms at each detector. To this end, each detector has its own algorithms to digitise and treat raw data based on actual detector performance and testing. The higher level output from each detector is then combined to provide a global particle track.

Output MAUS reconstruction objects are packaged in ROOT format for ease of analysis [38]. Data in ROOT files are subdivided into spills, a discrete block in the MAUS data structure. This structure is derived from the real-world MICE particle source using the ISIS proton synchrotron where each MICE target dip corresponds to one particle burst in MICE, referred to as a spill.

2.5.2 Simulation

Beam simulation within MAUS is initialised from a separate target simulation output file. Target simulation is performed in external software² which provides better accuracy in strong force interaction cross-sections and proton-nucleus interactions than is native to GEANT4. The resulting particles are handed over to a GEANT4 beamline simulation a short distance downstream of the target and are propagated through the MICE beamline by GEANT4 until reaching the D2 magnet, after which point the remaining detector simulation and particle propagation occurs in MAUS.

To ensure simulated MC measures a muon ensemble which resembles real data, appropriate selection of the dipole fields of D1 and D2 is required. Simulation studies have shown the resulting muon distribution within the

²Nominally target simulation occurs in MARS, though other software has been utilised.

tracker is largely independent of the MICE target simulation due to the momentum selection the dipole pair imposes. As a result, particle selection in D1 and D2 must be tuned to produce muon distributions in the tracker which match those measured in real data for a given beam. To better characterise the dipole field response as a function of current, several hall-probe surveys of the MICE dipole fields were conducted, with the resulting current-field relationship utilised in producing MC beams. As simulations produced utilising the survey fields can leave some discrepancy between simulated and real muon ensembles, corrections are then applied to the D1 and D2 field strengths by iteratively simulating and adapting these field parameters. This is done manually and ensures that the measurement of emittance change across the absorber considers equivalent distributions of particles for real and simulated events.

2.5.3 Reconstruction

During reconstruction, digitised detector output is converted into spacepoints, with the utilised subset of hit parameters dependent on the detector being reconstructed. For the Tracker and EMR, these spacepoints are then combined into tracks constituted by several trackpoints and the derived momentum parameters of the track. Global track reconstruction is then used to combine tracks and spacepoints from each detector corresponding to a single particle's propagation through the beamline, providing one complete track which spans the MICE detector suite.

Descriptions for each detector-case, in order of the beamline layout, are given below.

TOFs

TOF reconstruction first requires each slab hit being considered to have coincident readout from PMTs at either end of the slab, discarding hits where this is not true. PMT readout from both sides is then averaged to determine hit time and total energy deposition. After evaluating all slab hits in this manner, coincident combinations of hits in x and y slabs are considered as candidate spacepoints. Calibration corrections are applied to account for cable lengths and correct for "time walk" arising due to different PMT response for differing PMT charge depositions. This calibration is also used to inform the evaluation of spacepoints. Further spacepoint hit parameters can be derived as in section 2.4.2.

Cherenkov Detectors

Cherenkov reconstruction evaluates the times t_1, t_2 between which the charge deposition recorded by the ADC exceeds a threshold, loosely defining the

beginning and end of the Cherenkov pulse. The ADC count is then integrated between $t_1 - 8 \text{ ns} \leq t \leq t_2 + 16 \text{ ns}$, whereupon a total photoelectron count is obtained by subtraction of the pedestal and normalisation [39].

Trackers

The MICE tracker reconstructs helical or straight tracks for events with solenoid magnetic field active or inactive respectively. Each track is constructed using spacepoints at each tracker station, themselves assembled by combining hits from individual neighbouring SciFi planes. The tracker and its reconstruction algorithms and performance are covered in greater detail in Chapter 3.

KL

KL spacepoints are reconstructed where PMT readout has occurred on both sides of a slab, with slab segmentation providing positioning only in yz . Where a trigger contains spacepoints in multiple slabs, the global reconstruction combines adjacent spacepoints and produces an average y -position, weighted by charge deposited in each slab.

EMR

Similar to the KL reconstruction, each slab provides two spatial coordinates, either in xz for vertical or yz for horizontal slabs. The hit position is given by the centre of mass of the slab which recorded the hit, and by combining all recorded slab-hits through a polynomial fit, a complete track is reconstructed. Information on the density of slab hits through the EMR is also utilised, providing a PID variable for distinguishing between electron and muon tracks due to the electron's proclivity to shower.

2.5.4 Configuration Database

The MICE configuration database (CDB) has operated as an online MySQL database allowing users to download physical detector, magnet and cooling channel geometry files based on surveys at each stage of data-taking. This has been needed to accommodate the many changes in beamline geometries and absorber configurations over the various MICE steps, with updates logged in the CDB for centralised access/monitoring/recording. These geometries are integral to the reconstruction of events and allow MC simulation to be performed analogously to true data. In addition, magnet operating currents and changes to detector calibration are also stored on the CDB to inform field modelling and event reconstruction.

2.5.5 Controls, Monitoring and Data Acquisition

The controls and monitoring system used in MICE was based upon EPICS [40] (Experimental Physics and Industrial Control System) and integrated with the DATE [41] (Data Acquisition and Test Environment) system, originally developed for the ALICE experiment. This system allowed for straightforward beam selection, adjusting diffuser and beamline magnets automatically to the appropriate predetermined settings, whilst preparing the DAQ for data. Additionally, this system provided monitoring and alarm-handling for hardware components during operation, ensuring the environment in the MICE hall is stable.

The data acquisition system used by MICE was required to perform at an acquisition rate of 600 triggers per 1 ms spill, ensured by maintaining the Front End Electronics (FEE) digitisation time at under 500 ns. Once digitised and stored in buffer memory, data was then read out via optical link and written to hard-disk in the 1 s interval between this and the next spill.

Chapter 3

Tracker

The MICE scintillating-fibre tracker is the primary detector in the emittance evolution analysis, responsible for high-precision particle position and momentum reconstruction for helical muon tracks in a uniform solenoid field, thereby enabling the comparison of upstream and downstream beam properties. Additionally, straight track reconstruction provides track position and direction information for configurations without a magnetic field.

3.1 Tracker Construction

The tracker design consists of two identical instrumented cylinders, positioned upstream and downstream of the absorber, with the upstream tracker rotated 180° about the y -axis relative to the downstream tracker. The cylinder body was constructed mostly from carbon fibre, with detector planes assembled from scintillating fibres glued in a doublet layer with a Mylar coating. The assembled frame and fibres are shown in figure 3.1.

The instrumented components of each tracker were made up of 5 stations, with each station composed of 3 planes of scintillating fibres, referred to as planes U, V, and W. Each plane is composed of two layers offset with respect to each other (a "doublet layer") as shown in figure 3.2. Each plane in a station is rotated 120° with respect to its neighbours such that every point in the active tracker x - y plane is overlapped by at least one fibre from each of the three measurement planes. Using any combination of 2 planes, an x and y position can be deduced. The redundancy of the 3rd additional plane is significant in reducing the impact of hits detected due to noise as a required coincidence of the two neighbouring planes can be imposed.

Fibres were grouped into bunches of 7 to form a channel, with each channel connected to a Visual Light Photon Counter (VPLC) read out by Analog Front End with Timing (AFEIt) boards donated from the D0 experiment [43, 44]. VLPCs function as photo-multiplying silicon avalanche detectors




Figure 3.1: Left: Schematic of the carbon-fibre tracker frame. The carbon-fibre station frames where the scintillating fibre planes are glued into place are shown in green. The z -orientation and hence station ordering shown corresponds to the downstream tracker. Fibre layout is oriented such that the longitudinal tracker coordinate z_t increases as one moves from the fibre plane towards the station frame it is attached to. Right: A photograph of the tracker prior to installation. Wavelength shifting fibres have not yet been connected [42].

with high gain, high quantum efficiency and multi-photon counting capability, able to detect and differentiate between single numbers of photons within the visible light wavelength range. As the scintillation light produced from ionising radiation in the fibres was typically of shorter wavelength than the visible spectrum, wavelength-shifting fibres were used to connect each channel to its VLPC. Fibres were mirrored on the unconnected end to maximise their light yield.

The VLPCs required calibration to provide the sensitivity to the low photoelectron count produced by an incident muon and maintain operation with manageable levels of noise hits. To meet these sensitivity requirements, the VLPCs were liquid-Helium cooled within cryostats maintained at 9 K. [45]

Due to the rotated layout, the upstream tracker first detects the beam at station 5, whereas in the downstream tracker this happens first in station 1. In both trackers, station 1 is the station closest to the absorber. The tracker spacing between stations (20, 25, 30, and 35 cm) is optimised to reduce degeneracies in track reconstruction for muons traversing the tracker within the MICE beam momentum range.

3.2 Track Reconstruction

Track reconstruction for the scintillating fibre tracker detectors uses a 5-step process building higher-level objects using the output of each prior step, with the final output a complete track through the US or DS tracker.




Figure 3.2: (a) Arrangement of the doublet layers U, V, and W in each scintillating fibre station. The outer circle describes the 400 mm diameter warm solenoid bore whilst the inner circle shows the limit of the station's active tracking area. (b) Layout of each planes' doublet layer of scintillating fibres, with fibres in red illustrating a channel of 7 ganged fibres. Units are given in μm [42].

3.2.1 Digitisation

The scintillation light registered in every tracker channel which generated a non-zero Number of Photo-Electrons (NPE) for each event is passed to the digitisation routine. A minimum NPE requirement for each channel is applied to remove electronics noise, with this minimum usually requiring 2 PE. Channels which pass this selection are saved as a Digit, the simplest reconstruction object in the data structure. The NPE, raw Front-End Electronics (FEE) signals and calibration data are also stored in this object.

3.2.2 Clustering

Due to the geometry of the tracker planes, particles passing through the tracker can produce sufficient light in neighbouring channels to form digits in both. The clustering process recombines neighbouring digits such that two-channel hits are considered as a single object alongside one-channel hits. To achieve this, a clustering algorithm is applied, plane-by-plane, to all digits in that plane from largest to smallest NPE, combining the largest neighbouring digit, if any, to form a cluster. This constructs clusters such that they produce the maximum signal possible whilst containing no more than two digits each. Where more than two digits are next to one-another, the lowest signal digits are evaluated recursively as separate clusters by this process.

A second minimum NPE requirement can then be applied to clusters produced at this stage to further reduce detector noise, but this is not applied by default as low signal noise clusters are statistically probable and noise contribution to reconstructed tracks has been shown to be small [46].

3.2.3 Spacepoint Formation

Spacepoints are formed by combining coincident sets of clusters from neighbouring planes, producing objects with 3-dimensional position information. The stations are evaluated in turn, combining the one-dimensional measurements from up to 3 clusters over multiple planes to deduce two-dimensional x, y components given our knowledge of the tracker geometry. Where coincident clusters in two planes are combined to form a spacepoint, they will be referred to as "doublets", whilst three-plane spacepoints will be referred to as "triplets".

Triplet spacepoints are prioritised in the reconstruction over doublet spacepoints due to the additional power gained in rejecting noise-produced clusters due to the low probability of three coincidental noise clusters. This is further improved by utilising Kuno's conjecture, providing an efficient method to determine if overlapping channels form a viable triplet spacepoint. Kuno's conjecture states that the sum of α 's (the distance of the cluster from each plane center within that plane's coordinate system) for each of the three clusters is zero if they are coincident. An illustration of this is shown in figure 3.3.

In MAUS, the channel ID is used to determine the relative position of each cluster, corresponding to the central position of the channel or channels from which it is formed. In this arrangement, Kuno's conjecture is equivalent to specifying that the sum of channel numbers obtained should be equal to half the total number of channels. The probability of three coincident noise-produced clusters is then negligible.

Doublet spacepoints are constructed by analysing each unique pair of planes in turn, with a spacepoint accepted for every unique cluster combination which produces a spacepoint contained within the tracker station's fiducial radius.




Figure 3.3: Illustration of Kuno's conjecture along planes U, V, and W. Summing the signed distances from the center of each fibre to the plane centre, α_U, α_V and α_W , will equal zero for any three simultaneously coincident fibres.

3.2.4 Pattern Recognition

The pattern recognition (PR) step combines multiple spacepoints distributed throughout the tracker to verify whether they are produced by a viable single muon track. As the path the muon follows in the tracker is dependent on magnetic field strength, this process uses separate algorithms for straight tracks and helical tracks, each applied in the same manner but with a different linear least squares fit.

The list of spacepoints reconstructed through the tracker is iterated through combinatorially such that one spacepoint is used from each station. Where multiple spacepoints occur in one or more stations, all combinations are assessed. A minimum of three stations with spacepoints is required for straight tracks, a minimum of five stations with spacepoints required for helical tracks, with five preferred in both cases. Candidate tracks are then assessed for goodness of fit using a linear least squares fit, the output goodness-of-fit χ^2 of which is compared to select the best fit from all candidates. The resulting PR track, if found viable, is then used to seed the final stage of track reconstruction.

Additional requirements can be included to reject tracks with too large values of χ^2 or spacepoints which deviate too far from the best fit track, thereby further reducing the contribution of noise hits to the track reconstruction.

Straight Track Minimisation

For straight track reconstruction, the pattern recognition algorithm decomposes the track into two independent transverse planes, such as the x - z and y - z planes. The track model is then parametrised by straight line fits in both planes.

Helical Track Minimisation

Helical track reconstruction decomposes the fit into separate coordinate planes transverse and longitudinal to the magnetic field lines. In the transverse plane, a helical fit reduces to a circle fit, with a non-linear sinusoidal fit then required in the longitudinal plane. To avoid the non-linear sinusoidal fit, the circle fit is used to estimate the rotation through the helix arc between spacepoints. A linear straight-line fit is then performed in the plane of the circular arc path-length, s , with respect to z -position, the s - z plane.

3.2.5 Kalman Filter

The final stage of track reconstruction utilises a custom Kalman filter to perform the track fit [42], accounting for all errors and correlations of the given measurement. This algorithm is an optimal linear estimator, meaning that for a linear system no fitting algorithm can perform better. As the system is

only approximately linear for helical tracks, an extended Kalman filter was implemented which features an analytic prediction of each track point using an appropriate non-linear propagator function, whilst the covariance matrices are predicted using a first order expansion approximation around the predicted position. The helical track propagator function is parameterised from the following equations of motion:

$$\begin{aligned}
x' &= x + \frac{p_x}{p_t} R \sin(\Delta\theta) - \frac{p_y}{p_t} R(1 - \cos(\Delta\theta)), \\
y' &= y + \frac{p_y}{p_t} R \sin(\Delta\theta) + \frac{p_x}{p_t} R(1 - \cos(\Delta\theta)), \\
z' &= z + \Delta z, \\
p'_x &= p_x \cos(\Delta\theta) - p_y \sin(\Delta\theta), \\
p'_y &= p_y \cos(\Delta\theta) + p_x \sin(\Delta\theta), \\
p'_z &= p_z,
\end{aligned} \tag{3.1}$$

where the prime coordinates correspond to their propagated values at an arbitrary Δz distance along z , R denotes the radius of the helix, and $\Delta\theta$ describes the amount of rotation of the particle through the helical path. Both R and $\Delta\theta$ depend on the magnetic field strength, with

$$R = \frac{p_t}{qB_z} \tag{3.2}$$

and

$$\Delta\theta = \frac{cB_z q \Delta z}{p_z}. \tag{3.3}$$

For straight tracks, the propagation instead exploits a straight line extrapolation between stations.

The Kalman filter itself is made up of three stages: prediction, filtering, and smoothing. The prediction stage propagates an estimate of the track state between measurement planes, given the current state, the model of the system and the process noise. The filtering stage makes a weighted adjustment to the predicted state given the measurement and its errors. The smoothing stage occurs once the track state has been propagated to the final measurement plane, propagating the now optimal estimate for the final state backwards through the previous measurement planes. Each track state is adjusted using all the measurement information in this way.

The reconstruction accounts for physical processes such as energy loss and multiple Coulomb scattering in the tracker planes and absorber through modelling in the propagator and process noise. Similarly, measurement noise is modelled from the statistical spread of measurements within the measured channel. Pattern recognition parameters are used to provide the seed for the

Kalman fit, with the set of clusters associated with the track given as the measurement data. The resulting track contains the reconstructed position-momentum information for the particle at 5 trackpoints, each of which coincide with the corresponding tracker station.

3.2.6 Track Reconstruction Performance

The success of the track reconstruction routine can be evaluated with an analysis of residuals from simulated MC events. Given an MC beam of similar phase-space distribution to the real data beam, simulated events allow for a comparison of reconstructed phase-space variables against true phase-space variables for each track. The phase-space variables of each particle track can be evaluated at each of the tracker stations. However the most important are those at Station 1 as this is the closest measurement plane to the absorber and therefore the least susceptible to distorted cooling performance due to energy loss and MCS effects.

The simulated MC beam samples described in section 4.2 have been utilised to produce the following reconstruction residuals at Station 1 in TKU and TKD. The beam sample labels ‘3-140’, ‘4-140’, ‘6-140’, and ‘10-140’ denote beams with nominal upstream emittances of 3, 4, 6, and 10 mm respectively and momenta consistent with 140 ± 5 MeV/c. The residuals for trackpoint position reconstruction are shown in figures 3.4, 3.5, 3.6, and 3.7. Residual distributions are peaked around zero, with most events accurately reconstructed to within 1 mm.

The residuals for transverse momentum reconstruction in the upstream tracker are shown in figures 3.8 and 3.9. A bias exists in the p_y momentum reconstruction of ~ 1 MeV/c in TKU, with a smaller bias evident in p_x . Due to relative changes in phase-space population densities between beams with different absorber modules, a slight discrepancy in bias is expected and observed in the downstream transverse momentum reconstruction residuals, seen in figures 3.10 and 3.11. In each case, the mean bias for a given beam remains consistent to within 1 MeV/c with the averaged residual across all beams. A similar bias is evident in the longitudinal momentum reconstruction, seen in figures 3.12 and 3.12. Both TKU and TKD observe a mean bias of $\lesssim 1$ MeV/c, though with a much larger spread.

These reconstruction effects introduce a small measurement error which can impact the transverse emittance analysis in chapter 4 and motivates the corrections introduced there.




Figure 3.4: Reconstruction residuals (measured - true) for TKU S1 track-point x positions in mm. Each beam has an RMS of 0.4 mm.

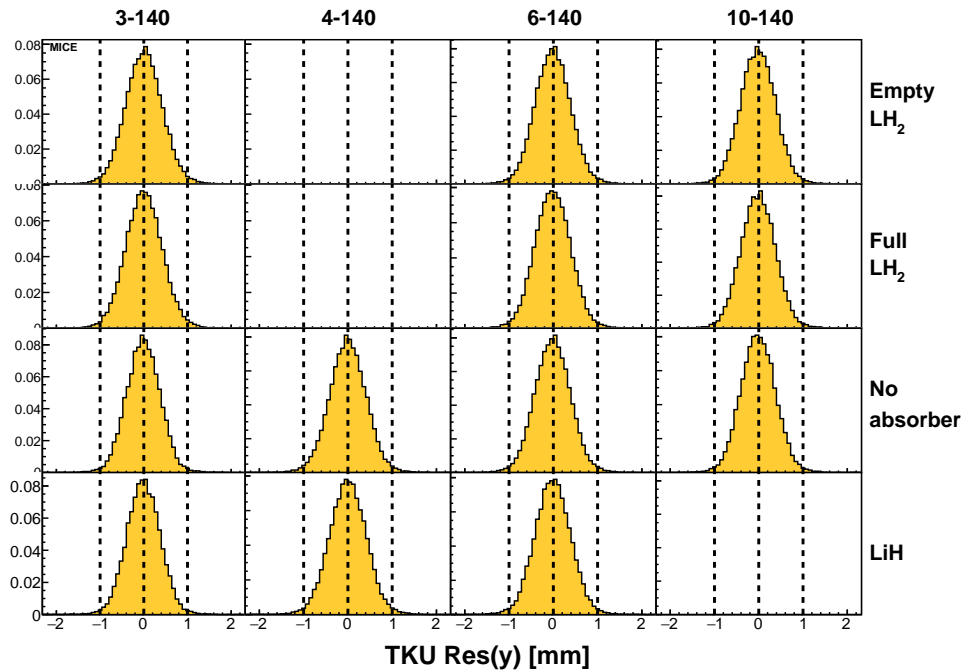


Figure 3.5: Reconstruction residuals (measured - true) for TKU S1 track-point y positions in mm. Each beam has an RMS of 0.4 mm.




Figure 3.6: Reconstruction residuals (measured - true) for TKD S1 track-point x positions in mm. Each beam has an RMS of 0.4 mm.




Figure 3.7: Reconstruction residuals (measured - true) for TKD S1 track-point y positions in mm. Each beam has an RMS of 0.4 mm.




Figure 3.8: Reconstruction residuals (measured - true) for TKU S1 track-point p_x momentum in MeV/c. Each beam has an approximate RMS of 1.3 MeV/c.




Figure 3.9: Reconstruction residuals (measured - true) for TKU S1 track-point p_y momentum in MeV/c. Each beam has an approximate RMS of 1.3 MeV/c.




Figure 3.10: Reconstruction residuals (measured - true) for TKD S1 track-point p_x momentum in MeV/c. Each beam has an approximate RMS of 1.5 MeV/c.




Figure 3.11: Reconstruction residuals (measured - true) for TKD S1 track-point p_y momentum in MeV/c. Each beam has an approximate RMS of 1.5 MeV/c.




Figure 3.12: Reconstruction residuals (measured - true) for TKU S1 track-point p_z momentum in MeV/c. Each beam has an approximate RMS of 2.3 MeV/c.

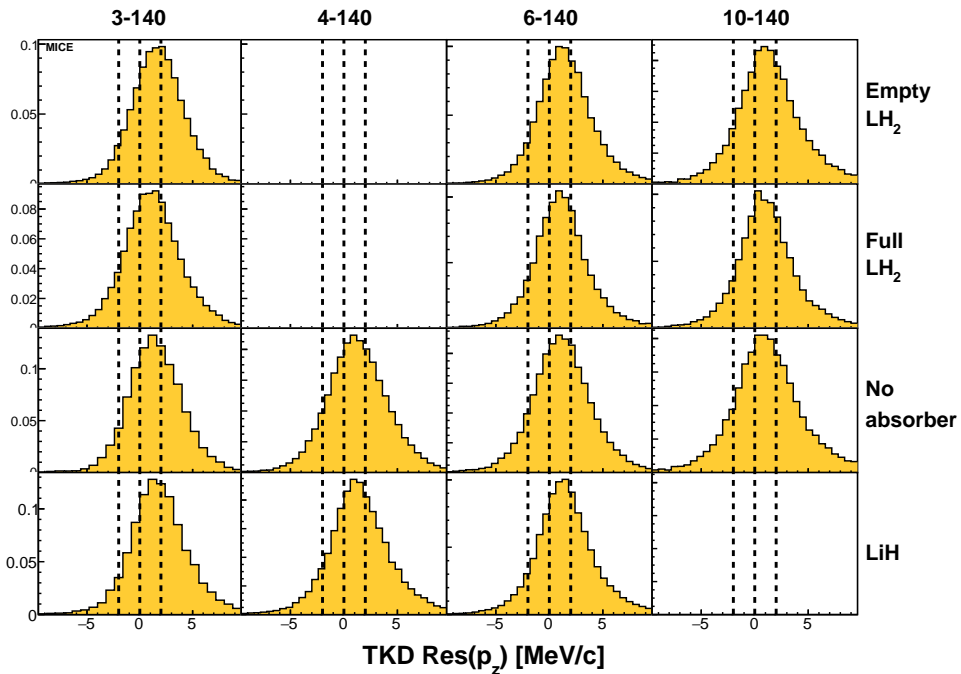


Figure 3.13: Reconstruction residuals (measured - true) for TKD S1 track-point p_z momentum in MeV/c. Each beam has an approximate RMS of 3.2 MeV/c.

3.3 Global Reconstruction

Global track reconstruction combines the various detector outputs to construct a track which traverses the MICE detector suite. Unlike individual detector tracks, reconstructed global track positions in space are given using the global coordinate system which spans the experiment, beginning at D2 and extending past the EMR. The global reconstruction consists of two steps: track matching and particle identification. A third step to perform a global track fitting based on the method utilised in the tracker reconstruction's Kalman filter was under development, but has yet to be implemented in MAUS.

3.3.1 Track Matching

The track matching stage determines whether hits in different detectors are caused by the same particle, propagating the reconstructed particle between detectors and looking for agreement in the position or timing. Due to the precision and accuracy in position-momentum phase-space provided by the tracker reconstruction, global track propagation is performed from the outermost tracker plane, Stations 5 in both TKU and TKD, towards the detector on either side. This procedure provides two separate tracks, upstream and downstream, which are then matched via a time-of-flight cut. The matching is performed for each feasible particle hypothesis as the particle ID has yet to be determined, with the correct track chosen through a separate PID algorithm in the next stage. If no track has been reconstructed in one or both of the trackers, no matching is performed.

The propagation between detectors uses a 4th order Runge-Kutta numerical integration method [47] to calculate and update particle properties at each step following the equations of motion

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

where q and c are the particle charge and speed of light respectively, using the electric and magnetic fields, E and B , provided by the geometry. Additionally, energy loss from material interactions is calculated from the Bethe-Bloch

equation [48]

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

where $K = 4\pi N_A r_e^2 m_e c^2 = 0.307075 \text{ MeV mol}^{-1} \text{ cm}^2$, N_A is Avagadro's constant, m_e and r_e are the electron mass and radius, the kinematic variables β and γ have their usual meanings, Z and A are the atomic and mass numbers of the interacting material, I is the mean excitation energy, $\delta(\beta\gamma)$ is the material density effect correction as a function of $\beta\gamma$, and W_{max} is the maximum energy transfer imparted to a free electron in a single collision, defined as

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

As the propagation routine steps through the beamline, the midpoint between the current and previous position is calculated and the material at that midpoint is obtained from the geometry, providing the material properties for the Bethe-Bloch energy loss. Where the distance to the nearest material boundary is lower than the standard step size, the step size is reduced significantly to minimise inaccuracies from underestimated straight-line step distance in high stopping power materials and to ensure the material during a step is uniform.

Track matching to TOFs 1 and 2 involves propagating to the central z position of either TOF and requiring agreement in the x and y coordinates. Requirements on δt between the tracker and TOFs is not implemented due to the trackers' insufficient timing resolution. Once a TOF1 match is established, an approximate expected travel time between TOFs 0 and 1 is calculated and used to perform matching to TOF0. An x - y coordinate matching is not attempted due to the large uncertainty resulting from the much larger propagation distance to TOF0. For the KL and EMR, y and x - y coordinate matching are used respectively as for TOFs 1 and 2. Cherenkov triggers are added to the tracks without checks due to their insufficient timing resolution for multiple track separation. Finally, track matching through the absorber requires both TOF1 and TOF2 hits and utilises a loose subluminal time-of-flight cut so as to ensure tracks have a realistic travel time without making assumptions on their interactions in the absorber.

3.3.2 Particle Identification

Particle identification for global tracks uses a combined set of detector output variables, producing confidence levels for each PID case given the measured variables. The PID variables used include the time-of-flight between TOF detectors,

the reconstructed momentum in TKU/TKD, and additional information from the Ckovs, KL and EMR detectors when present.

Each confidence level is assigned by comparison with probability density functions of the PID variables produced through Monte-Carlo simulation of the beam. The confidence level assigned to a particular PID hypothesis X is based on the calculation of log-likelihood values given by $\exp LL_X / \sum_i \exp LL_i$ for all i PID hypotheses. The PID hypothesis with the best confidence level is assigned to the track, with a successful identification requiring the PID hypothesis of the propagated global track to match the PID assigned by the identification routine. Assignment can also be configured to require the calculated confidence level exceed a configurable threshold.

3.4 Conclusion

The spectrometer solenoid tracker is the central detector within MICE, providing the most precise phase-space measurements from which global tracks are propagated through the rest of the beamline. The track reconstruction provides particle positions at each station to within ~ 1 mm and momentum information accurate at the 1 MeV/c level, given an accurate geometry model. Reconstruction biases exist in each phase-space variable for each station which motivate the corrections implemented in chapter 4.

Chapter 4

Emittance and Amplitude Analysis

MICE seeks to understand and characterise the cooling effect achieved from focusing muons onto an energy absorber using the high acceptance solenoid assembly which comprises the cooling channel. To do so, a number of differing nominal emittance-in beam samples which have passed through differing absorbers are defined here and their cooling performance is characterised by calculating particle amplitudes.

4.1 Beam Configurations

The beams analysed here will be referred to by their transverse normalised emittance, ε_{\perp}^n , and by their total momentum, p .

MICE took data over a number of beam configurations, tuned for one of several momentum bites at the upstream tracker reference plane and one of several input beam emittance values, shown in table 4.1

Each of the beam configurations was produced by double dipole momentum selection of the beam and selection of beam emittance with the diffuser (Sections 2.2.4 and 2.2.6).

Accounting for the two possible charges of the beam, there are then 32 possible beam configurations. As the 4 mm emittance was a late addition to

		Central Momentum (MeV/c)			
		140	170	200	240
Initial Emittance (mm)	3	3-140	3-170	3-200	3-240
	4	4-140	4-170	4-200	4-240
	6	6-140	6-170	6-200	6-240
	10	10-140	10-170	10-200	10-240

Table 4.1: Table of MICE beam configurations.

the MICE program, several of the earlier absorbers are without data in this configuration. This analysis considers muon beams selected from the 3-140, 4-140, 6-140, and 10-140 configurations.

The current settings for the magnet coils within the cooling channel were maintained throughout the collection of this data to provide the expected uniform solenoid field, with this cooling channel configuration denoted by the cooling channel tag ‘2017-02-6’.

4.2 Data Selection

The MICE Muon Beamline pre-selects incoming particles by momentum by varying the field in the pair of dipoles D1 and D2, with higher magnetic field selecting higher particle momentum. This produces a refined particle sample for event selection, allowing significant particle ID separation and sample selection to be made using the detector systems.

The event selection for the analysis produces two separate samples: the upstream sample whose selection is based entirely on detector measurements made upstream of the absorber module, and the downstream sample, selected as a subset of the upstream sample based on additional criteria in the downstream detector system. The number of events passing the selection criteria for each upstream sample are listed in tables 4.3 and 4.4, with cuts applied successively proceeding down the table. Events passing the downstream sample are listed in the same manner in tables 4.5 and 4.6. Simulated MC event samples are listed in tables 4.7, 4.8, 4.9, and 4.10.

A comparison of the events removed by this selection procedure for real data and MC events is shown in section 4.2.4.

4.2.1 Upstream Sample

The selection requirements of the upstream sample are defined as follows:

- One TOF1 Space point: Event has one, and only one, space point reconstructed in TOF1.
- One TOF0 Space point: Event has one, and only one, space point reconstructed in TOF0.
- One TKU track: Event has one, and only one, track reconstructed in TKU.
- TKU $\chi^2/N_{D.O.F.}$: Reconstructed track in TKU has a χ^2 per degree of freedom less than 8. Tracks with $\chi^2/N_{D.O.F.}$ above 8 are considered to have poor reconstruction quality and so may not reliably represent the beam optics.
- TKU Fiducial Volume: The reconstructed TKU track has a maximum radial displacement from the tracker centre less than 150 mm, calculated by projecting a helical trajectory between tracker stations. Tracks which exceed a 150 mm radius excursion may interact with the inner surface of the solenoid bore and lose energy or scatter, misrepresenting the cooling effect.
- TOF01: Event has a time-of-flight between TOF0 and TOF1 consistent with a muon within the dipole-selected momentum range. As increasing the emittance of a beam using the diffuser also increases momentum loss in the beam through the diffuser, higher emittance muon samples are dipole selected at a larger momentum so as to provide a final muon sample peaked around the same momentum value. As this results in varying velocities of particles upstream of the diffuser, time-of-flight selections are different for different beams.
- TKU Momentum: Event has a track in TKU with reconstructed momentum within 5 MeV/c of the designated beam momentum.
- TKU Momentum vs TOF01: Event has reconstructed time-of-flight consistent with the reconstructed momentum in TKU, within a given reconstruction error, accounting for propagated energy loss and the systematic offset between TOF and tracker reconstruction. Momentum from TOF01 is calculated as $\text{Mass} \times \beta\gamma$.
- Diffuser Aperture: The TKU track, extrapolated upstream to the diffuser, is required to have a radius at both upstream and downstream faces of the diffuser of less than 90 mm. Tracks which cross the diffuser at radii

greater than 90 mm may interact with the diffuser aperture, making them susceptible to poorly understood energy loss or scattering.

Events that fail to meet one or more of these requirements are rejected from the sample and are not included in the analysis. The TOF01 and TKU Momentum vs TOF01 selection criteria for each beam configuration are shown in table 4.2. TOF01 measurements are offset by the electron peak time-of-flight.

	Beam Configuration			
	3-140	4-140	6-140	10-140
TOF01 (ns)	1.5 to 6.0	1.5 to 6.0	1.5 to 5.5	1.5 to 4.5
TKU Momentum vs TOF01 (MeV/c)	26 ± 15	32 ± 15	35 ± 15	70 ± 25

Table 4.2: Upstream sample TOF01 and TKU Momentum vs TOF01 selection criteria for each beam configuration. The TKU Momentum vs TOF01 requirement is calculated as TOF01 momentum minus TKU momentum.

4.2.2 Downstream Sample

The selection requirements of the downstream sample are defined as follows:

- In Upstream Sample: Event is included in the upstream sample.
- One TKD track: Event has one, and only one, track reconstructed in TKD.
- TKD $\chi^2/N_{D.O.F}$: The reconstructed track in TKD has a χ^2 per degree of freedom less than 8.
- TKD Fiducial Volume: The reconstructed track in TKD has a maximum radial displacement from the tracker centre less than 150 mm, calculated by projecting a helical trajectory between tracker stations.
- TKD Momentum: The TKD track has reconstructed momentum of the designated beam momentum upstream $+30$ MeV/c, -50 MeV/c. The asymmetry in downstream momentum accommodates for the expected energy loss through the absorber.

Events that fail to meet one or more of these requirements are rejected from the downstream sample and are not considered in the downstream analysis.

	Upstream Sample Selection, Reconstructed Data					
	2017-02-6 3-140 LH ₂ Empty	2017-02-6 3-140 LH ₂	2017-02-6 3-140 No Absorber	2017-02-6 3-140 LiH	2017-02-6 4-140 No Absorber	2017-02-6 4-140 LiH
All events	1.02×10^6	1.87×10^6	8.21×10^5	4.98×10^5	7.65×10^5	9.26×10^5
One TOF1 Space point	9.83×10^5	1.80×10^6	7.82×10^5	4.75×10^5	7.32×10^5	8.89×10^5
One TOF0 Space point	7.73×10^5	1.46×10^6	6.05×10^5	3.72×10^5	5.54×10^5	6.89×10^5
One TKU track	5.45×10^5	1.03×10^6	2.32×10^5	1.43×10^5	3.80×10^5	4.76×10^5
TKU $\chi^2/N_{D.O.F.}$	5.18×10^5	9.79×10^5	2.16×10^5	1.33×10^5	3.51×10^5	4.41×10^5
TKU Fiducial Volume	5.17×10^5	9.75×10^5	2.14×10^5	1.32×10^5	3.49×10^5	4.38×10^5
TOF01	2.90×10^5	5.25×10^5	1.19×10^5	7.47×10^4	2.09×10^5	2.64×10^5
TKU Momentum	1.14×10^5	2.09×10^5	4.14×10^4	2.63×10^4	9.29×10^4	1.18×10^5
TKU Momentum vs TOF01	1.13×10^5	2.08×10^5	4.12×10^4	2.61×10^4	9.26×10^4	1.17×10^5
Diffuser Aperture	1.03×10^5	1.90×10^5	3.54×10^4	2.24×10^4	8.51×10^4	1.08×10^5
Upstream Sample	1.03×10^5	1.90×10^5	3.54×10^4	2.24×10^4	8.51×10^4	1.08×10^5

Table 4.3: Upstream sample selection for number of reconstructed events, shown to 3 significant figures. Shown for 2017-02-6 3-140 and 4-140 samples.

	Upstream Sample Selection, Reconstructed Data						
	2017-02-6 6-140 LH ₂ Empty	2017-02-6 6-140 LH ₂	2017-02-6 6-140 No Absorber	2017-02-6 6-140 LiH	2017-02-6 10-140 LH ₂ Empty	2017-02-6 10-140 LH ₂	2017-02-6 10-140 No Absorber
All events	1.06×10^6	1.82×10^6	7.46×10^5	6.27×10^5	2.00×10^6	3.67×10^6	1.41×10^6
One TOF1 Space point	1.01×10^6	1.75×10^6	7.12×10^5	5.99×10^5	1.85×10^6	3.43×10^6	1.30×10^6
One TOF0 Space point	7.85×10^5	1.40×10^6	5.34×10^5	4.61×10^5	1.36×10^6	2.66×10^6	9.37×10^5
One TKU track	5.37×10^5	9.51×10^5	3.65×10^5	3.17×10^5	7.19×10^5	1.40×10^6	4.95×10^5
TKU $\chi^2/N_{D.O.F.}$	4.93×10^5	8.60×10^5	3.36×10^5	2.91×10^5	6.64×10^5	1.28×10^6	4.58×10^5
TKU Fiducial Volume	4.89×10^5	8.53×10^5	3.33×10^5	2.89×10^5	6.30×10^5	1.21×10^6	4.33×10^5
TOF01	3.04×10^5	5.32×10^5	2.05×10^5	1.80×10^5	4.03×10^5	8.09×10^5	2.76×10^5
TKU Momentum	1.24×10^5	2.20×10^5	8.43×10^4	7.37×10^4	1.22×10^5	2.46×10^5	8.40×10^4
TKU Momentum vs TOF01	1.24×10^5	2.19×10^5	8.39×10^4	7.33×10^4	1.21×10^5	2.44×10^5	8.31×10^4
Diffuser Aperture	1.14×10^5	2.00×10^5	7.63×10^4	6.66×10^4	8.30×10^4	1.66×10^5	5.69×10^4
Upstream Sample	1.14×10^5	2.00×10^5	7.63×10^4	6.66×10^4	8.30×10^4	1.66×10^5	5.69×10^4

Table 4.4: Upstream sample selection for number of reconstructed events, shown to 3 significant figures. Shown for 2017-02-6 6-140 and 10-140 samples.

	Downstream Sample Selection, Reconstructed Data					
	2017-02-6 3-140 LH ₂ Empty	2017-02-6 3-140 LH ₂	2017-02-6 3-140 No Absorber	2017-02-6 3-140 LiH	2017-02-6 4-140 No Absorber	2017-02-6 4-140 LiH
In Upstream Sample	1.03×10^5	1.90×10^5	3.54×10^4	2.24×10^4	8.51×10^4	1.08×10^5
One TKD track	1.01×10^5	1.85×10^5	3.48×10^4	2.19×10^4	8.28×10^4	1.04×10^5
TKD $\chi^2/N_{D.O.F}$	9.86×10^4	1.81×10^5	3.41×10^4	2.15×10^4	8.12×10^4	1.02×10^5
TKD Fiducial Volume	9.82×10^4	1.79×10^5	3.40×10^4	2.10×10^4	7.97×10^4	9.84×10^4
TKD Momentum	9.82×10^4	1.79×10^5	3.40×10^4	2.10×10^4	7.97×10^4	9.83×10^4
Downstream Sample	9.82×10^4	1.79×10^5	3.40×10^4	2.10×10^4	7.97×10^4	9.83×10^4

Table 4.5: Downstream sample selection for number of reconstructed events, shown to 3 significant figures. Shown for 2017-02-6 3-140 and 4-140 samples.

	Downstream Sample Selection, Reconstructed Data						
	2017-02-6 6-140 LH ₂ Empty	2017-02-6 6-140 LH ₂	2017-02-6 6-140 No Absorber	2017-02-6 6-140 LiH	2017-02-6 10-140 LH ₂ Empty	2017-02-6 10-140 LH ₂	2017-02-6 10-140 No Absorber
In Upstream Sample	1.14×10^5	2.00×10^5	7.63×10^4	6.66×10^4	8.30×10^4	1.66×10^5	5.69×10^4
One TKD track	1.05×10^5	1.85×10^5	7.20×10^4	6.21×10^4	6.44×10^4	1.28×10^5	4.57×10^4
TKD $\chi^2/N_{D.O.F}$	1.03×10^5	1.82×10^5	7.06×10^4	6.11×10^4	6.32×10^4	1.26×10^5	4.48×10^4
TKD Fiducial Volume	9.83×10^4	1.74×10^5	6.78×10^4	5.82×10^4	5.48×10^4	1.12×10^5	3.92×10^4
TKD Momentum	9.82×10^4	1.74×10^5	6.77×10^4	5.82×10^4	5.47×10^4	1.12×10^5	3.91×10^4
Downstream Sample	9.82×10^4	1.74×10^5	6.77×10^4	5.82×10^4	5.47×10^4	1.12×10^5	3.91×10^4

Table 4.6: Downstream sample selection for number of reconstructed events, shown to 3 significant figures. Shown for 2017-02-6 6-140 and 10-140 samples.

4.2.3 Simulation

Events are simulated using the custom-built MAUS software. The upstream section of beamline up to D2 is simulated using G4Beamline, with the resulting particle events handed over to MAUS. The same sample selection is applied to simulated events to retrieve upstream and downstream MC samples, shown in tables 4.7, 4.8, 4.9 and 4.10.

Magnet Tuning

A discrepancy in the incoming momenta of the simulated beams has been eliminated by artificially tuning the simulated dipole fields. Where a discrepancy between data and MC was previously present, the simulated dipole field has been modified with a ~ 1 to 2% larger field than suggested by measurements obtained from the hall probe dipole survey. Figure 4.1 (left) shows the position and momentum distributions of an upstream sample taken from a nominal simulated 10-140 beam prior to this tuning process, with the resulting distributions after an applied 2.12% D2 field correction shown in figure 4.1 (right). In both cases, the event sampling uses the selection described in section 4.2.1, showing all events which pass all selection criteria except for the requirements on reconstructed TKU momentum. The tuning process ensures the upstream sample muon ensembles for reconstructed data and simulated events occupy similar proportions of phase-space and are therefore meaningfully comparable.




Figure 4.1: Reconstructed momentum distributions at TKU Station 1 for an upstream muon sample with all selection requirements other than the momentum selection requirement applied. Plots momentum distributions (left) before and (right) after dipole tuning, with real data events in black and MC events in yellow.

	Upstream Sample Selection, Reconstructed Simulation					
	2017-02-6 3-140 LH ₂ Empty	2017-02-6 3-140 LH ₂	2017-02-6 3-140 No Absorber	2017-02-6 3-140 LiH	2017-02-6 4-140 No Absorber	2017-02-6 4-140 LiH
All events	4.98×10^5	4.87×10^5	4.03×10^5	4.01×10^5	5.41×10^5	5.55×10^5
One TOF1 Space point	4.44×10^5	4.34×10^5	3.53×10^5	3.51×10^5	4.83×10^5	4.96×10^5
One TOF0 Space point	4.13×10^5	4.03×10^5	3.26×10^5	3.25×10^5	4.49×10^5	4.61×10^5
One TKU track	2.90×10^5	2.83×10^5	1.35×10^5	1.35×10^5	3.03×10^5	3.11×10^5
TKU $\chi^2/N_{D.O.F.}$	2.76×10^5	2.69×10^5	1.26×10^5	1.26×10^5	2.80×10^5	2.87×10^5
TKU Fiducial Volume	2.74×10^5	2.67×10^5	1.22×10^5	1.22×10^5	2.77×10^5	2.84×10^5
TOF01	1.28×10^5	1.25×10^5	6.04×10^4	6.04×10^4	1.36×10^5	1.40×10^5
TKU Momentum	4.76×10^4	4.64×10^4	2.03×10^4	2.03×10^4	5.58×10^4	5.74×10^4
TKU Momentum vs TOF01	4.73×10^4	4.61×10^4	2.02×10^4	2.02×10^4	5.56×10^4	5.72×10^4
Diffuser Aperture	4.39×10^4	4.27×10^4	1.75×10^4	1.75×10^4	5.06×10^4	5.19×10^4
Upstream Sample	4.39×10^4	4.27×10^4	1.75×10^4	1.75×10^4	5.06×10^4	5.19×10^4

Table 4.7: Upstream sample selection for number of reconstructed MC events, shown to 3 significant figures. Shown for 2017-02-6 3-140 and 4-140 samples.

	Upstream Sample Selection, Reconstructed Simulation						
	2017-02-6 6-140 LH ₂ Empty	2017-02-6 6-140 LH ₂	2017-02-6 6-140 No Absorber	2017-02-6 6-140 LiH	2017-02-6 10-140 LH ₂ Empty	2017-02-6 10-140 LH ₂	2017-02-6 10-140 No Absorber
All events	5.99×10^5	5.93×10^5	5.99×10^5	5.99×10^5	1.08×10^6	1.06×10^6	1.07×10^6
One TOF1 Space point	5.35×10^5	5.30×10^5	5.35×10^5	5.35×10^5	9.59×10^5	9.47×10^5	9.55×10^5
One TOF0 Space point	4.99×10^5	4.94×10^5	4.99×10^5	4.98×10^5	8.94×10^5	8.83×10^5	8.90×10^5
One TKU track	3.28×10^5	3.25×10^5	3.29×10^5	3.29×10^5	4.63×10^5	4.57×10^5	4.61×10^5
TKU $\chi^2/N_{D.O.F.}$	3.00×10^5	2.96×10^5	3.01×10^5	3.00×10^5	4.29×10^5	4.23×10^5	4.27×10^5
TKU Fiducial Volume	2.95×10^5	2.92×10^5	2.96×10^5	2.96×10^5	3.87×10^5	3.82×10^5	3.86×10^5
TOF01	1.49×10^5	1.48×10^5	1.49×10^5	1.50×10^5	2.19×10^5	2.17×10^5	2.19×10^5
TKU Momentum	5.62×10^4	5.55×10^4	5.60×10^4	5.60×10^4	5.32×10^4	5.24×10^4	5.30×10^4
TKU Momentum vs TOF01	5.57×10^4	5.50×10^4	5.55×10^4	5.55×10^4	5.13×10^4	5.07×10^4	5.13×10^4
Diffuser Aperture	5.02×10^4	4.95×10^4	5.00×10^4	5.00×10^4	3.55×10^4	3.49×10^4	3.54×10^4
Upstream Sample	5.02×10^4	4.95×10^4	5.00×10^4	5.00×10^4	3.55×10^4	3.49×10^4	3.54×10^4

Table 4.8: Upstream sample selection for number of reconstructed MC events, shown to 3 significant figures. Shown for 2017-02-6 6-140 and 10-140 samples.

	Downstream Sample Selection, Reconstructed Simulation					
	2017-02-6 3-140 LH ₂ Empty	2017-02-6 3-140 LH ₂	2017-02-6 3-140 No Absorber	2017-02-6 3-140 LiH	2017-02-6 4-140 No Absorber	2017-02-6 4-140 LiH
In Upstream Sample	4.39×10^4	4.27×10^4	1.75×10^4	1.75×10^4	5.06×10^4	5.19×10^4
One TKD track	4.28×10^4	4.16×10^4	1.71×10^4	1.69×10^4	4.88×10^4	4.94×10^4
TKD $\chi^2/N_{D.O.F}$	4.19×10^4	4.05×10^4	1.65×10^4	1.64×10^4	4.75×10^4	4.82×10^4
TKD Fiducial Volume	4.18×10^4	4.02×10^4	1.64×10^4	1.62×10^4	4.67×10^4	4.70×10^4
TKD Momentum	4.18×10^4	4.02×10^4	1.64×10^4	1.62×10^4	4.67×10^4	4.69×10^4
Downstream Sample	4.18×10^4	4.02×10^4	1.64×10^4	1.62×10^4	4.67×10^4	4.69×10^4

Table 4.9: Downstream sample selection for number of reconstructed MC events, shown to 3 significant figures. Shown for 2017-02-6 3-140 and 4-140 samples.

	Downstream Sample Selection, Reconstructed Simulation						
	2017-02-6 6-140 LH ₂ Empty	2017-02-6 6-140 LH ₂	2017-02-6 6-140 No Absorber	2017-02-6 6-140 LiH	2017-02-6 10-140 LH ₂ Empty	2017-02-6 10-140 LH ₂	2017-02-6 10-140 No Absorber
In Upstream Sample	5.02×10^4	4.95×10^4	5.00×10^4	5.00×10^4	3.55×10^4	3.49×10^4	3.54×10^4
One TKD track	4.55×10^4	4.47×10^4	4.57×10^4	4.51×10^4	2.67×10^4	2.62×10^4	2.70×10^4
TKD $\chi^2/N_{D.O.F}$	4.46×10^4	4.40×10^4	4.44×10^4	4.42×10^4	2.62×10^4	2.58×10^4	2.63×10^4
TKD Fiducial Volume	4.19×10^4	4.18×10^4	4.21×10^4	4.20×10^4	2.29×10^4	2.34×10^4	2.33×10^4
TKD Momentum	4.19×10^4	4.18×10^4	4.21×10^4	4.20×10^4	2.28×10^4	2.32×10^4	2.32×10^4
Downstream Sample	4.19×10^4	4.18×10^4	4.21×10^4	4.20×10^4	2.28×10^4	2.32×10^4	2.32×10^4

Table 4.10: Downstream sample selection for number of reconstructed MC events, shown to 3 significant figures. Shown for 2017-02-6 6-140 and 10-140 samples.

4.2.4 Real versus Simulated Sample Selection

To assess the sample selection procedure and compare the simulated event sample against the real data sample, distributions of upstream and downstream samples are plotted showing only events which pass all selection requirements for the sample excluding the requirement on the variable shown in each figure. Each figure is composed of one distribution from each data sample, with rows of the same absorber type and columns of the same nominal upstream emittance. As the TKU momentum vs TOF01 selection effectively acts as a pre-selection to ensure the particle is a muon and is largely replicated by the individual TOF01 and TKU momentum selection requirements, the effect this has on all events which have both a single TOF01 measurement and a single reconstructed TKU track is instead shown in figure 4.2. The time-of-flight of each event is shown offset by the electron time-of-flight, hence the electron population is evident at 0 ns. The muon events occupy the central portion of each distribution, whereas the pion population is evidently reconstructed at later time-of-flight. The combined effect of the individual TKU momentum and TOF01 requirements on events within the 2D phase-space are shown in red.

Figures 4.3 to 4.11 show the effect of each remaining selection requirement on the upstream samples, whilst figures 4.12 to 4.14 show those on the downstream samples. Due to a discrepancy between data and MC in the relative particle populations which reach the tracker, the relative number of space points reconstructed in either TOF differs between data and MC. As the sample impurity is very low, this does not affect the analysis.





Figure 4.2: Reconstructed TKU momentum versus TOF01 time-of-flight for all events in (top) data and (bottom) MC which reconstruct a single track in TKU and a single spacepoint in TOFs 1 and 2. The upper and lower bounds of the TKU momentum vs TOF01 selection is shown by the blue curves in each distribution. The TKU momentum and TOF01 requirements are shown in red.




Figure 4.3: Upstream sample selection requiring 1 reconstructed spacepoint in TOF1, showing real data events in black and MC events in yellow.




Figure 4.4: Upstream sample selection requiring 1 reconstructed spacepoint in TOF0, showing real data events in black and MC events in yellow.




Figure 4.5: Upstream sample selection requiring 1 reconstructed track in TKU, showing real data events in black and MC events in yellow.




Figure 4.6: Upstream sample selection requiring a $\chi^2 > 8$ for reconstructed TKU tracks, showing real data events in black and MC events in yellow.




Figure 4.7: Upstream sample selection requiring max fiducial radius less than 150 mm for reconstructed TKU tracks, showing real data events in black and MC events in yellow.




Figure 4.8: Upstream sample selection requiring a TOF01 time-of-flight consistent with the muon hypothesis, showing real data events in black and MC events in yellow. Dotted lines are included to guide the eye.




Figure 4.9: Upstream sample selection requiring a reconstructed momentum of 140 ± 5 for tracks in TKU, showing real data events in black and MC events in yellow.




Figure 4.10: Upstream sample selection requiring a global track propagated upstream which passes through the US virtual diffuser face at radius < 90 mm, showing real data events in black and MC events in yellow.




Figure 4.11: Upstream sample selection requiring a global track propagated upstream which passes through the DS virtual diffuser face at radius < 90 mm, showing real data events in black and MC events in yellow.




Figure 4.12: Downstream sample selection requiring 1 reconstructed track in TKD, showing real data events in black and MC events in yellow.




Figure 4.13: Downstream sample selection requiring a $\chi^2 > 8$ for reconstructed TKD tracks, showing real data events in black and MC events in yellow.




Figure 4.14: Downstream sample selection requiring max fiducial radius less than 150 mm for reconstructed TKD tracks, showing real data events in black and MC events in yellow.




Figure 4.15: Downstream sample selection requiring a reconstructed momentum of $140 + 30 \text{ MeV}/c$, $-50 \text{ MeV}/c$ for tracks in TKD, showing real data events in black and MC events in yellow.

4.2.5 Hybrid Simulation

The full beamline simulation employed for production of the samples described in section 4.2.3 is computationally expensive for large numbers of muon events as this includes particle production and propagation through much of the non-instrumented upstream beamline. To provide significantly larger samples of muons for systematic studies, a hybrid MC simulation has been employed, where the reconstructed upstream real data sample has been used to produce a multivariate Kernel Density Estimation (KDE) distribution of particle properties at TKU Station 5. The resulting distribution is sampled to provide (x, p_x, y, p_y, p_z) muon properties at TKU Station 5, in addition to muon-consistent TOF01 information, with the resulting events propagated through the remainder of the MICE beamline. This procedure can provide a significantly larger muon sample than the original data sample, with simulation times much lower than for the full MC simulation.

Several additional ‘truth’ samples can then be considered alongside the reconstructed upstream and downstream event samples. A ‘truth’ sample has selection requirements based on the true values of particle phase-space variables, rather than their reconstructed values. The number of events passing the selection criteria for each hybrid MC sample are shown in tables 4.11, 4.12, 4.13, and 4.14. These hybrid MC samples are used to characterise reconstruction effects and uncertainties, discussed in greater detail in sections 4.4 and 4.5. Selection criteria upstream of TKU station 5 are not applied as particles are considered to have met these criteria by virtue of being included in the KDE data sample.

The upstream truth sample is constructed by requiring that the event is included in the upstream reconstructed sample, that the MC truth muon passes through the active region of all 5 TKU stations, and that the MC truth muon in TKU has a maximum radial displacement from the tracker centre less than 150 mm, calculated by projecting a helical trajectory between tracker stations using the MC truth kinematic variables.

The downstream truth sample is constructed by requiring that the event is included in the upstream truth sample, that the MC truth muon passes through the active region of all 5 TKD stations, that the MC truth muon in TKD has a maximum radial displacement from the tracker centre less than 150 mm, calculated by projecting a helical trajectory between tracker stations using the MC truth kinematic variables, and that the MC truth muon in TKD has momentum of the designated beam momentum upstream $+30 \text{ MeV}/c, -50 \text{ MeV}/c$.

	Upstream Sample Selection, Hybrid MC			
	2017-02-6 3-140 LH ₂ Empty	2017-02-6 4-140 No Absorber	2017-02-6 6-140 LH ₂ Empty	2017-02-6 10-140 LH ₂ Empty
All Events	9.81×10^5	9.81×10^5	9.81×10^5	9.81×10^5
One TKU track	9.40×10^5	9.66×10^5	9.08×10^5	8.59×10^5
TKU $\chi^2/N_{D.O.F.}$	9.29×10^5	9.54×10^5	8.98×10^5	8.50×10^5
TKU Fiducial Volume	9.29×10^5	9.53×10^5	8.97×10^5	8.43×10^5
TKU Momentum	6.61×10^5	6.87×10^5	5.94×10^5	5.05×10^5
Upstream Sample	6.61×10^5	6.87×10^5	5.94×10^5	5.05×10^5

Table 4.11: Upstream sample selection for number of reconstructed hybrid simulation events, shown to 3 significant figures. Shown for 2017-02-6 3, 6, and 10-140 LH₂ and 4-140 no absorber samples.

	Downstream Sample Selection, Hybrid MC			
	2017-02-6 3-140 LH ₂ Empty	2017-02-6 4-140 No Absorber	2017-02-6 6-140 LH ₂ Empty	2017-02-6 10-140 LH ₂ Empty
In Upstream Sample	6.61×10^5	6.87×10^5	5.94×10^5	5.05×10^5
One TKD track	6.25×10^5	6.51×10^5	5.43×10^5	3.97×10^5
TKD $\chi^2/N_{D.O.F}$	6.12×10^5	6.35×10^5	5.31×10^5	3.88×10^5
TKD Fiducial Volume	6.10×10^5	6.23×10^5	5.07×10^5	3.46×10^5
TKD Momentum	6.10×10^5	6.22×10^5	5.07×10^5	3.45×10^5
Downstream Sample	6.10×10^5	6.22×10^5	5.07×10^5	3.45×10^5

Table 4.12: Downstream sample selection for number of reconstructed hybrid simulation events, shown to 3 significant figures. Shown for 2017-02-6 3, 6, and 10-140 LH₂ and 4-140 no absorber samples.

Upstream MC Truth Sample Selection, Hybrid MC				
	2017-02-6 3-140 LH ₂ Empty	2017-02-6 4-140 No Absorber	2017-02-6 6-140 LH ₂ Empty	2017-02-6 10-140 LH ₂ Empty
In Upstream Sample	6.61×10^5	6.87×10^5	5.94×10^5	5.05×10^5
TKU MC Truth Stations	6.61×10^5	6.87×10^5	5.94×10^5	5.05×10^5
TKU MC Truth Fiducial Volume	6.61×10^5	6.87×10^5	5.94×10^5	5.04×10^5
Upstream MC Truth Sample	6.61×10^5	6.87×10^5	5.94×10^5	5.04×10^5

Table 4.13: Upstream MC truth sample selection for number of reconstructed hybrid simulation events, shown to 3 significant figures. Shown for 2017-02-6 3, 6, and 10-140 LH₂ and 4-140 no absorber samples.

	Downstream MC Truth Sample Selection, Hybrid MC			
	2017-02-6 3-140 LH ₂ Empty	2017-02-6 4-140 No Absorber	2017-02-6 6-140 LH ₂ Empty	2017-02-6 10-140 LH ₂ Empty
In Upstream Truth Sample	6.61×10^5	6.87×10^5	5.94×10^5	5.04×10^5
TKD MC Truth Stations	6.28×10^5	6.58×10^5	5.52×10^5	4.09×10^5
TKD MC Truth Fiducial Volume	6.26×10^5	6.45×10^5	5.24×10^5	3.59×10^5
TKD MC Truth Momentum	6.26×10^5	6.44×10^5	5.23×10^5	3.57×10^5
Downstream MC Truth Sample	6.26×10^5	6.44×10^5	5.23×10^5	3.57×10^5

Table 4.14: Downstream MC truth sample selection for number of reconstructed hybrid simulation events, shown to 3 significant figures. Shown for 2017-02-6 3, 6, and 10-140 LH₂ and 4-140 no absorber samples.

4.3 Calculation of Amplitudes

Cooling and heating in MICE are characterised by measuring the change in particle amplitude. As defined in equation 1.18, the transverse amplitude of a particle in an elliptical beam is calculated as

$$A_{\perp} = \varepsilon_{\perp}(p - \bar{p})^T \Sigma^{-1}(p - \bar{p}), \quad (4.1)$$

where $(p - \bar{p})^T \Sigma^{-1}(p - \bar{p})$ is the squared normalised distance of each particle from the beam centroid in phase-space. This requires the calculation of the beam's 4D covariance matrix, Σ .

To avoid an induced bias from comparison of particles in the beam with the beams own centre, the sample is split into half samples. Instead, each half-sample uses the covariance matrix of the other sample to calculate its own particle amplitudes. This removes the bias from common statistical deviations in the calculated covariance matrix and each particle amplitude. The sample for which particle amplitudes are being calculated is referred to as the test sample, while the sample which provides the covariance matrix is referred to as the reference sample.

In this analysis, particle amplitude is calculated using an iterative process which accounts for the long tails in the amplitude distributions by sequentially removing large amplitude events from the event sample. This iterative algorithm first produces a binned amplitude distribution using the entire reference and test samples, as outlined above. The particle amplitudes of events in the reference sample are also calculated using the reference sample covariance matrix. A largest allowed amplitude bin is defined initially as one less than the largest amplitude bin in the reference sample, with subsequent iterations defined as one less than the previous largest allowed bin. The reference sample covariance matrix and amplitudes are then recalculated after having removed all events in larger bins than the largest allowed amplitude bin from the reference sample. This recalculation and removal alters the beam centre with each iteration and so is repeated until no further events are recalculated with larger than allowed amplitude. The updated covariance matrix is then used to recalculate binned amplitudes for the test sample, with the amplitude of all events in larger bins than the largest allowed amplitude bin stored and those events removed from the test sample, providing the number of events from the test sample which contribute to the smallest disallowed amplitude bin. The recalculation of the reference and test samples and the removal of events from each sample is repeated with the largest allowed amplitude bin reduced at each repetition, until all events are removed from the test sample. The entire algorithm is then repeated, swapping the test and reference samples.

An example subset of the resulting 2D RMS ellipses produced by this iterative covariance matrix recalculation are shown in figure 4.16. The RMS ellipses of each half-sample match very closely at each recalculation stage shown.





Figure 4.16: Example RMS ellipses resulting from iterative recalculation of the covariance matrix. The two half-samples are coloured in shades of red and green, with darker shades as the amplitude cut approaches the distribution centre.


The resulting distributions describe the amplitude particle density (PDFs) and can be combined to form cumulative density functions (CDFs). Example PDF and CDF distributions of the true single particle amplitudes for a simulated beam are shown in figures 4.17 and 4.18. Here, particle amplitudes are calculated from the true phase-space variables of muons at Station 1 in TKU and TKD which satisfy upstream and downstream MC truth selection criteria. For a 4 mm nominal emittance beam passed through a lithium hydride absorber, the equilibrium emittance is close to the initial emittance and so little change between US and DS distributions is expected, as shown in figure 4.17. For a simulated beam with larger initial nominal emittance, an increase in the number of low amplitude muons is expected, as shown for the 6-140 LiH simulated beam in figure 4.18.



(a) Amplitude PDF distribution.

(b) Amplitude CDF distribution.

Figure 4.17: Example MC truth amplitude PDF and CDF distributions for a simulated 4-140 beam passed through a lithium hydride absorber. Normalised bin contents of the upstream and downstream distributions are denoted by red and green markers respectively, with an arbitrary smooth curve added to represent the possible underlying data distribution. PDFs are normalised against the largest US bin, CDFs are normalised against the largest DS bin.



(a) Amplitude PDF distribution.

(b) Amplitude CDF distribution.

Figure 4.18: Example MC truth amplitude PDF and CDF distributions for a simulated 6-140 beam passed through a lithium hydride absorber. Normalised bin contents of the upstream and downstream distributions are denoted by red and green markers respectively, with an arbitrary smooth curve added to represent the possible underlying data distribution. PDFs are normalised against the largest US bin, CDFs are normalised against the largest DS bin.

4.4 Effects of Detector Reconstruction on Amplitude

The main contributions to track reconstruction biases arise from beam impurity, track reconstruction inefficiency, and resolution effects or biases in phase-space variable reconstruction. Beam impurity, where additional non-muon events are included in the upstream or downstream samples, causes an overestimated number of events in an amplitude bin. Conversely, reconstruction inefficiency results from failing to reconstruct tracks in one or more detectors, leading to an underestimation in bin events, and can only contribute to uncertainties in the downstream sample as the cooling measurement is defined with respect to events measured upstream. Combined, impurity and inefficiency can artificially augment or deplete areas of the phase-space. Resolution effects can lead to migration of events from one bin to another by mis-reconstructing their position in phase-space. These biases in the track reconstruction performance can be studied and accounted for in many cases by utilising true variables from simulated MC events. The hybrid simulation procedure and sample selection outlined in section 4.2.5 provide the basis for such corrections.

4.4.1 Track Reconstruction Inefficiency

Inefficiencies occur where the detector fails to properly reconstruct a muon track which should have been included in the event sample. Inefficiency in the trackers can result from extreme scattering in the tracker planes leading to failed reconstruction of the track entirely or misreconstruction, resulting in the improper exclusion of the track from the event sample. Scattering in the tracker stations is dominant over Lorentz force curvature for low p_t tracks, leading to poor p_z reconstruction and hence increased probability of rejection from the event sample. Conversely, high amplitude events are also at larger risk of rejection from the sample due to track deformation in the non-linear non-uniform fields.

To quantify this contribution to the measurement uncertainty, the additional 'truth' samples defined in section 4.2.5 can be considered alongside the reconstructed upstream and downstream event samples. This allows the characterisation of tracker inefficiency by evaluating the ratio of the downstream reconstructed sample over the downstream truth sample as a function of true phase-space variables. These inefficiencies are shown as a function of 2D phase-space in Fig. 4.19.

4.4.2 Beam Impurity

Impurity in the beam can in principle originate from noise in the tracker readout, leading to reconstructed clusters, spacepoints and tracks that do not




Figure 4.19: Reconstruction inefficiencies in the tracker as a function of different 2D phase-space variables for a simulated 10-140 no absorber beam.

originate from a beam particle. Noise contribution to track reconstruction can be estimated with a data-driven measurement and has been shown in previous studies to have little impact on track reconstruction, with the average noise per channel per event evaluated at 0.18% upstream and 0.06% downstream [46]. Impurities can also arise from pions or electrons misidentified as muons within the beam sample, but such events are rare, compatible with zero in other studies [49].

4.4.3 Reconstruction Resolution

Resolution effects and biases in the reconstruction can arise due to mischaracterisation of the solenoid field or a misalignment of the detectors or fields relative to the model of the cooling channel geometry assumed by the event reconstruction. A known bias in the track reconstruction algorithms is also present, as discussed in section 3.2.6, and this motivates the corrections presented in section 4.4.4. This bias corresponds to a mean misreconstruction of p_x and p_y by < 0.5 MeV/c and < 1 MeV/c respectively in TKU, and < 0.5 MeV/c and < 1 MeV/c respectively in TKD. These mean biases are smaller than the reconstruction precision, shown in figures 4.20, 4.21, 4.22 and 4.23, and are small on the scale of the reconstructed distributions, shown comparing MC reco and MC truth distributions in figures 4.24 and 4.25.





Figure 4.20: Reconstruction residuals of x and y in TKU, calculated as reconstructed minus true phase-space variables.





Figure 4.21: Reconstruction residuals of p_x and p_y in TKU, calculated as reconstructed minus true phase-space variables.





Figure 4.22: Reconstruction residuals of x and y in TKD, calculated as reconstructed minus true phase-space variables.


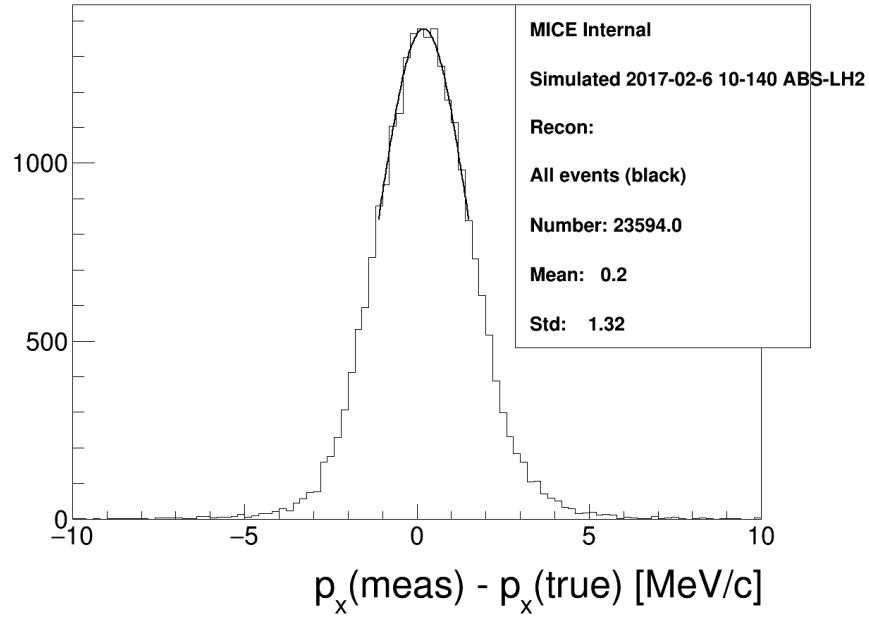


Figure 4.23: Reconstruction residuals of p_x and p_y in TKD, calculated as reconstructed minus true phase-space variables.




Figure 4.24: Comparison of reconstructed and true x , y , p_x , and p_y phase-space variables in TKU station 1. Black points show reconstructed variables for MC events, yellow histogram shows true variables for MC events.




Figure 4.25: Comparison of reconstructed and true x , y , p_x , and p_y phase-space variables in TKD station 1. Black points show reconstructed variables for MC events, yellow histogram shows true variables for MC events.

Where mismodelling of the magnetic field occurs, reconstruction differences between the actual and model fields will still produce good values for a track's χ^2 per degree of freedom, but the reconstructed momenta p_t and p_z scale with the B_z field term, from equations 3.2 and 3.3 used by the helical track reconstruction algorithm. Hence, a well-fitted track may still evince a bias in the reconstructed momentum. As such, a +1% discrepancy in the field corresponds to a +1% error in momentum reconstruction.

Differences between the actual and model fields cannot be corrected for via simulation and so have been validated by direct measurements of the solenoid field. Hall probes stationed at several positions along the cooling channel have been used to record field values over the lifetime of the experiment, with the deviation of each probe's readings from its mean over the data-taking period shown in figure 4.26. These measurements show a maximum deviation of less than 0.02% in SSU and 0.03% in SSD. It's important to note that knowing the field in this way does not solve all biases. If there are position or angular alignment differences between the data and the reconstruction model, biases may still exist in the reconstructed particle properties. This will be considered in the systematic uncertainties evaluated in section 4.5.2.




Figure 4.26: Hall probe readings over time from several hall probes positioned along the MICE cooling channel.

4.4.4 Corrections

Due to the detector reconstruction effects outlined above, reconstructed data measurements will not necessarily match the true amplitude distribution and so the goal of the applied correction is to retrieve the underlying true distribution through an unfolding procedure. By utilising the hybrid MC sample selection described in section 4.2.5, a set of three amplitude distributions are generated for each tracker. Each distribution returns a vector of events per bin describing:

- \mathbf{A}^t , the amplitude distribution of the truth event sample calculated from true phase space variables.
- \mathbf{A}^{rt} , the amplitude distribution of the reconstructed event sample calculated from true phase space variables.
- \mathbf{A}^r , the amplitude distribution of the reconstructed event sample calculated from reconstructed phase space variables.

From these samples, we can calculate an efficiency and purity vector \mathbf{E} , and a normalised migration matrix \mathbf{M} which describes the crossing probability between amplitude bins. We define

$$A_i^t = E_i A_i^{rt} \quad (4.2)$$

and

$$A_i^{rt} = \sum_j M_{ij} A_j^r \quad (4.3)$$

such that

$$E_i = \frac{A_i^t}{A_i^{rt}} \quad (4.4)$$

and

$$A_i^t = E_i \sum_j M_{ij} A_j^r. \quad (4.5)$$

We can define a matrix \mathbf{m} as the un-normalised migration matrix, calculated for each m_{ij} as $A_i^r \cup A_j^{rt}$, the number of events contained in both A_i^r and A_j^{rt} . The normalised migration matrix, \mathbf{M} , then satisfies

$$M_{ij} = \frac{m_{ij}}{A_i^r} \quad (4.6)$$

such that a distribution \mathbf{A}^r with reconstructed events in bin i has an $M_{ii} = 1$ if all of these events remain in the i -th bin in \mathbf{A}^{rt} .

Applying the transformations as described in Eq. 4.5 to the reconstructed amplitude distribution retrieves the true amplitude distribution. The efficiency and purity vector \mathbf{E} and normalised migration matrix \mathbf{M} are calculated from a

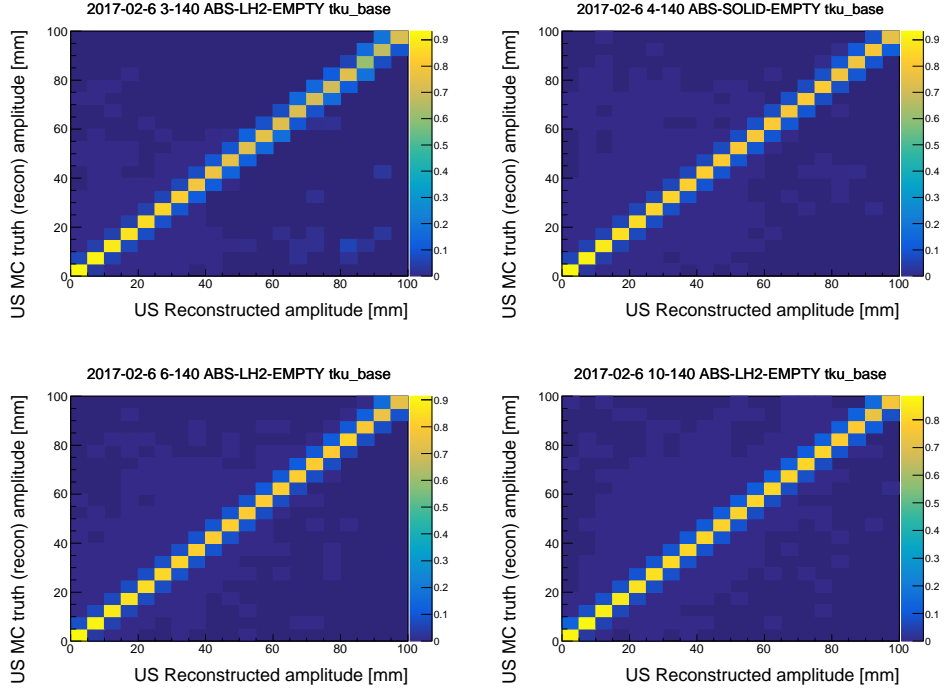


Figure 4.27: US migration matrices for 3, 4, 6, and 10-140 beams describing the probability of TKU reconstruction in each amplitude bin compared to an event’s true amplitude bin.

hybrid MC sample for each nominal emittance beam setting independently and applied to the corresponding reconstructed data and MC sample distributions, providing corrected upstream and downstream amplitude distributions.

The migration matrices for the 3, 4, 6, and 10-140 hybrid MC beams are shown for the upstream event sample in Fig. 4.27 and for the downstream event sample in Fig. 4.28. The probability of an event being reconstructed in the same amplitude bin as its true amplitude bin corresponds to the diagonal elements of \mathbf{M} , and is the most likely case, but in each beam there is a $\gtrsim 10\%$ probability for event migration for all bins and hence a non-trivial correction to be made. This effect is more significant in the downstream sample due to the larger reconstruction biases shown in section 4.4.3. While the probability of an event migrating to a neighbouring amplitude bin is relatively high, migrations through more than one bin occur significantly less frequently indicating significant errors in reconstruction are rare. The increased number of non-diagonal entries present in the 3 mm migration matrix at large amplitudes occurs largely from statistical variations present due to the very low number of events at these amplitudes in the 3 mm beam.

The efficiency correction factors for each downstream MC truth amplitude bin are shown in figure 4.29. As the upstream MC truth sample requires the particle be contained in the upstream reconstructed sample, and as very few

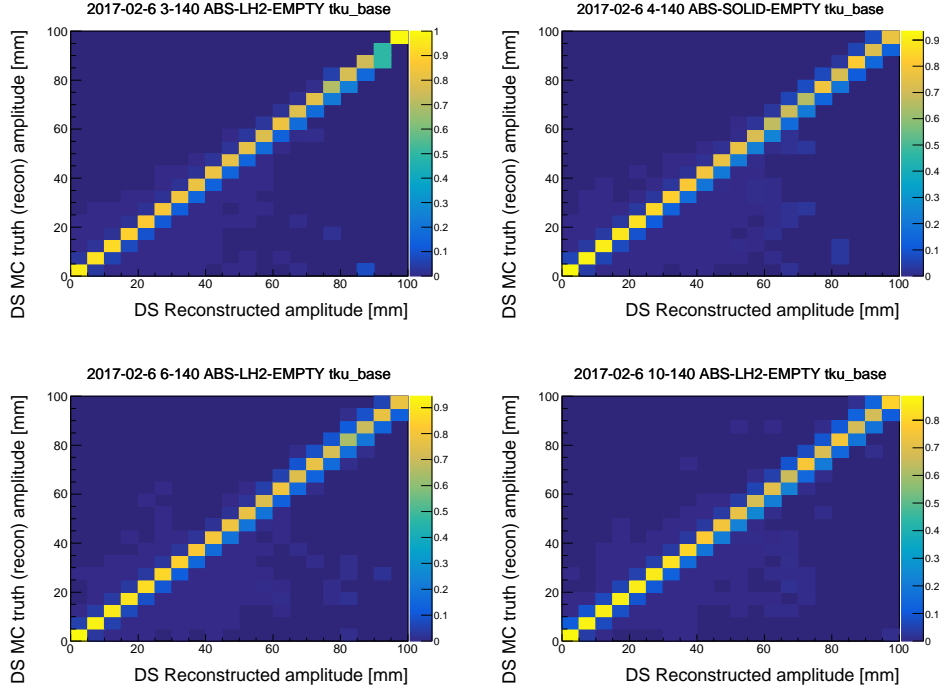




Figure 4.28: DS migration matrices for 3, 4, 6, and 10-140 beams describing the probability of TKD reconstruction in each amplitude bin compared to an event’s true amplitude bin.


events are removed by the upstream MC truth sample selection requirements, the upstream efficiency correction factors are effectively 1 in every case and hence are not shown. The effect of applying the full correction procedure to each sample from which it was calculated is shown in figures 4.30 and 4.31. In each case, the correction appropriately restores the low amplitude bins to match the MC truth amplitude distribution. At large amplitudes where the total number of events goes to zero, the ratio becomes very large, however this accounts for very few events. This is more evident in the downstream corrections as the beams observe significant particle loss at amplitudes greater than 30 mm due to aperture scraping through the cooling channel, resulting in much fewer events in this region. As the efficiency is very close to 1 in the upstream sample, inefficiency contributions to the corrections upstream are insignificant and the PDFs are effectively corrected from the migration matrix only. For the downstream sample, a larger inefficiency correction is required which increases the overall contribution of this correction procedure.




DS MC Truth Amplitude [mm]



DS MC Truth Amplitude [mm]



DS MC Truth Amplitude [mm]



DS MC Truth Amplitude [mm]

Figure 4.29: Inefficiency correction factors for each DS MC truth amplitude bin for (top to bottom) 3, 4, 6, and 10-140 beams.

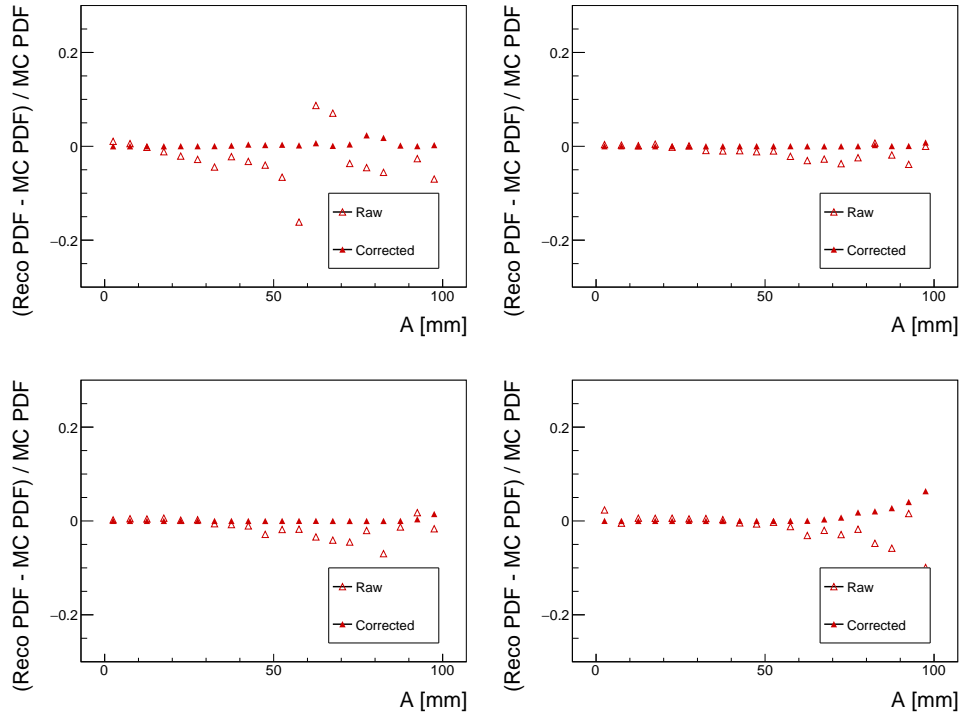


Figure 4.30: Hybrid MC upstream amplitude PDF resolutions for (top left) 3-140, (top right) 4-140, (bottom left) 6-140, and (bottom right) 10-140 beams before and after applying the efficiency and migration correction.

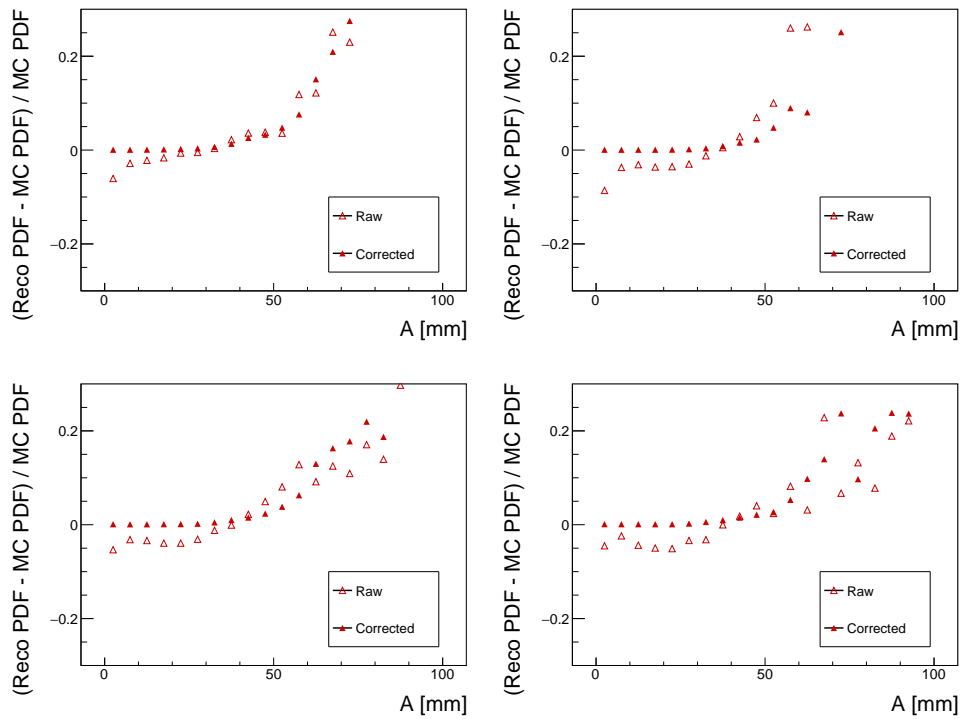


Figure 4.31: Hybrid MC downstream amplitude PDF resolutions for (top left) 3-140, (top right) 4-140, (bottom left) 6-140, and (bottom right) 10-140 beams before and after applying the efficiency and migration correction.

4.5 Calculation of Uncertainties

The corrections outlined above recover the ideal distributions given the well-understood biases in the MAUS reconstruction algorithm, but do not account for statistical or systematic uncertainties, the latter of which dominate the measurement uncertainty.

4.5.1 Statistical Uncertainties

Statistical uncertainties in the measured change in emittance arise from the finite sample size of the muons selected for study. As the upstream sample is predefined as the selected muons to be studied, the statistical uncertainty for each bin of this distribution is zero, ignoring any statistical uncertainty from finite sampling of the underlying beam distribution. Instead, only stochastic processes contribute statistical uncertainty to the measurement; hence the statistical uncertainties considered occur with respect to the resulting scattering and energy loss effects within the absorber and their impact on the downstream sample.

To quantify the statistical uncertainty downstream, first consider the number of events in both the i^{th} bin of the upstream $\mathbf{A}^{\mathbf{F}}$ distribution and the j^{th} bin of downstream $\mathbf{A}^{\mathbf{r}}$ distribution, $A_i^{\mathbf{r}US} \cup A_j^{\mathbf{r}DS}$, forming a migration matrix with elements \mathcal{E}_{ij} . Defining the reconstruction of an event from the i^{th} bin upstream in the j^{th} bin downstream as a successful Bernoulli trial, then the probability of success p is the ratio

$$p = \frac{A_i^{\mathbf{r}US} \cup A_j^{\mathbf{r}DS}}{A_i^{\mathbf{r}US}}, \quad (4.7)$$

and the statistical error on the number of events in \mathcal{E}_{ij} is binomially distributed. Each bin uncertainty $\Delta\mathcal{E}_{ij}$ is given by the width of the 68% confidence interval, and the uncertainty of the j^{th} bin of $\mathbf{A}^{\mathbf{r}DS}$ is then the sum in quadrature over all upstream bins,

$$PDF \text{ Error} = \sqrt{\sum_i (\Delta\mathcal{E}_{ij})^2}. \quad (4.8)$$

The largest relative statistical error on PDF bins below 30 mm in amplitude (the approximate aperture scraping limit) is 2.5%.

The statistical uncertainty in the cumulative distribution is the sum in quadrature of the error on each bin which contributes to the cumulative bin considered. For the n^{th} cumulative bin,

$$CDF \text{ Error} = \sqrt{\sum_{j \leq n} \sum_i (\Delta\mathcal{E}_{ij})^2}. \quad (4.9)$$

4.5.2 Systematic Uncertainties

Systematic errors in the analysis can be introduced in the reconstruction methods or sample selection and can produce a bias in the measurement. These errors are separated into two types: detector systematic errors and performance systematic errors. Each type is calculated uniquely, with both combined in the final error by summing in quadrature.

Detector Systematic Errors

Detector systematic errors are caused by a misunderstanding of the physical placement, composition or operation of the detector, leading to an improper description of the detector or its place within the beamline. By reconstructing events using an inappropriate geometry, the reconstructed phase-space distributions can be systematically biased. The effect of such uncertainties can be studied by simulating events with a systematically deviated geometry while reconstructing with the nominal geometry and comparing the resulting measurement against the nominal geometry simulation. The uncertainty corresponding to each of the following reconstruction errors is analysed with this method in isolation assuming independent systematic uncertainties, with the total contribution found by summing the contributions in quadrature. These contributions are:

- +3 mrad horizontal rotation of TKU
- +3 mrad horizontal rotation of TKD
- +3 mm horizontal displacement of TKU
- +3 mm horizontal displacement of TKD
- +3% SSU Centre coil current
- +3% SSD Centre coil current
- +5% SSU End 1 coil current
- +5% SSU End 2 coil current
- +5% SSD End 1 coil current
- +5% SSD End 2 coil current
- +50% increase in tracker glue density

Owing to the symmetry of the MICE cooling channel, it is assumed that the systematics are two-sided, such that an opposite sign systematic shift

corresponds to an equivalent opposite sign uncertainty. The horizontal position and rotation shifts correspond to the uncertainty on the tracker and its sub-components' positions in the beamline. Although regular surveys of the MICE hall were made after each alteration to the experiment configuration, shifts under magnet powering can occur. Additionally, reconstruction uncertainties from mischaracterisation of the magnet coil currents or solenoid field strength are represented in the coil powering uncertainties, with the tracker glue density uncertainty accounting for the potential variation in the ratio of silicon beads to glue in the tracker plane glue mixture.

Systematic errors on the reconstruction are contained in the correction factors and are included by comparison of each systematically deviated simulation's corrections when applied to the nominal geometry amplitude distribution. For each corrected amplitude bin, A_i , the error is given by $A_i^{sys} - A_i^{ref}$, where \mathbf{A}^{sys} is the systematically varied simulation's correction applied to the nominal distribution, \mathbf{A}^{ref} is the nominally corrected reference distribution, and errors from each systematic are added in quadrature per bin.

Performance Systematic Errors

Performance systematic errors arise due to the impact of event selection and beam quality on the channel performance. The choice of the selection criteria used can impact the measured performance of the detector and misrepresent the cooling effect. To account for this, the following causes of error are considered for both real and simulated events:

- -2 mm TKD fiducial radius selection requirement
- +0.3 TKD $\chi^2/N_{D.O.F}$ threshold selection requirement

The systematic shown on the fiducial radius selection accounts for the possible inclusion of muons which scrape the beamline apertures, thereby scattering or losing energy and biasing the measured performance. Such events should not be included in the sample and are nominally removed by the requirement that the maximum calculated excursion from the tracker axis using the reconstructed phase-space variables is below 150 mm. In practice, mis-reconstruction of phase-space variables may include additional events which should be excluded. A variation of 2 mm corresponds to a misreconstruction of p_{\perp} by ~ 2 MeV/c. Similarly, the sample selection requirement which removes poorly reconstructed muons may preferentially reject muons in subsections of phase-space. An increase in the reconstruction goodness-of-fit threshold is tested.

These systematics are evaluated by comparing the performance of a re-sampled downstream sample with varied selection requirements against the

nominal sample performance. The performance uncertainties are characterised by comparison of the upstream to downstream bin migrations which describe the cooling performance of the channel for a given beam. For a row-normalised¹ US to DS migration matrix, \mathcal{M} , $\mathcal{M}_{ij}^{err} = \mathcal{M}_{ij}^{sys} - \mathcal{M}_{ij}^{ref}$ provides each bin of the error matrix, with errors from each systematic added in quadrature per bin. The error matrix is then multiplied by the bin contents of each reconstructed upstream reference sample as $n_j^{err} = \sum_i \mathcal{M}_{ij}^{err} n_i^{tku}$ to give the systematic variation in downstream PDF bin populations for that sample.

Phase-Space Variables Reconstruction

To ensure reasonable performance of the experiment in the case of the above systematic uncertainties, the resulting phase-space variable distributions can be compared against the nominal case. Figures 4.32, 4.33, 4.34, 4.35 show the resulting variation in reconstructed beam distributions in the TKU reference plane compared to the reconstructed data and MC events for the same beam. Similarly, figures 4.37, 4.38, 4.39, 4.40 show the resulting variation in reconstructed beam distributions in the TKD reference plane compared to the reconstructed data and MC events for the same beam.

The 1D distributions for the nominal and systematically varied hybrid MC samples largely reflect the real data distributions in each case. Some common deviations which derive from the sampling procedure are present, notably in the case of the p_x and p_y distributions in TKU, as well as the max fiducial radius in TKD, but the overall shape of these distributions is still in good agreement with MC and reconstructed data. For most 1D representations, the variation between systematic MC productions is small on the scale of the distribution being considered. The most obvious variations are seen when considering the χ^2 per degree of freedom in TKU and TKD, where some samples show large changes induced in the goodness of fit. However, the perturbations produced by each sample are expected to be small in a 1D representation, with larger contributions when considering a multidimensional distribution such as that represented by particle amplitudes.

¹Such that the sum of elements in each row equals 1.




Figure 4.32: Comparison of reconstructed TKU x-position distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.33: Comparison of reconstructed TKU y-position distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.34: Comparison of reconstructed TKU p_x momentum distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.35: Comparison of reconstructed TKU p_y momentum distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.36: Comparison of reconstructed TKU χ^2 distributions at the reference plane with TKU systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.37: Comparison of reconstructed TKD x-position distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.38: Comparison of reconstructed TKD y-position distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.

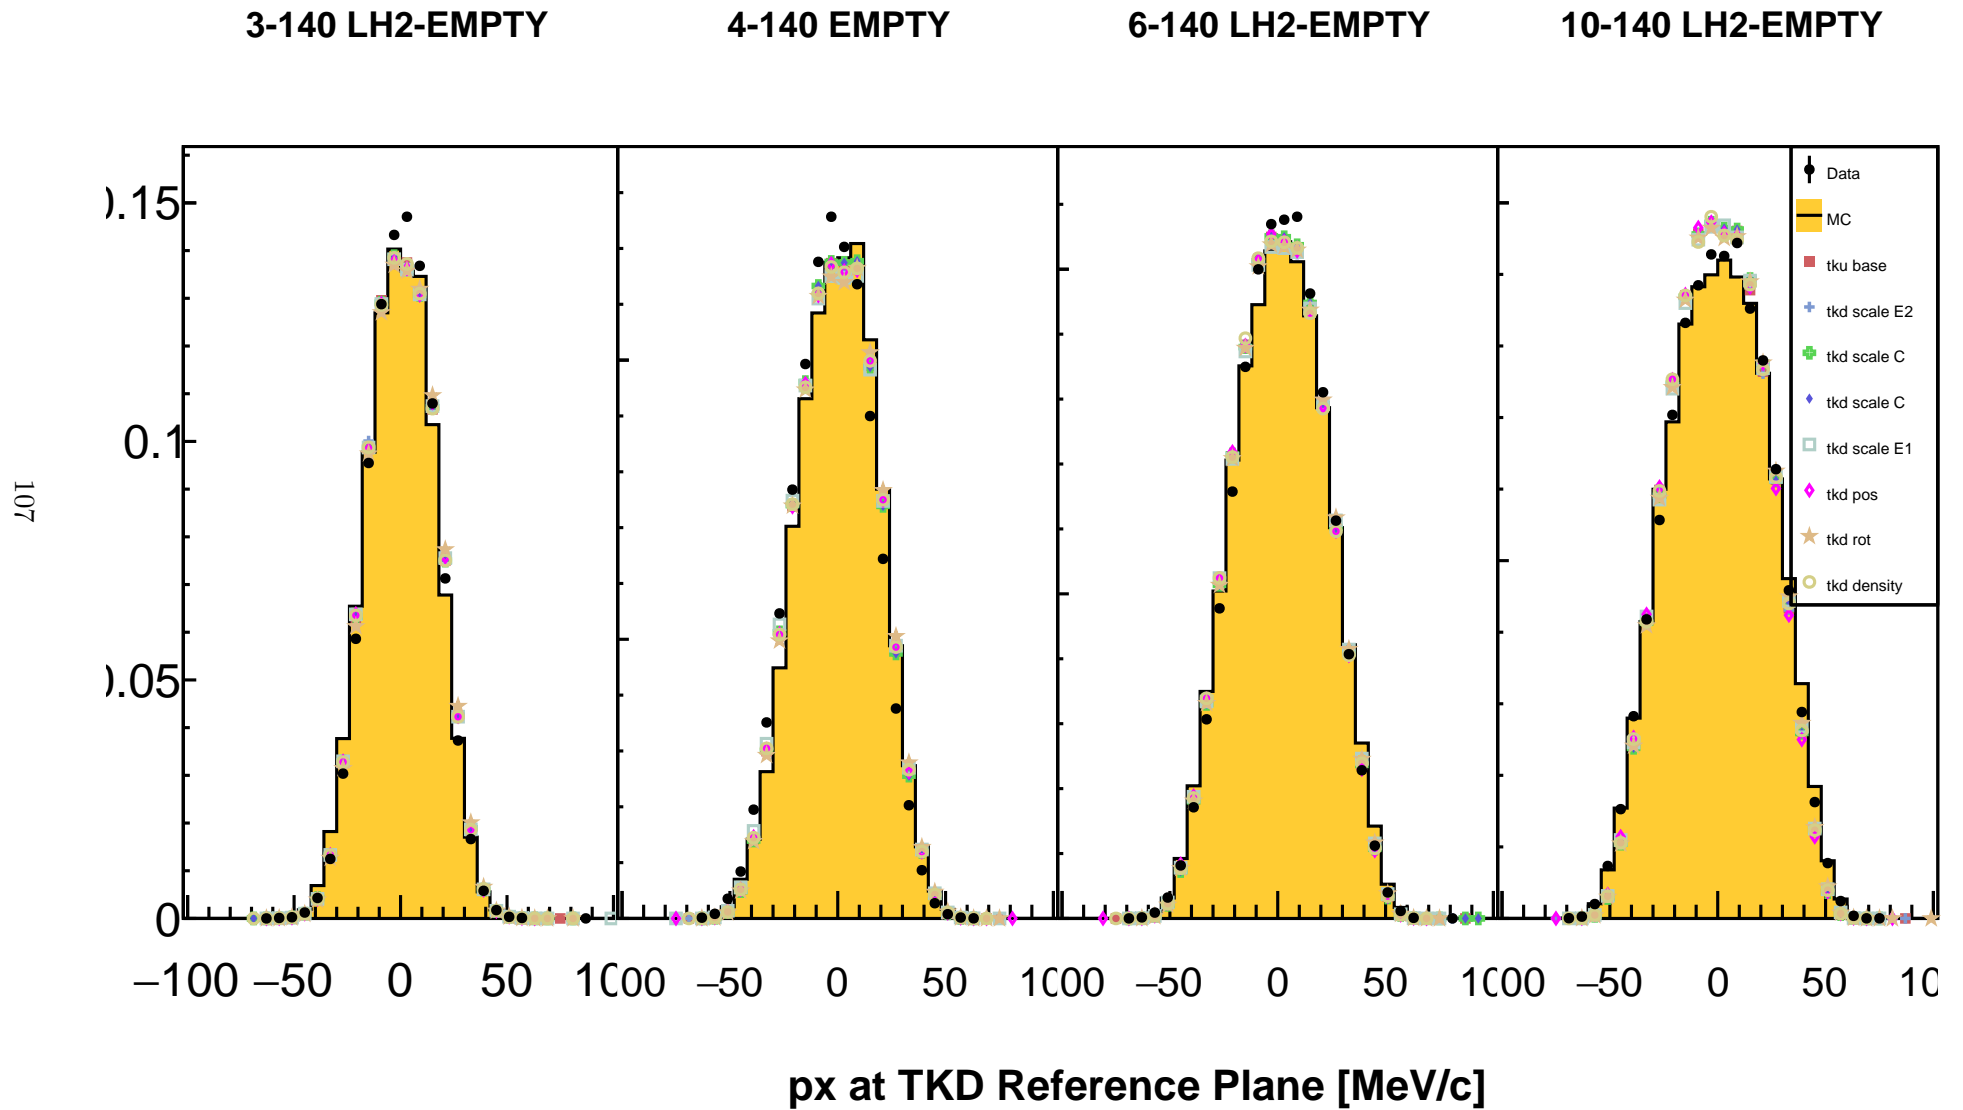


Figure 4.39: Comparison of reconstructed TKD p_x momentum distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.40: Comparison of reconstructed TKD p_y momentum distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.41: Comparison of reconstructed TKD χ^2 distributions at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.




Figure 4.42: Comparison of TKD maximum radial excursion at the reference plane with TKD systematics. Data is shown in black, MC shown in yellow, with other coloured points representing the base hybrid MC and different systematically varied reconstructions.

Correction Uncertainties

The corrections described in section 4.4.4 are directly altered by the systematic reconstruction uncertainties described above. The effect this has on the resulting corrections is studied with the use of systematically shifted MC geometries and shown directly here.

The change induced in the downstream efficiency correction through each downstream systematic varied simulation is shown in figure 4.43. At large amplitudes, above 30 mm, the variation in the correction is dominated by statistical uncertainty. As the efficiency correction upstream is small, the systematic uncertainty there is not shown.

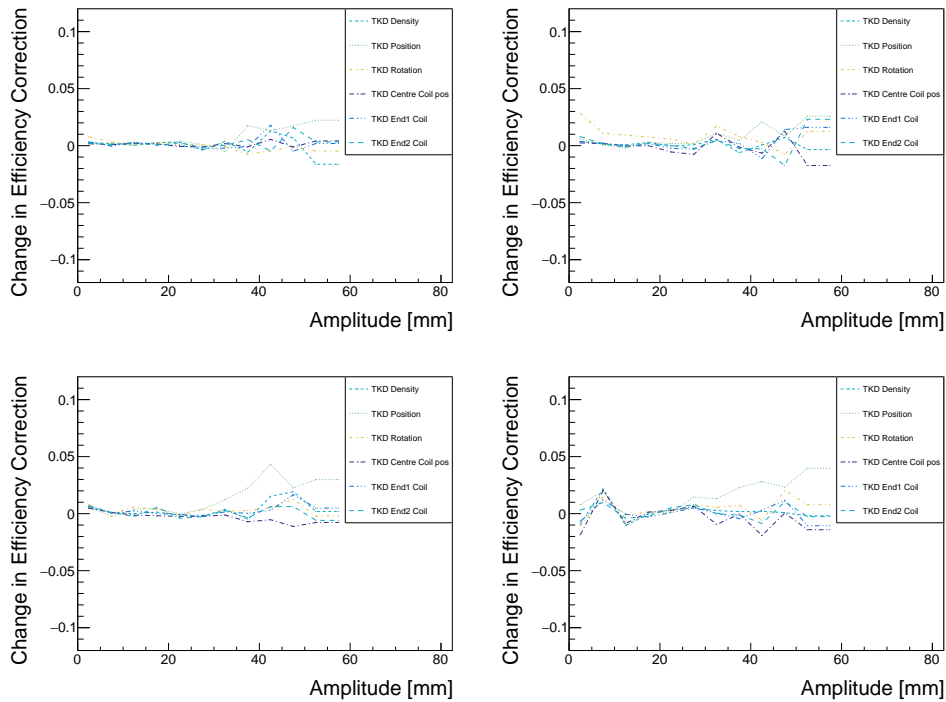


Figure 4.43: Change in the downstream efficiency correction as a result of each systematic variation of the simulation geometry, shown for 3-140 (top left), 4-140 (top right), 6-140 (bottom left), and 10-140 (bottom right) beams.

Similar plots showing the change induced in the leading diagonal terms of the calculated migration matrix of each beam by each systematic variation are shown in figures 4.44 and 4.45. For low amplitudes, the diagonal correction terms experience variations below the 5% level.

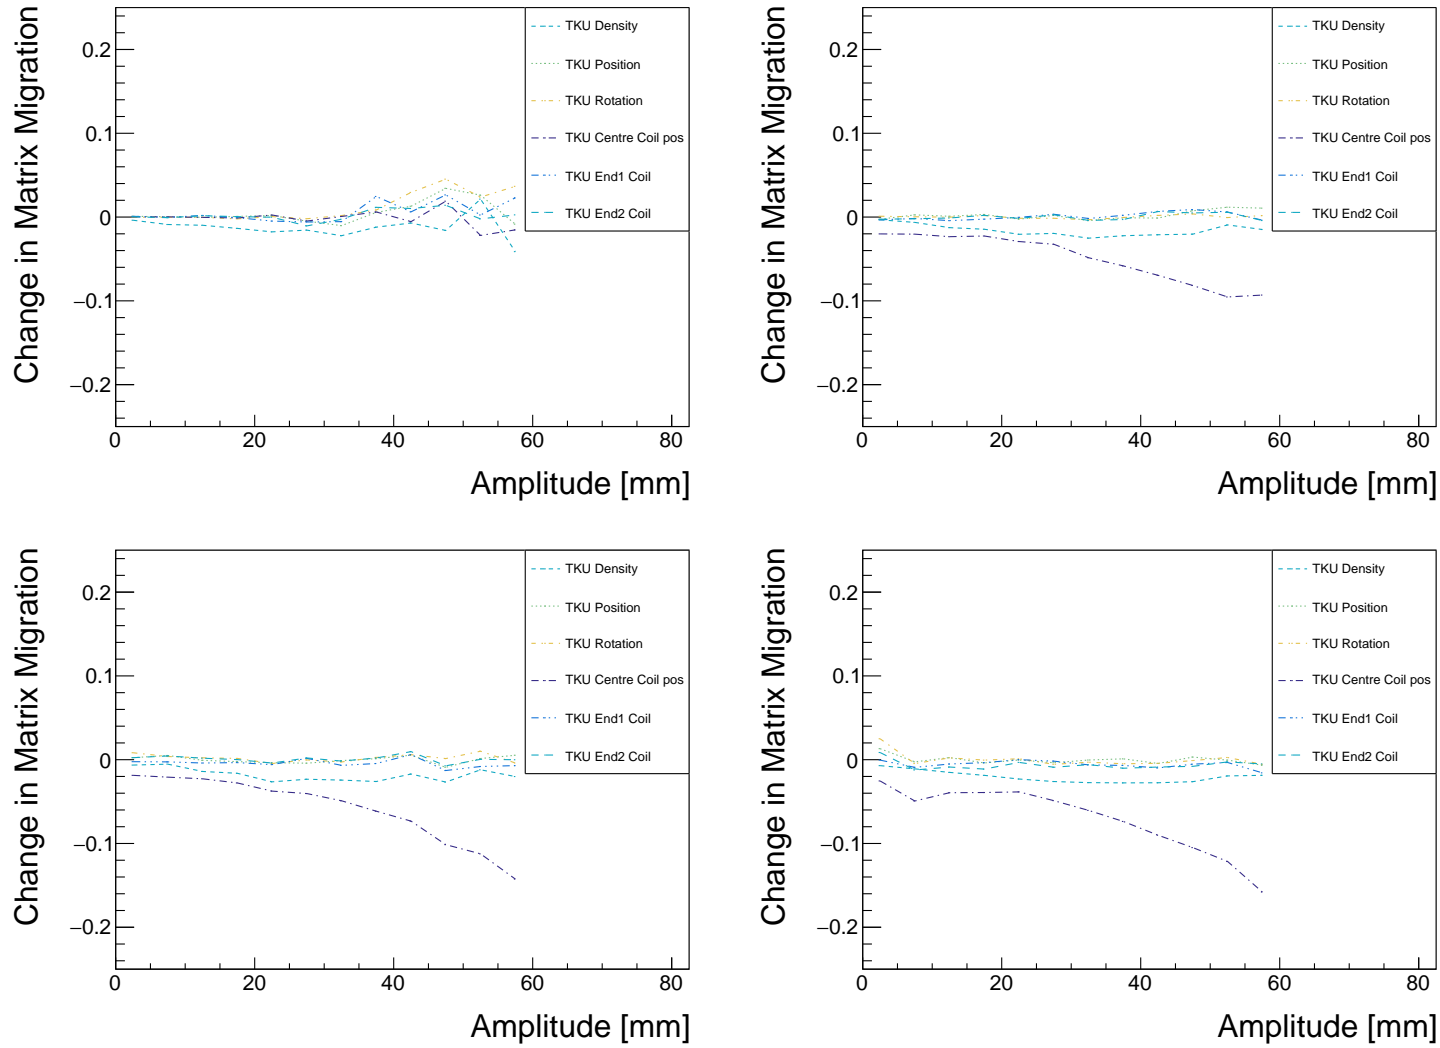


Figure 4.44: Change of the diagonal terms in the upstream migration matrix as a result of each systematic variation of the simulation geometry, shown for 3-140 (top left), 4-140 (top right), 6-140 (bottom left), and 10-140 (bottom right) beams. Errors are zero by definition currently as there are no performance systematics on the upstream distributions.

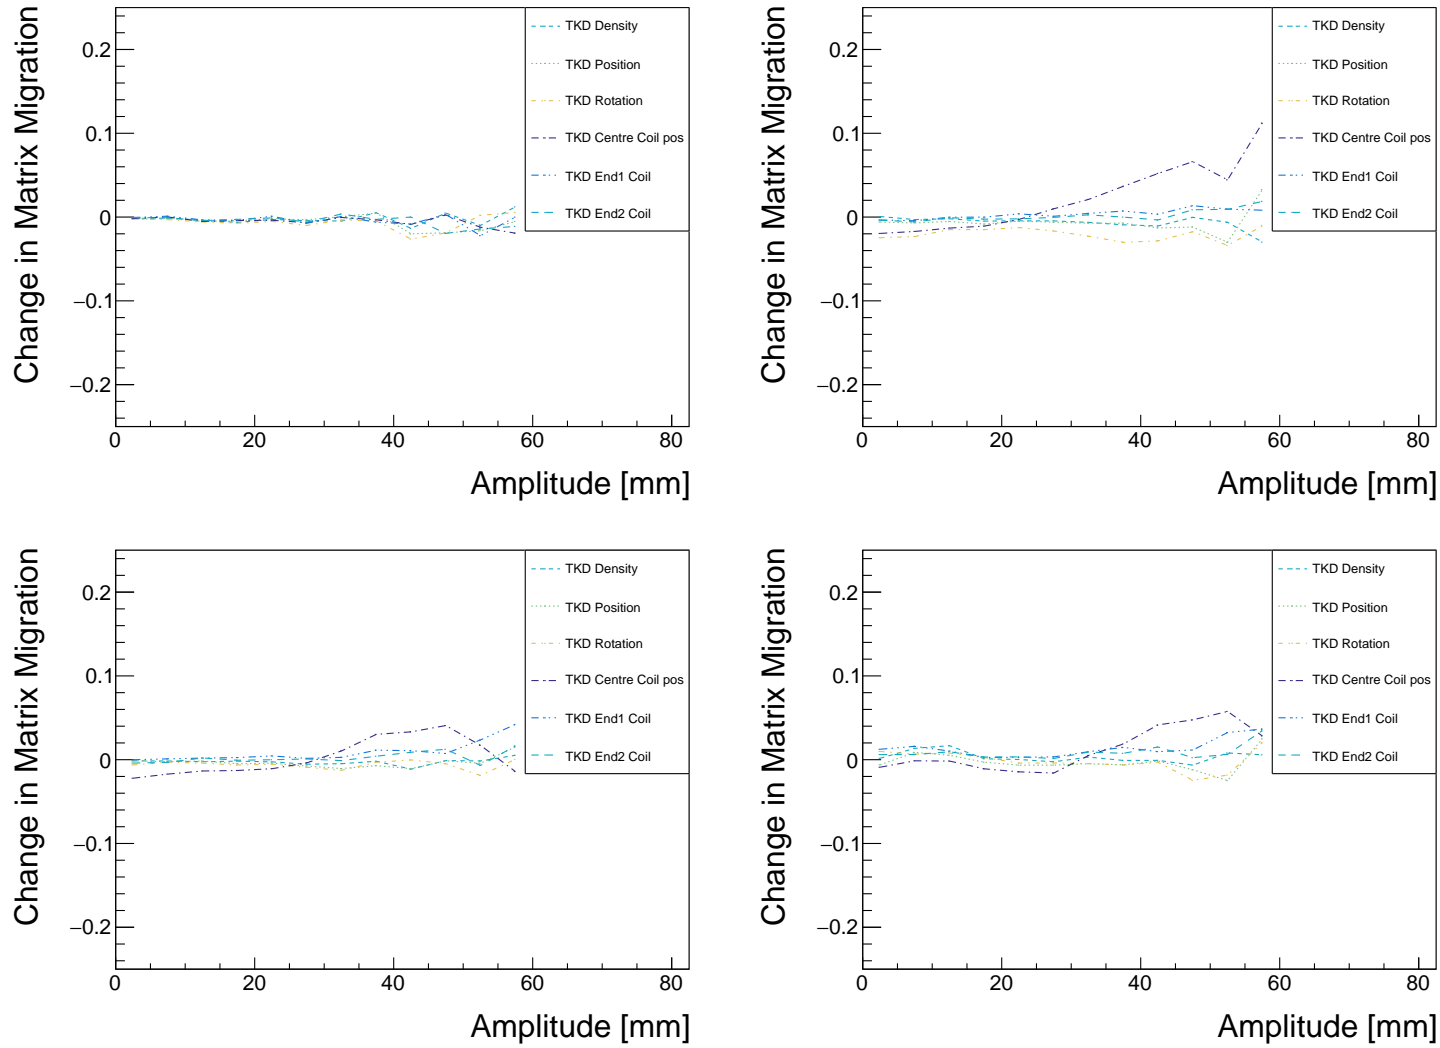


Figure 4.45: Change of the diagonal terms in the downstream migration matrix as a result of each systematic variation of the simulation geometry, shown for 3-140 (top left), 4-140 (top right), 6-140 (bottom left), and 10-140 (bottom right) beams.

4.6 Results

The distributions of calculated particle amplitudes are discussed here, shown with corrections applied. The relative sizes of these corrections are also discussed, and the pre- and post-corrected distributions are shown.

4.6.1 Amplitude PDFs and CDFs

The resulting amplitude particle distributions (PDFs) and cumulative distributions, integrated from zero, (CDFs) are shown for reconstructed data events in figures 4.46, 4.47, 4.48 and 4.49. Simulated events are shown in figures 4.50, 4.51, 4.52 and 4.53. Upstream distributions are shown in red, with downstream distributions in green. The widths of the coloured bands represent the combined statistical and systematic errors for each distribution.

The higher emittance beams, produced by the presence of additional diffuser segments in the beam, show increasing spread in upstream amplitudes, with more extensive distribution tails. As the 4 mm nominal emittance beam enters close to the equilibrium emittance for the channel, the downstream distribution for the 4-140 LiH beam (figure 4.46) closely resembles the upstream distribution. For higher nominal emittance beams in the presence of an absorber (6-140 LiH in figure 4.46, 6-140 LH2 and 10-140 LH2 beams in figure 4.47), the downstream distributions show a clear increase in the relative number of low amplitude muons. This effect is more pronounced in the CDF distributions in figures 4.48 and 4.49, showing an increase in cumulative density for muons below the aperture scraping limit. This cooling is also evident in the simulated distributions. Where the beam traverses the empty liquid hydrogen vessel or an empty vacuum, no cooling is observed. The 3-140 distributions are below the equilibrium emittance and hence experience an overall increase in emittance as the heating term is dominant. This causes low amplitude muons in the upstream distribution to move to higher amplitudes downstream. Again, this heating is replicated in the simulated distributions.




Figure 4.46: Normalised amplitude distributions for reconstructed data events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Blue dashed lines indicate the approximate aperture scraping limit.




Figure 4.47: Normalised amplitude distributions for reconstructed data events traversing the LH₂ empty and LH₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Blue dashed lines show the approximate aperture scraping limit. Blue dashed lines indicate the approximate aperture scraping limit.




Figure 4.48: Normalised cumulative amplitude distributions for reconstructed data events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Blue dashed lines indicate the approximate aperture scraping limit.




Figure 4.49: Normalised cumulative amplitude distributions for reconstructed data events traversing the LH₂ empty and LH₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Blue dashed lines indicate the approximate aperture scraping limit.




Figure 4.50: Normalised amplitude distributions for reconstructed MC events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions.




Figure 4.51: Normalised amplitude distributions for reconstructed MC events traversing the LH₂ empty and LH₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions.




Figure 4.52: Normalised cumulative amplitude distributions for reconstructed MC events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions.




Figure 4.53: Normalised cumulative amplitude distributions for reconstructed MC events traversing the LH₂ empty and LH₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions.

4.6.2 Corrected vs Uncorrected Distributions

A comparison of the uncorrected amplitude distribution populations with their corrected values is shown in figures 4.54 and 4.55. As the corrections required for upstream distributions at low amplitudes are very small, the corrected and uncorrected distributions match closely. The downstream correction is small relative to the ionisation cooling effect. This cooling is clear even where a correction has not been applied. A greater correction effect is shown in the downstream distributions, largely owing to the increased contribution from the efficiency correction downstream.

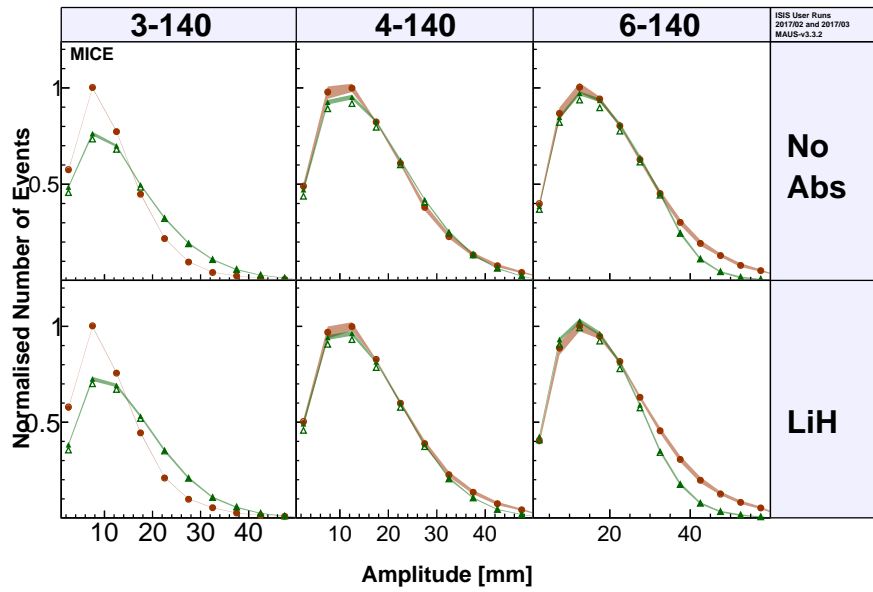


Figure 4.54: (Hollow) Uncorrected and (filled) corrected amplitude distributions for reconstructed data events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Where hollow markers are not visible, they coincide with the filled markers.




Figure 4.55: (Hollow) Uncorrected and (filled) corrected amplitude distributions for reconstructed data events traversing the LH₂ empty and LH₂ full absorber states. Coloured bands show combined statistical and systematic errors for (red) upstream and (green) downstream distributions. Where hollow markers are not visible, they coincide with the filled markers.

4.6.3 PDF and CDF Ratios

From equation 1.18, it is clear that the normalised RMS emittance of the beam is proportional to the mean of the particle amplitude distribution. For a beam well-described by a multivariate Gaussian distribution, amplitudes are distributed as

$$f(A_{\perp}) = \frac{A_{\perp}}{4(\varepsilon_{\perp}^n)^2} \exp\left(\frac{-A_{\perp}}{2\varepsilon_{\perp}^n}\right) \quad (4.10)$$

and hence the upstream and downstream amplitude distributions $f^u(A_{\perp})$, $f^d(A_{\perp})$ are related to the upstream and downstream emittances $\varepsilon_{\perp}^{n|u}$, $\varepsilon_{\perp}^{n|d}$ by

$$\frac{f^d(A_{\perp})}{f^u(A_{\perp})} = \left(\frac{\varepsilon_{\perp}^{n|u}}{\varepsilon_{\perp}^{n|d}}\right) \exp\left[-\frac{A_{\perp}}{2}\left(\frac{1}{\varepsilon_{\perp}^{n|d}} - \frac{1}{\varepsilon_{\perp}^{n|u}}\right)\right]. \quad (4.11)$$

Hence, ratios greater than one demonstrate an increased muon density in the beam's core corresponding to an increased emittance.

The ratio of downstream over upstream PDFs and CDFs for both reconstructed data and simulated events are shown in figures 4.56, 4.57, 4.58 and 4.59. This increased core density is shown for the 6 and 10 mm emittance absorber-in beams, indicating a clear cooling signal in both simulated events and data. Ratios below one in the region above 30 mm indicate muons lost from striking the beam pipe or outside the downstream tracker fiducial volume. Heating is visible in the 3 mm beams where the ratio of cumulative distributions is below one, also consistent in both simulated events and data. In contrast, the control samples taken with no absorber both show no consistent heating, with beam losses from aperture scraping emerging at large amplitude. However, this is not the case for the 3 mm beam, where beam heating is observed even with no absorber present. This is due to poor matching with the cooling channel, resulting in non-linear emittance growth.

Whilst the PDF distributions show a bin-by-bin increase for the number of events at low amplitudes and a corresponding decrease for muon density at larger amplitudes, the CDF distributions make explicit where this effect arises due to ionisation cooling and where particle loss at higher amplitudes due to aperture scraping begins, evidenced by the fall in cumulative muon ratio below one shown for high nominal emittance beams. For the 3 mm beam, this loss at high amplitudes is not seen, as without the diffuser inserted into the beam, muons are well-contained throughout the entire cooling channel.

This result provides an important confirmation of the ionisation cooling principle, shown for the first time in a 'solenoid mode' cooling channel, where the magnetic field is operated with constant polarity across the absorber. This result is important in validating the muon ionisation cooling model as a viable tool for achieving high brightness muon beams. In such future facilities, a




Figure 4.56: Ratios of the amplitude PDF distributions for reconstructed data and MC events traversing the no absorber and LiH absorber states. The y-axis for the 3-140 beam is extended to accommodate the larger variation in PDF ratio. Coloured bands show combined statistical and systematic errors for (blue) measured data and (red) simulated MC. Blue dashed lines indicate the approximate aperture scraping limit.




Figure 4.57: Ratios of the amplitude PDF distributions for reconstructed data and MC events traversing the LH₂ empty and LH₂ full absorber states. Coloured bands show combined statistical and systematic errors for (blue) measured data and (red) simulated MC. Blue dashed lines indicate the approximate aperture scraping limit.




Figure 4.58: Ratios of the amplitude CDF distributions for reconstructed data and MC events traversing the no absorber and LiH absorber states. Coloured bands show combined statistical and systematic errors for (blue) measured data and (red) simulated MC. Blue dashed lines indicate the approximate aperture scraping limit.




Figure 4.59: Ratios of the amplitude CDF distributions for reconstructed data and MC events traversing the LH₂ empty and LH₂ full absorber states. Coloured bands show combined statistical and systematic errors for (blue) measured data and (red) simulated MC. Blue dashed lines indicate the approximate aperture scraping limit.

combination of 'solenoid mode' and 'flip mode' field arrangements may provide the most efficient route to meeting its cooling requirements.

Chapter 5

Canonical Angular Momentum

Canonical angular momentum is an important beam property for ionisation cooling. Designs for an extended ionisation cooling channel are typically composed of many repeating absorbers with intervening RF accelerating cavities and solenoidal magnetic channels. For muons propagating through an extended ionisation cooling channel, the canonical angular momentum of the muon beam directly impacts the expected cooling performance with each pass, with larger canonical angular momentum resulting in less emittance reduction and a higher equilibrium emittance [13, 50]. Furthermore, the horizontal and vertical emittances are coupled by the canonical angular momentum, such that the transverse emittance is larger for beams with larger canonical angular momentum, as is the beams spot size.

Though a beam may enter a cooling channel with relatively little canonical angular momentum, this property is not constant under the non-conservative forces present during muon interactions with the absorber. Where the sign of the B_z field component remains constant throughout the extended cooling channel, this results in a net increase in the canonical angular momentum of the beam with each absorber pass. This can be managed or avoided entirely by alternating the sign of the B_z field component across the absorber, as discussed in section 2.3.3. To ensure the required cooling performance of future facilities can be attained, the canonical angular momentum growth in the beam must be both understood and managed.

5.1 Canonical Angular Momentum Build-up

For particles propagating through a cylindrically symmetric solenoid, canonical angular momentum is a constant of the motion and a conserved quantity. This conservation constraint is associated with the symmetry of the system, such

that the Hamiltonian is independent of azimuthal angle [11, 51].

The canonical angular momentum of single particles inside a cylindrically symmetric solenoid with its long axis oriented along the z -axis is given by

$$L_{\text{canon}} = r(p_\theta + q\mathcal{A}_\phi) = xP_y - yP_x + q(x\mathcal{A}_y - y\mathcal{A}_z) \quad (5.1)$$

where \mathcal{A} is the magnetic vector potential, r is the radial distance from the z -axis, q is the particle charge [13]. Inside a cylindrically symmetric solenoid, the magnetic field is only a function of z , $B(z)$. To satisfy Ampere's Law, the vector potential must point in the direction of the polar angle, ϕ , around the long axis: $\vec{B} = \nabla \times [\mathcal{A}_\phi(r, z)e_\phi]$. For a solenoid focusing system such as the one in MICE, a paraxial approximation of \mathcal{A}_ϕ is given by [13]

$$\mathcal{A}_\phi \approx \frac{r}{2}B(z) - \frac{r^3}{16}B''(z) + \mathcal{O}(r^5). \quad (5.2)$$

Equation 5.1 can be decomposed into the contribution from the particle kinematics,

$$L_{\text{kin}} = xP_y - yP_x, \quad (5.3)$$

and the contribution from the magnetic field,

$$L_{\text{field}} = qr\mathcal{A}_\phi. \quad (5.4)$$

The net canonical angular momentum is represented in the beam dynamics by the dimensionless parameter \mathcal{L} which arises naturally in the beam covariance matrix given in equation 1.13. The parameter \mathcal{L} is related to the mean canonical angular momentum, $\langle L_{\text{canon}} \rangle$, by

$$\langle L_{\text{canon}} \rangle \approx \langle xP_y - yP_x + \kappa r^2 P_z \rangle = 2mc\varepsilon_n \mathcal{L}, \quad (5.5)$$

with

$$\kappa \equiv \frac{qB(z)}{2P_z} \approx \frac{q\mathcal{A}_\phi}{P_z r} = \phi'. \quad (5.6)$$

The remaining beam covariance quantities are given by the Courant-Snyder optics parameters α , β , and γ [13]. Moreover, the dimensionless parameter \mathcal{L} directly relates the envelope transverse beta function, β_\perp , to the single particle beta function, β_p , by

$$\beta_\perp = \beta_p \sqrt{1 + \mathcal{L}^2} \quad (5.7)$$

for a matched beam. We see that for a beam with large canonical angular momentum, it is more difficult to focus and hence cool than an equivalent beam of lower canonical angular momentum.

A beam entering a solenoid focusing system with no kinetic or canonical

angular momentum will pick up kinetic angular momentum in the fringe solenoid field, but will lose this again on exiting the solenoid field. Throughout, the canonical angular momentum will remain zero, with the field and kinetic terms from eqns. 5.3 and 5.4 cancelling. Although a beam under conservative forces will not change its canonical angular momentum, the presence of an absorber in the path of beam leads to non-conservative scattering and energy loss which may also change the kinetic angular momentum of the beam. The change in the kinetic angular momentum of a beam as it passes through material is proportional to the change in its longitudinal momentum,

$$\Delta L_{kin}/L_{kin} \approx \Delta p_z/p_z. \quad (5.8)$$

As the particle kinetic angular momentum has changed, the canonical angular momentum is no longer zero and the particle will exit the solenoid having acquired some additional kinetic angular momentum. This effect motivated the introduction of ‘flip’ mode cooling channel operation; as the magnetic field component B_z is responsible for the Larmor oscillation which imparts the beam with a large kinetic angular momentum, if B_z across the absorber is zero, then ΔL_{kin} is zero. Hence flipping the solenoid field over successive absorber passes ensures the field at the absorber remains close to zero, and therefore the change in mechanical, and hence canonical, angular momentum stays close to zero. Alternatively, a larger number of absorber passes can be made in solenoid mode with occasional field flipping, inducing an opposite-sign change in mechanical angular momentum for opposite-sign field ‘solenoid’ mode operation [13].

5.2 Analysis

To evaluate the canonical angular momentum growth that the beam undergoes in solenoid mode, the analysis considers the particles in US and DS samples discussed in sections 4.2.1, 4.2.2, and 4.2.3. The canonical angular momentum of single particles in MICE is calculated from the reconstructed particle properties using equation 5.1, with \mathcal{A}_ϕ given by equation 5.2 to first order. The field map model for the cooling channel provides B_z values at each trackpoint.

5.2.1 Single Particle Distributions

Single particle distributions of the reconstructed L_{kin} and L_{field} components for each DS beam sample are shown for events at TKU station 1 in figures 5.1 and 5.2, with DS beam samples at TKD station 1 shown in figures 5.3 and 5.4. The same sample is used upstream and downstream to explicitly show the change in canonical angular momentum of those muons which remain in the sample after passing through the absorber. As the field polarity remains

constant throughout the channel, the contributions from L_{field} upstream and downstream are both positive-definite. The reconstructed L_{kin} contributions vary around zero, but are more frequently negative. As the lower nominal emittance beams with fewer diffuser irises inserted into the beamline observe less scattering before reaching the tracker, these beams show an initial L_{kin} closer to zero on average with fewer events in the negatively skewed tail. Similarly, the L_{field} contribution, which grows proportionally with r^2 , shows increasingly more events at large positive values for higher nominal emittance beams as the beam occupies a larger area in x - y phase-space.

The reconstructed single particle distributions of canonical angular momentum at TKU and TKD station 1 are shown in figures 5.5 and 5.6 respectively. A comparison of the two sets of distributions show a positive skew which arises most significantly in those beams where material has been inserted within the AFC module. The "no absorber" beams which have no added material between trackers to propagate through show little change between US and DS trackers.

To better show the difference between canonical momenta upstream and downstream, the change in canonical angular momentum of single muons passing through each absorber module configuration, calculated to first order using the MICE solenoid field model, is shown in figure 5.7. For the no absorber and empty LH₂ case, the distributions remain largely symmetric around zero, whereas a significant positive skew is induced by the presence of the LiH disk and full LH₂ vessel, indicating canonical angular momentum growth.




Figure 5.1: Normalised single particle distributions of reconstructed kinetic angular momentum components, L_{kin} , of the DS muon sample at TKU station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.




Figure 5.2: Normalised single particle distributions of reconstructed field angular momentum components, L_{field} , of the DS muon sample at TKU station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.




Figure 5.3: Normalised single particle distributions of reconstructed kinetic angular momentum components, L_{kin} , of the DS muon sample at TKD station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.




Figure 5.4: Normalised single particle distributions of reconstructed field angular momentum components, L_{field} , of the DS muon sample at TKD station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.




Figure 5.5: Normalised single particle distributions of reconstructed canonical angular momentum, L_{canon} , of the DS muon sample at TKU station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.




Figure 5.6: Normalised single particle distributions of reconstructed canonical angular momentum, L_{canon} , of the DS muon sample at TKD station 1 for 3, 4, 6, and 10-140 beams. Data events are shown in black, MC events in yellow.




Figure 5.7: Normalised single particle distributions of the change in canonical angular momentum of the DS muon sample across the absorber for 3, 4, 6, and 10-140 beams. Reconstructed data events are shown in black, reconstructed MC events in yellow.

5.2.2 Tracker Evolution

The evolution of the mean canonical angular momentum of each beam through the tracker can be calculated by evaluating the mean L_{kin} and L_{field} terms. Statistical errors on the mean are given by the standard error of the sample, $\frac{\sigma}{\sqrt{n}}$. A systematic error is calculated for each beam, employing the systematically varied hybrid MC samples listed in section 4.5.2. As no reconstruction correction is applied, the error on the mean from each systematic sample is assigned simply as the error for the measurement at that station, with the total systematic obtained from errors summed in quadrature.

The kinetic and field angular momentum terms at each tracker station are shown in figures 5.8 and 5.9, with the combined canonical angular momentum in figure 5.10. Each figure shows reconstructed data in black, reconstructed MC in red, and MC truth in lilac. The band represents the combined systematic and statistical error of the sample. Although the systematic errors on L_{kin} and L_{field} are small compared to the measured component, the mean canonical angular momentum of each beam is much smaller than the two component terms as they largely cancel in each case, and so the systematic errors become more significant on the scale of the measured change across the absorber. This is particularly true for the 10-140 beams, where a larger variation in L_{kin} when varying the centre coil yields a larger systematic uncertainty.

A reconstruction bias is evident in each tracker, where the central stations observe a lower canonical angular momentum than those at the extremities of either tracker, apparent in both reconstructed data and reconstructed MC events. This effect is believed to arise due to the track reconstruction's use of a single averaged B_z value for each tracker. A comparison of the reconstructed MC and MC truth canonical angular momentum at each station shows this as a reconstruction effect, as the MC truth beam instead exhibits a small station-by-station increase as the beam experiences scattering in the SciFi planes, with a large increase measured for beams which propagate through material in the AFC. More canonical momentum growth is measured for the two absorbers, with a reduced effect produced by scattering from the empty LH₂ vessel.

Evaluating the covariance matrix of the beam at each detector station, equation 5.5 can be used to calculate the dimensionless parameter \mathcal{L} which represents the net canonical momentum of the beam optics as the beam propagates through the tracker. By constructing such a covariance matrix, the evolution of \mathcal{L} through the trackers is obtained, shown in figure 5.11. The behaviour here mirrors that seen in figure 5.10 as expected, showing significant growth for beams which can scatter on absorber material in the cooling channel, and little change for the no absorber case.

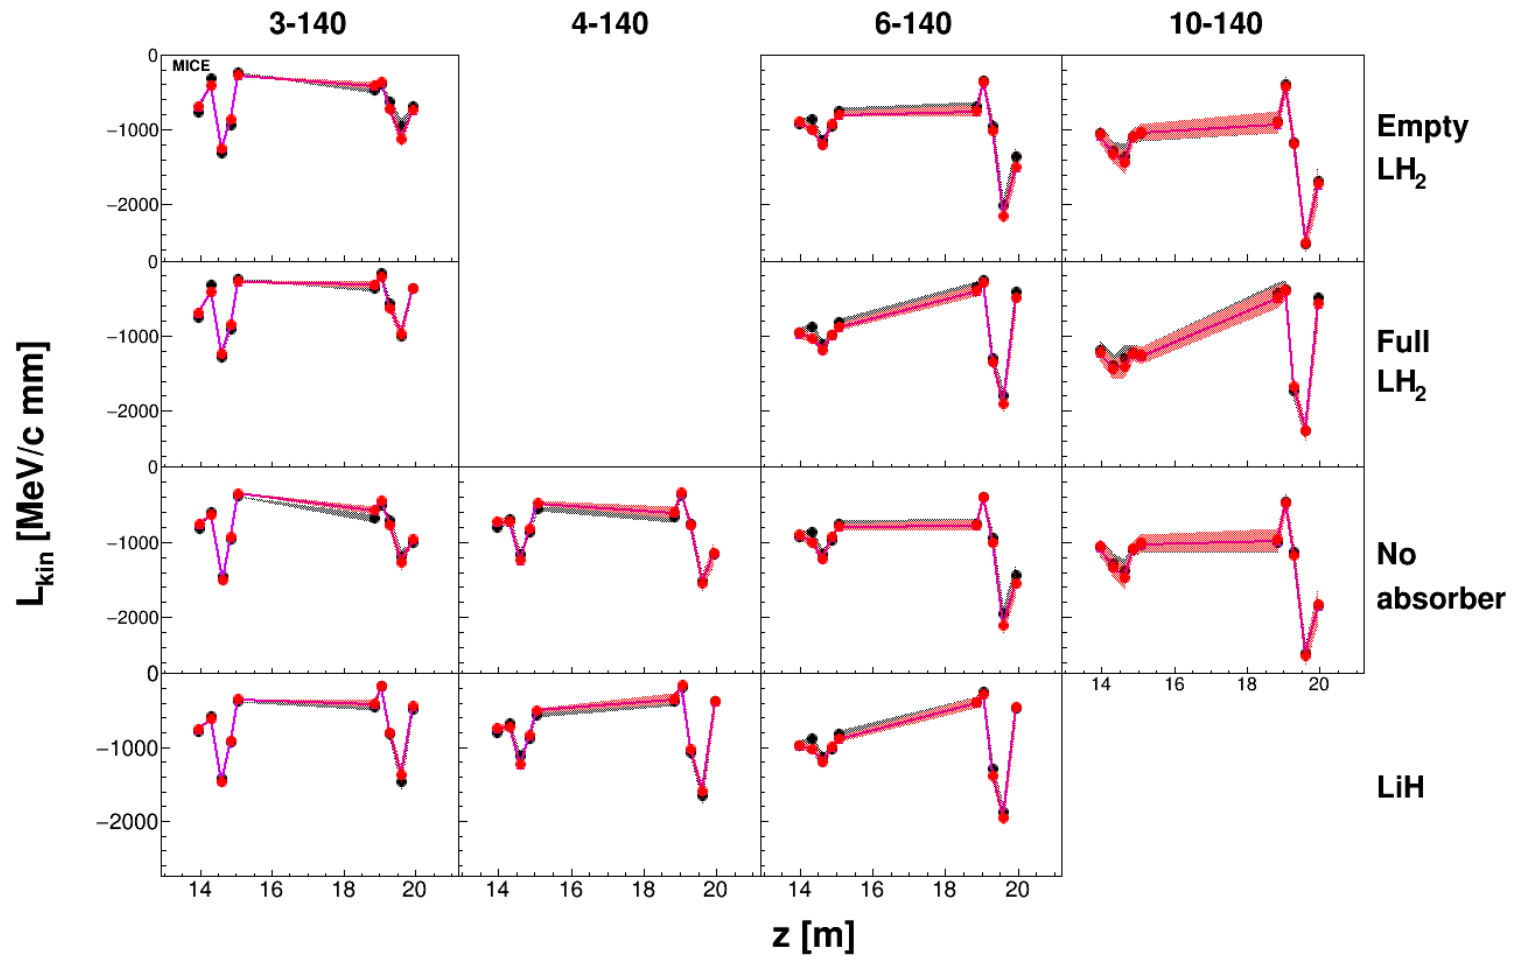


Figure 5.8: Evolution of the mean kinetic angular momentum, L_{kin} , of the beam at each tracker station for beams 3, 4, 6, and 10-140. Black points show reconstructed data events, red points show reconstructed MC events, and purple shows MC truth variables for events in the MC sample.

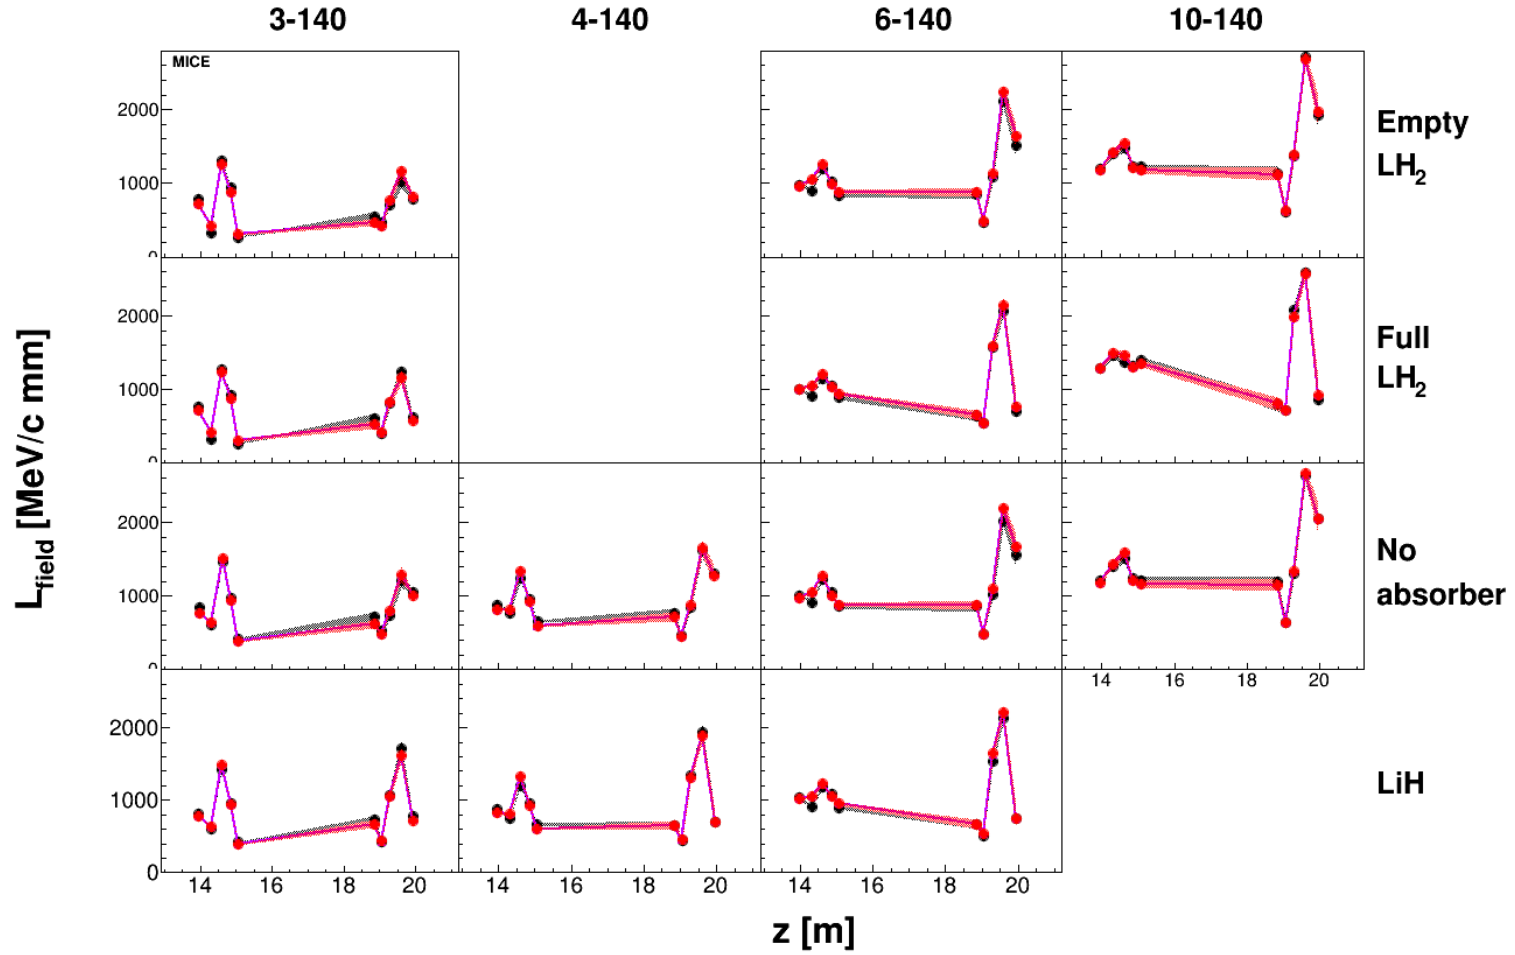


Figure 5.9: Evolution of the mean field angular momentum, L_{field} , of the beam at each tracker station for beams 3, 4, 6, and 10-140. Black points show reconstructed data events, red points show reconstructed MC events, and purple shows MC truth variables for events in the MC sample.

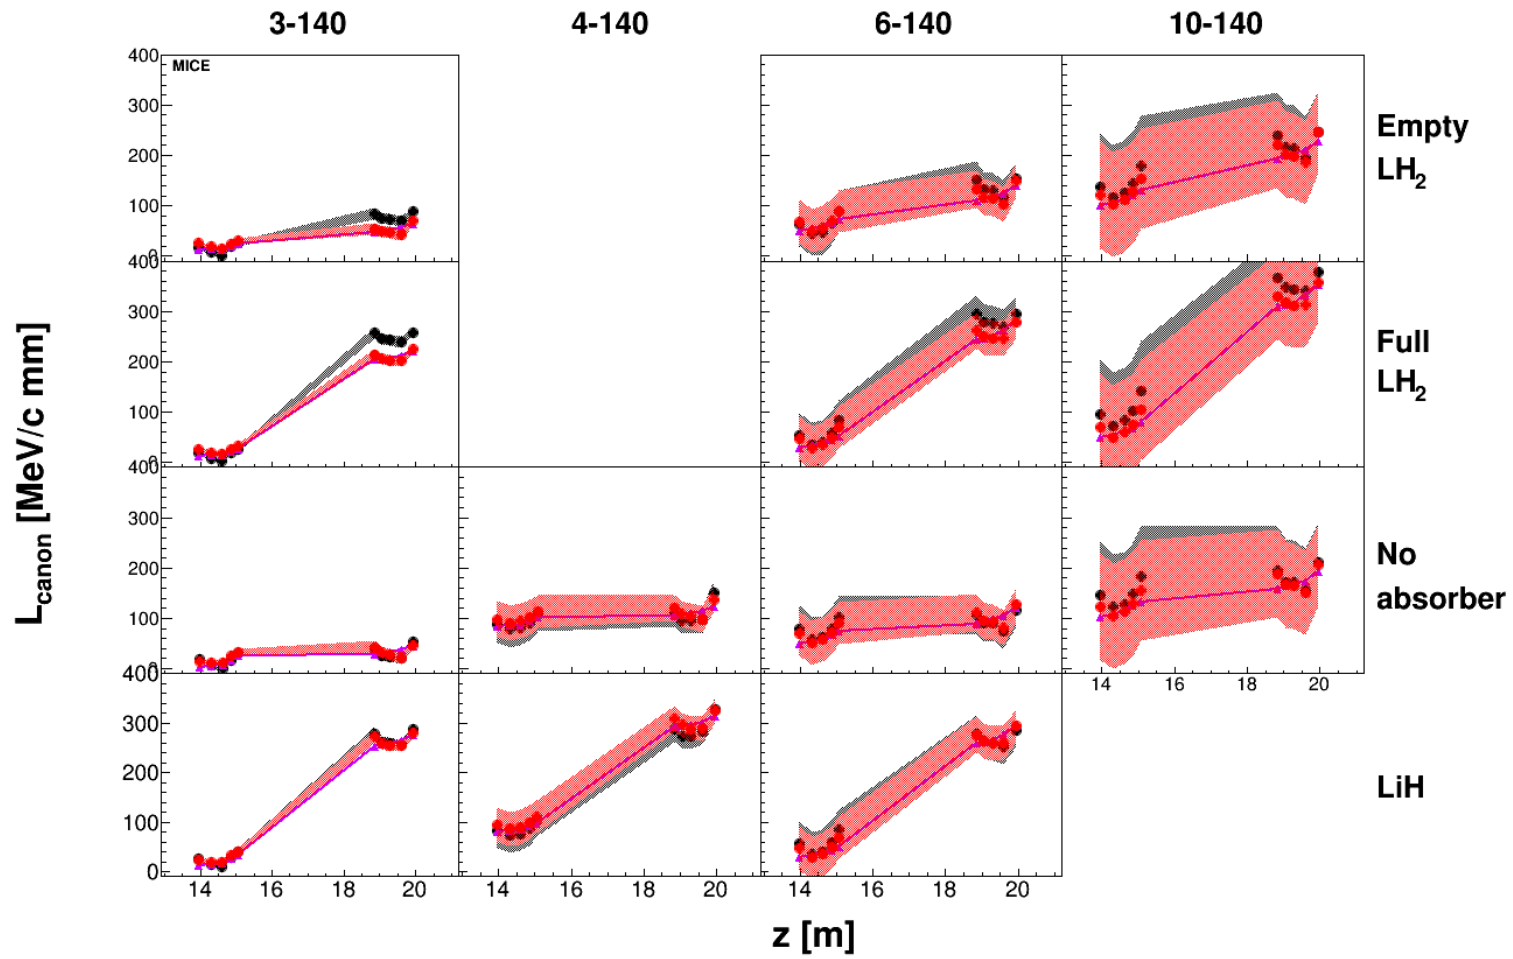


Figure 5.10: Evolution of the mean canonical angular momentum of the beam at each tracker station for beams 3,4,6, and 10-140. Black points show reconstructed data events, red points show reconstructed MC events, and purple shows MC truth variables for events in the MC sample.




Figure 5.11: Evolution of \mathcal{L} through the cooling channel, calculated at each tracker station for beams 3,4,6, and 10-140. Black points show reconstructed data events, red points show reconstructed MC events, and purple shows MC truth variables for events in the MC sample.

5.3 Summary

This work presents the first measurement of muon canonical angular momentum change across the ionization cooling channel in MICE. For this measurement, the solenoidal field polarity was maintained throughout the channel, resulting in a significant measured increase where an absorber was present. This measurement agrees with expectation, as simulated events show the same trend in angular momentum growth as reconstructed data.

Understanding the build-up of canonical angular momentum over ‘solenoid’ mode absorber passes and validating the simulation tools used to model this effect is an important step in enabling a future neutrino factory or muon accelerator facility, as this will guide the design choices and extended cooling channel layouts which may be considered. One alternative to a facility with entirely ‘flip’ mode absorbers or entirely ‘solenoid’ mode absorbers is to alternate the B_z field component only occasionally, with longer segments of alternating sign solenoid field on the order of tens of metres long. This arrangement allows for an increase in canonical angular momentum over short segments which is then reduced back to close to zero over the succeeding opposite-sign segment. Work is on-going to prepare a collaborative publication which compares the canonical angular momentum growth observed in ‘flip’ and ‘solenoid’ modes to better inform this topic.

Chapter 6

Conclusions and Future Work

6.1 Conclusion

The Muon Ionization Cooling Experiment completed data-taking in 2018, having undertaken a significant data-collection program. Analysis of a subset of this data has been presented here. This work has been undertaken by the student, with the support of the MICE collaboration.

The emittance change for a range of muon beams which traverse different absorber media through the cooling channel has been characterised through the calculation of distributions of single particle amplitudes. The resulting distributions show the first measurement of muon ionization cooling in a ‘solenoid’ mode cooling channel, where the on-axis magnetic field points in the same direction throughout the channel and a large non-zero field is present across the absorber. Reconstruction effects have been considered and corrected for, with systematic uncertainties considered with this in mind. The corrected distributions of amplitude show the expected cooling effect for 6 and 10 mm nominal emittance beams, with equilibrium emittance maintained for 4 mm nominal emittance beams and heating observed in the 3 mm nominal emittance beams. An increase in muon density of approximately 20% is observed in the lowest amplitude bin for a 10 mm nominal emittance beam with the presence of an LH₂ absorber. For a 6 mm nominal emittance beam, the increase in muon density of the lowest amplitude bin is approximately 6% with the presence of the lithium hydride disk or 12% with the filled LH₂ vessel. The 4 mm nominal emittance beams show a change in muon density consistent with 0 in the lowest amplitude bin, while the 3 mm nominal emittance beams observe a decrease of between 15% and 35% in the lowest amplitude bin. This effect is seen in both reconstructed real data and simulated events.

The canonical angular momentum of particles within the cooling channel has been calculated, with the measured change across the absorber presented. The effect observed is of a similar size for all beams traversing a given absorber

material and is large given the canonical angular momentum of the beam upstream of the absorber. The largest effect is measured with the presence of the lithium hydride disk or filled LH₂ vessel, with almost no change shown where no absorber is present. This provides solid motivation for the use of distinct ‘flip’ and ‘solenoid’ mode cooling channels in order to manage this effect.

6.2 Future Work

Work is on-going to prepare this research for publication, with an internal review process within the MICE collaboration currently underway. This will address where inconsistencies in MC and data distributions arise and how they are dealt with.

A large amount of field-on data remains unanalysed, including data-sets spanning a range of momenta between 170 MeV/c and 240 MeV/c. An emittance change analysis for these higher momentum beams could be conducted. As the ionization cooling effect is less pronounced at these higher momenta, additional care may be required in curating the beam samples to provide an appropriate upstream sample which is well understood.

Part II

DUNE

The Deep Underground
Neutrino Experiment

Chapter 7

Background

7.1 Neutrino Discovery

The first suggestion of the neutrino arose from experiments on radiation in the early 20th century. Early investigations into beta radiation by James Chadwick showed that, unlike the sharp narrow energy spectra of alpha and gamma emissions, the energy of the emitted beta (β) decay electron formed a broad, continuous distribution [52]. This result directly challenged the fundamental principle of energy conservation, thought to be universally true.

Chadwick's anomalous result and its interpretation remained in contradiction to this core conservation principle of physics until 1930, when Wolfgang Pauli suggested a 'desperate remedy': the existence of a light, neutral fermion which he called the neutron [53]. This name was later revised to the neutrino (the Italian equivalent of "little neutral one") by Enrico Fermi following the discovery of the large neutral nucleon by Chadwick which Chadwick also named the neutron. It was also Fermi who, in 1933, first formalised a theory of beta decay which underpins the four-fermion interaction [54]:

$$n \rightarrow p + e^{-} + \bar{\nu}_e \quad (7.1)$$

with p and n as protons and neutrons, $e^{+/-}$ as positrons/electrons, and $\nu_e/\bar{\nu}_e$ as neutrinos/anti-neutrinos. Fermi's theory, a precursor to the theory of weak interaction, posited the spontaneous production of a $\bar{\nu}_e$ and e^{-} pair, similar to spontaneous photon emission in gamma decay. His theory also allowed for the absorption of neutrinos via inverse β decay as

$$\bar{\nu}_e + p \rightarrow e^{+} + n, \quad (7.2)$$

proposed as an experimental neutrino detection method by Wang Ganchang in 1942 [55]. It took another 14 years before this was realised by Cowan and Reines with the first direct observation of inverse beta-decay, utilising a nuclear




Figure 7.1: Measurements of the hadron production cross-section around the Z resonance, with curves indicating the predicted hadronic production cross-sections for two, three, and four neutrino species [59].

reactor source, water target, and liquid scintillators for detection [56]. Their initial planned experimental apparatus which included a nuclear explosion and a detector in free-fall did not make it into the final design.

It would not be long before the discovery of a second flavour of neutrino, produced by a team working at the Brookhaven National Laboratory [57]. A high-energy proton beam provided by the Brookhaven Alternating Gradient Synchrotron was collided with a fixed beryllium target to provide a pion source, with an aluminium spark chamber placed 70 ft downstream behind 13.5 m of steel shielding. Their experiment observed a distinct muon-associated neutrino flux, ν_μ , produced from pions decaying via

$$\pi^\pm \rightarrow \mu^\pm + \nu_\mu/\bar{\nu}_\mu \quad (7.3)$$

which interacted in the detector, producing muon events. With the discovery of the tau (τ) lepton in a series of experiments between 1974 and 1977 at SLAC [58], the existence of an accompanying third neutrino flavour was immediately implied. Indeed, at this time it was feared a huge number of light neutrino species were still to be discovered, as the standard theory would accommodate any number of fermion families, requiring only that the number of lepton and quark families are equal. Stanford Linear Collider (SLC) and Large Electron Positron Collider (LEP) measurements of the Z -boson resonance/decay width from 1989 onwards indicated the τ neutrino would be the third and final light neutrino, determining the number of light neutrino species to be 2.9840 ± 0.0082 through their contribution to the decay width [59]. However, direct observation of this neutrino would have to wait until the construction of the DONUT (Direct Observation of NU Tau) experiment, which measured 4 ν_τ interactions

in 1997 (with a background of 0.34 events) produced via charmed meson decays [60].

7.2 Neutrinos in the Standard Model

In the Standard Model (SM) description, neutrinos are neutral, massless leptons which accompany the charged leptons, e^\pm , μ^\pm , and τ^\pm . While the charged leptons can interact via the electromagnetic force, neutrino interactions are restricted solely to weak interaction, with their interactions prescribed by the helicity of the particle: -1 for ν , and $+1$ for $\bar{\nu}$, or equivalently, left-handed for neutrinos and right-handed for anti-neutrinos. The helicity of a particle is the projection of its spin, \vec{s} , on its direction of motion, \vec{p} . Due to the massless nature of the SM Dirac neutrino¹, the helicity of an SM neutrino is equivalent to its chirality and is a fixed property. For non-massless particles, the helicity is a constant of the motion, but is not Lorentz invariant, whilst the chirality *is* Lorentz invariant, but is *not* a constant of motion; a propagating massive left-handed fermion may evolve into a right-handed fermion or vice-versa.

It appears however that the SM description of a Dirac neutrino is insufficient to adequately describe all observed phenomena. Due to the handed chiral coupling of the weak interaction which interacts with/produces only L-chiral particles or R-chiral anti-particles, Dirac neutrinos cannot have mass without the addition of a right-handed sterile neutrino (and left-handed sterile anti-neutrino). An extension of the standard model exists which introduces Majorana neutrinos, where the left-handed neutrino couples with the right-handed anti-neutrino such that it may be its own anti-particle. This would require an extension to the Higgs field to include a triplet carrying weak hypercharge of ± 2 , or once again the addition of a sterile heavy neutrino - in both cases, this requires the introduction of Beyond Standard Model (BSM) physics.

¹A Dirac fermion is one with distinct particle and anti-particle states.




Figure 7.2: Projection of spin, \vec{s} , on the direction of particle motion, \vec{p} , for a particle, q , showing right- and left-handed helicities.

7.3 Neutrinos Beyond The Standard Model

7.3.1 Flavour Oscillations

Experimental observations of neutrino fluxes from solar and atmospheric sources, beginning with Ray Davis's Homestake experiment in the late 1960's [61], and confirmed by Kamiokande in 1989 [62], along with measurements from SAGE and GALLEX, have long provided evidence of BSM physics within the neutrino sector. Solar modelling of fusion processes in the sun suggested a discrepancy between the number of neutrinos being observed to those expected of approximately one-third. Solar and neutrino physicists made efforts to refine their theories for neutrino production and detection respectively in the hope of resolving the anomaly, but neither side found success [63].

The solution was to be found elsewhere: Pontecorvo had proposed a theory of neutrino oscillation as far back as 1957 [64], in which he postulated it was possible for ν and $\bar{\nu}$ mixing. Seeing the discovery of distinct electron and muon neutrino species, this work was expanded by Maki, Nakagawa and Sakata [65–67] to include a theoretical description of neutrino flavour oscillations in vacuum. To verify this theory, the Sudbury Neutrino Observatory (SNO) water Cherenkov detector was constructed, utilising a deuterium (D_2O) heavy water detector medium to provide charged- and neutral- weak current sensitivity via the processes

$$\nu_e + d \rightarrow e^- + p + p \quad (7.4)$$

and

$$\nu_i + d \rightarrow \nu_i + p + n, \quad (7.5)$$

respectively, in addition to the elastic scattering process

$$\nu_i + e^- \rightarrow \nu_i + e^-, \quad (7.6)$$

which is suppressed for ν_μ and ν_τ . SNO data on charged current (CC) interactions, sensitive only to electron neutrinos, showed the same deficit produced in prior experiments. However in contrast, neutral current (NC) interactions provided sensitivity to all neutrino fluxes, independent of flavour, and described a total ν flux exactly consistent with that predicted by solar models. This result provided the first direct evidence of neutrino flavour oscillations, suggesting ν_e were oscillating into other flavours in transit. Unfortunately this result stands in direct contradiction to the SM. For such oscillations to occur, neutrinos are required to have non-zero masses as the mechanism for flavour mixing necessitates that the neutrino mass eigenstates must be distinct from the flavour eigenstates and must be non-degenerate. This allows each flavour eigenstate, ν_e, ν_μ, ν_τ , to be expressed as a superposition of the mass

eigenstates, $|\nu_i\rangle$,

$$|\nu_\alpha\rangle = \sum_i^N U_{\alpha i} |\nu_i\rangle, \quad (7.7)$$

where $|\nu_\alpha\rangle$ are the resulting flavour eigenstates and $U_{\alpha i}$ are the elements of a unitary mixing matrix. Then, as neutrinos propagate, each mass state observes a separate phase advance, leading to interference which provokes flavour mixing.

Neutrino flavour oscillations have been well-characterised by the Pontecorvo-Maki-Nakagawa-Sakata matrix (PMNS matrix), a unitary mixing matrix, analogous to the Cabibbo-Kobayashi-Maskawa (CKM) mixing matrix from the quark sector, which describes the flavour change which occurs as neutrinos propagate from a source. In general, any unitary three-by-three matrix has nine degrees of freedom; however in the case of the PMNS matrix, it is possible to absorb 5 of those real parameters as relative phases between states. The PMNS matrix is usually parameterised using the three weak mixing angles, θ_{12} , θ_{23} , θ_{13} , and a CP-violating phase, δ_{CP} , appearing as

$$\underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & c_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix}}_{\text{Atmospheric}} \underbrace{\begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta_{CP}} & 0 & c_{13} \end{pmatrix}}_{\text{Reactor}} \underbrace{\begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\text{Solar}}, \quad (7.8)$$

where c_{ij} and s_{ij} are used to denote $\cos\theta_{ij}$ and $\sin\theta_{ij}$ respectively. Each sub-matrix shown within the PMNS matrix takes its primary measurements from a different experimental source. The left-hand matrix, composed of functions of θ_{23} , describes the oscillation of ν_μ to ν_τ and is measured from atmospheric oscillations. The center matrix, composed of functions of θ_{13} and δ_{CP} , describes ν_μ to ν_e and is measured from reactor experiments. The right-hand matrix is composed from functions of θ_{12} only and is measured from solar oscillations of ν_e to ν_μ and ν_e to ν_τ . When these mixing matrix terms are propagated through the evolution of a time-dependant Schrodinger equation solution, an oscillation probability of

$$P(\alpha \rightarrow \beta) = \delta_{\alpha\beta} - 4 \sum_{i>j} \text{Re}(U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^*) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) + 2 \sum_{i>j} \text{Im}(U_{\alpha i}^* U_{\alpha j} U_{\beta i} U_{\beta j}^*) \sin \left(\frac{\Delta m_{ij}^2 L}{2E} \right) \quad (7.9)$$

is obtained [68], with Δm_{ij}^2 the mass difference between states i and j squared, L the length traversed in propagation, and E the neutrino energy. These Δm_{ij}^2 mass differences, and hence the absolute neutrino masses, are then necessary for a non-zero oscillation probability.

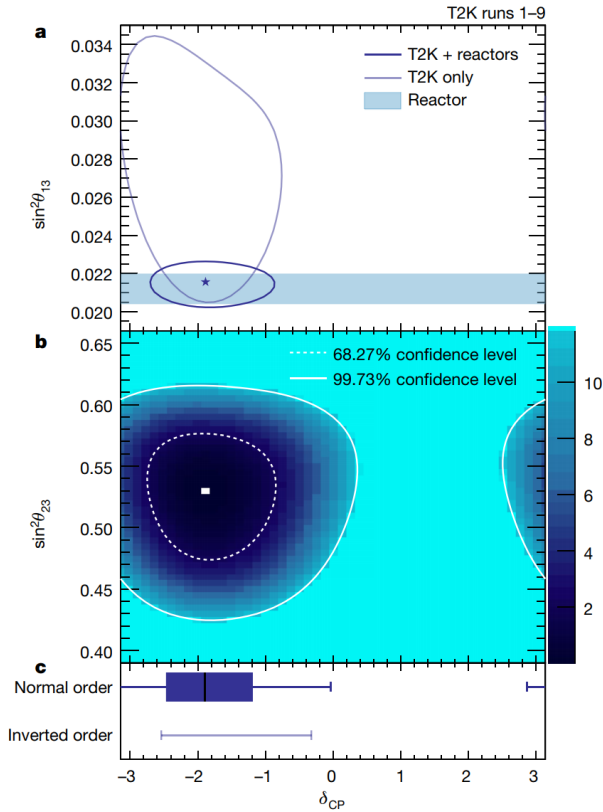


Figure 7.3: The most recent constraints on the PMNS oscillation parameters δ_{CP} , $\sin^2 \theta_{13}$, and $\sin^2 \theta_{23}$ from the T2K collaboration. **a**: Two-dimensional confidence intervals at the 68.27% confidence level for δ_{CP} versus $\sin^2 \theta_{13}$ in the preferred normal ordering hierarchy. The star shows the best-fit point of the T2K+reactors fit in the preferred normal mass ordering. **b**: Two-dimensional confidence intervals at the 68.27% confidence level for δ_{CP} versus $\sin^2 \theta_{23}$ from the T2K+reactors fit in the preferred normal ordering hierarchy. The colour scale gives $-2 \times$ the log likelihood for each parameter value. **c**: One-dimensional confidence intervals at the 68.27% (box) and 99.73% (whiskers) level on δ_{CP} from the T2K+reactors fit for normal and inverted hierarchies. Only the 99.73% interval is shown for inverted hierarchy [69].

CP Violation

It is within the PMNS mixing matrix that an opportunity for charge-parity (CP) violation arises, contained within the $\sin \theta_{13} e^{-i\delta_{CP}}$ terms. Each carries a complex phase δ_{CP} of undetermined value with the potential to induce particle anti-particle differences in the oscillation probabilities between flavour states. Such differences are already observed in the quark sector, but at significantly lower levels than required to fully explain the matter-antimatter imbalance which is present in the observable universe.

Measurements from the Tokai to Kamioka (T2K) and NO ν A collaborations provide the current best constraints on δ_{CP} , with recent T2K results [69] excluding values of $\delta_{CP} = 0$ or π at the 95% confidence level, as shown in




Figure 7.4: The effect of different δ_{CP} values on $\nu_\mu \rightarrow \nu_e$ (left) and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ (right) appearance probabilities as a function of neutrino energy, for $\delta_{CP} = -\pi/2$ (blue), 0 (red), and $\pi/2$ (green). Probabilities are given at a baseline of 1300 km using the normal mass hierarchy. Taken from [72].

figure 7.3. These two values define the points of CP conservation and so a measurement differing from zero or π would represent the discovery of CP violation in the leptonic sector, providing a possible explanation for the matter-antimatter asymmetry in the universe. These T2K results conflict somewhat with recent NO ν A results [70], with both experiments disfavouring the other's region of best fit.² As such, more accurate measurements of δ_{CP} are necessary to provide definitive results at the 5σ level, prompting the construction of next generation neutrino oscillation experiments such as DUNE and HK. The impact of different values of δ_{CP} on a DUNE-like $\nu_\mu \rightarrow \nu_e$ oscillation experiment is shown in figure 7.4.

MSW Matter effect

The oscillation probabilities shown in equation 7.9 describe those occurring within a vacuum. For a neutrino propagating through material, its oscillation probabilities are modified by the presence of electrons, each presenting an opportunity for electro-weak interactions. Whilst μ and τ neutrinos are restricted to neutral current interactions, due to the abundance of electrons in normal matter, electron neutrinos may interact via both neutral and charged current interactions. Such CC coherent forward scattering provides an excess Hamiltonian potential V_e of strength

$$V_e = \pm\sqrt{2}G_F N_e, \quad (7.10)$$

²Some work has suggested this discrepancy may be explained by non-standard neutrino interactions [71].

where G_F is the Fermi coupling constant, N_e is the number density of electrons in material, and the sign is determined by the (anti)matter type. This introduces a relative phase-shift between neutrino flavours, modifying the mass eigenstates, which should be observed in the measured mixing angles of any sufficiently long-baseline neutrino oscillation experiment in matter. Similar Hamiltonian perturbations arising from NC interactions produce a phase-shift common to all flavours and which are hence irrelevant to the oscillation probability. This mass-dependent oscillation phenomenon is known as the MSW effect, after Mikheyev, Smirnov and Wolfenstein, or the matter effect.

Mass Hierarchy

The dependence of neutrino flavour oscillations on the square of the mass differences between neutrino masses, Δm_{ij}^2 , means that whilst these mass-squared differences are well-measured from oscillation experiments, the absolute masses and sign of each mass splitting are not and *cannot* be measured using the vacuum formulation of neutrino oscillations. Recent combined results from PDG [12] for the mass-squared splittings give a positive-definite³ best fit value for the smallest splitting of $\Delta m_{21}^2 = 7.53 \times 10^{-5} \text{ eV}^2$. Ambiguity over the sign of the two remaining mass-squared differences exists, but current results place $|\Delta m_{31}^2| \approx |\Delta m_{32}^2| = 2.5 \times 10^{-3} \text{ eV}^2$. This ambiguity constitutes the neutrino mass hierarchy problem, and presents two possible orderings: "Normal" ordering, where $m_1 < m_2 < m_3$, and "inverted" ordering, with $m_3 < m_1 < m_2$. This naming reflects the mass ordering of the charged leptons, as ν_1 has the largest ν_e component and ν_3 the smallest. The two cases are shown in Fig. 7.5.

While the sign of the mass-squared splittings is inaccessible from observations of vacuum oscillations, the MSW matter effect introduces an additional dependence on the sign of Δm^2 . This allows sufficiently long-baseline neutrino experiments such as NO ν A to be sensitive to the mass hierarchy, with recent results weakly favouring normal ordering [70].

Mass hierarchy determination can also be based on constraints found elsewhere: Precision measurements of the ν_e mass are underway, with the KATRIN experiment's first data run setting an upper limit of 1.1 eV at the 90% confidence level and further data-taking ongoing. Beyond this, neutrinoless double beta decay ($0\nu\beta\beta$) experiments which assume neutrinos are Majorana fermions conduct searches for this phenomenon, with an expected rate proportional to the square of the effective neutrino mass, $m_{ee} = ||U_{e1}|^2 m_1 + |U_{e2}|^2 m_2 e^{2\alpha i} + |U_{e3}|^2 m_3 e^{2\beta i}|$, and can thereby set a confidence limit on the effective mass through non-observation (given neutrinos *are* Majorana). Similarly, cosmological modelling

³ $m_2^2 > m_1^2$

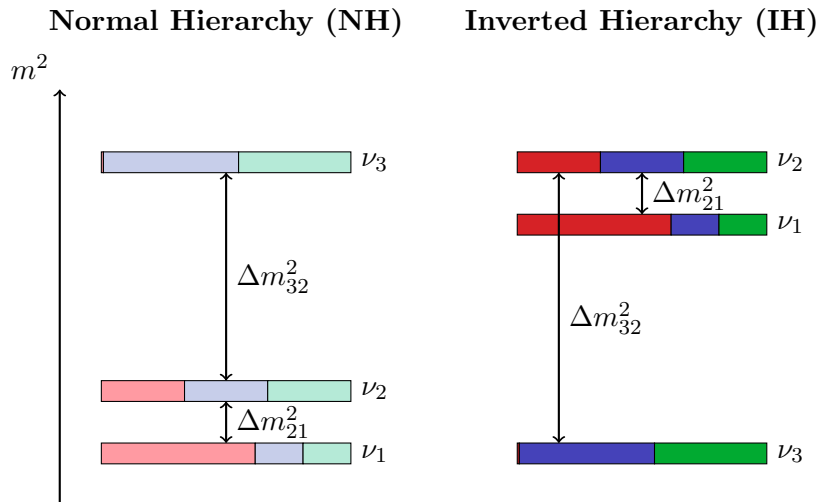


Figure 7.5: A visual representation of the two possible neutrino mass orderings: (left) the normal hierarchy (NH) and (right) inverted hierarchy (IH). The relative proportions of red (ν_e), blue (ν_μ), and green (ν_τ) segments in each mass eigenstate box corresponds to the relative probability of finding each flavour eigenstate in that mass eigenstate.

of the mass dispersion in the universe based on a cosmological neutrino mass $m_{cosmo} = \sum m_i$ may set limits which exclude either hierarchy.

7.4 Current Neutrino Sector Parameters

A combined analysis of the results from many recent neutrino experiments can, if well-understood, provide the best picture of our current understanding of the neutrino sector and the parameters which govern it. Figure 7.6 shows an example of the confidence intervals for fits of $\sin^2 \theta_{23}$ and $|\Delta m_{31}^2|$ from a number of experiments, from which a combined result can be assessed. The results of

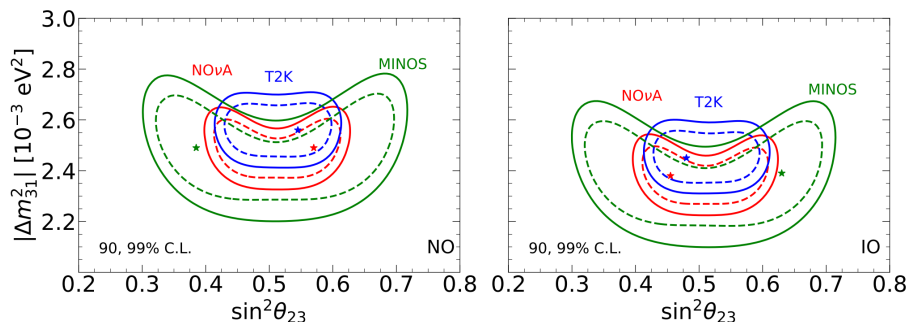


Figure 7.6: Two-dimensional confidence intervals at the 90% and 99% confidence level for $\sin^2 \theta_{23}$ versus $|\Delta m_{31}^2|$ for normal ordering (left) and inverted ordering (right). Each shows confidence intervals from T2K (blue), NO ν A (red), and MINOS (green) data, with best fit values indicated by stars [73].

Oscillation Parameter	Best Fit $\pm 1\sigma$	3σ range
$\Delta m_{21}^2 [10^{-5} \text{ eV}^2]$	$7.50_{-0.20}^{+0.22}$	6.94 – 8.14
$ \Delta m_{31}^2 [10^{-3} \text{ eV}^2]$ (NH)	$2.55_{-0.03}^{+0.02}$	2.47 – 2.63
$ \Delta m_{31}^2 [10^{-3} \text{ eV}^2]$ (IH)	$2.45_{-0.03}^{+0.02}$	2.37 – 2.53
$\sin^2 \theta_{12}/10^{-1}$	3.18 ± 0.16	2.71 – 3.69
$\sin^2 \theta_{23}/10^{-1}$ (NH)	5.74 ± 0.14	4.34 – 6.10
$\sin^2 \theta_{23}/10^{-1}$ (IH)	$5.78_{-0.17}^{+0.10}$	4.33 – 6.08
$\sin^2 \theta_{13}/10^{-2}$ (NH)	$2.200_{-0.062}^{+0.069}$	2.000 – 2.405
$\sin^2 \theta_{13}/10^{-2}$ (IH)	$2.225_{-0.070}^{+0.064}$	2.018 – 2.424
δ_{CP}/π (NH)	$1.08_{-0.12}^{+0.13}$	0.71 – 1.99
δ_{CP}/π (IH)	$1.58_{-0.16}^{+0.15}$	1.11 – 1.96

Table 7.1: Summary of neutrino oscillation parameters determined from a recent global analysis, showing fit results for normal hierarchy (NH) and inverted hierarchy (IH) [73].

a recent global reassessment of neutrino oscillation parameters [73] are shown in table 7.1.

While T2K and NO ν A are still statistics-limited in their uncertainties and will continue to take data in the short term, larger steps forward are planned in the near future with construction of two next-generation neutrino oscillation experiments underway. The first of these is Hyper-Kamiokande (HyperK) [74], a large-scale water-Cherenkov neutrino detector with a tank containing a planned 260,000 metric tons of ultrapure water - an increased fiducial volume approximately 10 times larger than that of Super-Kamiokande. Construction of the detector will take place at the host site in the Kamioka mine in Japan, with plans to locate the tank at a depth of 650 m underground.

The second of these planned next-generation experiments is the Deep Underground Neutrino Experiment (DUNE), which will be discussed in greater detail in the following chapter.

Chapter 8

The DUNE Experiment

The Deep Underground Neutrino Experiment (DUNE) is an international next generation long-baseline neutrino oscillation experiment, with facilities based at two U.S. sites. The Fermi National Accelerator Laboratory (Fermilab), based outside of Chicago, Illinois, will operate the proton accelerator which drives the neutrino beam and house the near detector (ND) complex, whilst the Sanford Underground Research Facility (SURF) in Lead, South Dakota, will host the large underground far detector (FD) complex. The arrangement of these facilities is shown in figure 8.1. Construction of the facilities involved is currently underway, with plans to utilise an upgrade of the Fermilab accelerator complex to provide what will be the world’s most intense neutrino beam.

The detector complexes themselves will utilise novel liquid argon time projection chamber (LArTPC) technology to provide full event reconstruction, with additional complementary detector systems in place at the near detector site. The near and far detector systems are described in detail in sections 8.3 and 8.4. The Long-Baseline Neutrino Facility (LBNF) will be providing infrastructure for the near and far complexes, including cavern excavations at SURF and civil construction of the near detector systems, in addition to preparing and operating the proton beamline. The DUNE project will provide the four 10 kt fiducial mass far detector modules.




Figure 8.1: The planned facilities and baseline of the DUNE and LBNF projects, depicting the near site at Fermilab and the far site at SURF. [72]

8.1 DUNE Physics Goals

The DUNE collaboration has defined a clear set of primary and ancillary physics goals. The primary physics goals focus on fundamental questions in neutrino and astroparticle physics, and are as follows [72]:

- Comprehensive measurements of $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation parameters, including:
 - Measurement of the charge-parity (CP) violating phase, δ_{CP} , at a sensitivity of better than three standard deviations (3σ) over more than 75% of the possible range of values.
 - Measurement of the θ_{23} mixing angle and determination of the octant in which the angle lies.
 - Determination of the neutrino mass hierarchy, i.e. the sign of Δm_{31}^2 .
- Search for evidence of proton decay and other baryon number violating processes.
- Detect and measure intra-galactic core-collapse supernovae ν_e flux occurring within DUNE's lifetime.

The ancillary science goals will supplement the key goals driving the experiment, and provide a rich complementary science program. They are:

- Search for BSM neutrino oscillation measurements, including non-standard neutrino interactions, sterile neutrino signatures, Lorentz invariance violation, CPT violation.
- Measurement of tau neutrino appearance.
- Measurement of atmospheric neutrino oscillations.
- Near detector-led studies of neutrino interaction physics, including measurements of interaction cross-sections, nuclear effects such as neutrino final-state interactions, and of nucleon structure.
- Measurement of the weak mixing angle via $\sin^2 \theta_W$.
- Search for dark matter signatures.

8.2 The DUNE Beamline

To produce and direct the neutrino beam towards the near and far detector sites, several key facilities are required along the beamline. A side-on view of the proposed LBNF beamline up to the near detector hall is shown in




Figure 8.2: The Neutrino beamline and DUNE near detector hall at Fermilab, IL. [72]

figure 8.2. The proton beam which drives the LBNF neutrino beam will be provided by the Fermilab Main Injector proton accelerator. This facility will benefit from the Proton Improvement Plan II (PIP-II) upgrade, a leading-edge superconducting linear proton accelerator under construction at Fermilab which will deliver between 1.0 and 1.2 MW of proton beam power, achieving an estimated 1.1×10^{21} protons-on-target per year in a wide energy band of 60 GeV to 120 GeV [72]. Additionally, a second upgrade of the accelerator complex aims to provide an increase in beam power to 2.4 MW by 2035.

After extraction from the main injector ring, the beamline will rise over an embankment and be directed down towards the target hall. The embankment allows placement of the target hall complex above ground, limiting the excavation required and allowing easy access for target installation and maintenance. Proton collisions at the target will then produce a secondary charged particle beam, consisting mainly of pions and kaons, which are directed towards a helium-filled decay pipe of length 194 m. This structure allows the particles additional travel time to decay into neutrinos prior to reaching the ND hall. Following the decay pipe, the beam passes through the absorber hall and muon alcove where hadron and muons are absorbed by an extended section of aluminium, steel and concrete shielding [75]. A final distance of 304 m of rock will separate the absorber hall from the near detector hall.

8.2.1 The Target and Horns

The proposed target utilises a 1.8 m long, 16 mm diameter graphite rod fitted with a liquid helium cooling jacket to manage the extreme energy deposition from the megawatt proton beam. To arrive at this target design, a balance was

made between improved particle production from extended interaction length and increased scattering and absorption of produced particles in graphite downstream of their production point. Additional considerations included radiation hardness of the material (radiological damage), robustness to thermal damage, and ease of target replacement. Further optimisations are currently being made to maximise FD neutrino flux.

A set of staggered magnetic focusing horns surround, or are positioned downstream from, the target, acting as focusing devices for secondary particles. Each horn is a magnetic lens which refocuses particles, primarily charged pions and kaons, toward the decay pipe by producing a pulsed magnetic field contained within the inert gas volume inside the horn. Resulting from LBNF optimizations, the design will utilise 3 horns, referred to as Horns A, B and C. Horn A will enclose the target and cooling jacket, while the upstream ends of horns B and C will be situated 2.956 m and 17.806 m respectively from the upstream end of the target. Their proposed layout within the target hall is shown in figure 8.3. The design of each horn's shape and their placement was optimized using a genetic machine learning algorithm to provide greater sensitivity to signals of CP violation by increasing FD neutrino flux and reducing the wrong-sign neutrino background [75]. The horn systems are powered by an electrical current supplied via aluminium striplines connected to each horn conductor, with the stripline design setting an upper limit on the operating current range of each horn. The horns can be operated with forward horn current (FHC) or reverse horn current (RHC), focusing positively or negatively charged hadrons to produce neutrino or anti-neutrino dominated beams respectively.




Figure 8.3: Planned design and layout of the magnetic focusing horns in the target hall. The beam here goes from left to right. [75]

8.3 Near Detector Complex

Located 304 m downstream of the end of the decay pipe, the near detector hall will house three complementary subdetectors which form the complete near detector itself, namely the Liquid Argon TPC (ND-LAr), the Gaseous Argon TPC (ND-GAr), and the System for on-Axis Neutrino Detection (SAND).¹ The three subdetectors are staggered along the direction of the neutrino beam, as shown in figure 8.4.

The near detector plays an essential role in understanding the neutrino beam DUNE will utilise for any long-baseline oscillation analysis. To provide a control for the far detector measurements and achieve the precision the experiment requires, several key tasks will be carried out by the near detector: measuring and monitoring the beam, constraining systematic uncertainties and providing essential input towards the neutrino interaction model. The ND achieves this by measuring unoscillated $\nu_\mu/\bar{\nu}_\mu$ and $\nu_e/\bar{\nu}_e$ energy spectra.

As the reconstructed energy spectrum is an unresolved convolution of cross-section, flux, and energy response, the ND has to independently constrain each component such that it can be well-modelled in both the near and far detectors. Models of the detectors, beam, and interactions are used to account for differences between the detectors, such as their event acceptance, and the underlying fluxes they sample. Because these requirements fall on the near detector, it must outperform the FD, providing a higher detection efficiency, superior identification of charged and neutral particles, and better energy reconstruction. Additionally, to fulfil these specifications, both the ND-LAr and ND-GAr subdetectors are designed to move to take data off the beam axis, while SAND stays on-axis as a beam monitor to ensure consistent operation of the neutrino beam. This capability is referred to as the DUNE Precision Reaction-Independent Spectrum Measurement (DUNE-PRISM).

8.3.1 ND-LAr

The first of the ND subdetectors which observes the beam, and the one which shares the most similarity with the FD, is the modularised LArTPC, ND-LAr. Due to the intense neutrino flux and high event rate expected at such close proximity to the neutrino source, ND-LAr will operate in a high-multiplicity environment with an expected $\mathcal{O}(50)$ events per $10 - \mu\text{s}$ proton spill. To accommodate these conditions, the ND-LAr will be constructed out of a matrix of modularised optically isolated LArTPCs, each based on ArgonCube technology, occupying $5 \text{ m} \times 7 \text{ m} \times 3 \text{ m}$ and containing 67 t of fiducial mass

¹ND-GAr has previously been referred to as the Multi-Purpose Detector or MPD. The trio of detectors is sometimes affectionately referred to LArgon, GArgon and NArgon, short for 'Liquid Argon', 'Gaseous Argon' and 'Not Argon' respectively.




Figure 8.4: Illustrated layout of the subdetectors within the near detector hall. The beam points right to left here, passing in order through ND-LAr, ND-GAr, then SAND [76].

liquid argon. This mass will consist of 35 smaller modules that allow for isolated reconstruction of each neutrino-interaction event, with the complete fiducial volume large enough to provide the required hadronic shower containment and statistics for a near-far analysis.

ND-LAr shares the same target nucleus as the FD and relies on the same fundamental detection principles employed in the FD (-Single Phase), but adapted to accommodate the much larger intensity expected. In doing so, the sensitivity to nuclear effects and detector-driven systematic uncertainties in a near-far oscillation analysis is largely reduced. The detector is expected to collect 1×10^8 ν_μ -CC events per year on axis.

8.3.2 ND-GAr

Positioned directly downstream of the ND-LAr, the ND-GAr's design utilises a high-pressure gaseous argon TPC (HPgTPC) surrounded by an electromagnetic calorimeter (ECAL) and enclosed in a 0.5 T magnetic field, in addition to an iron and plastic scintillator muon system. Based closely on the design of the ALICE TPC but operating at a pressure of 10 atm, the HPgTPC will be capable of reconstructing particles at lower thresholds than is achievable in the LArTPC designs and is expected to provide improved understanding of the relationship between true and reconstructed energy.

Due to the smaller relative size of the ND-LAr subdetector than the FD modules, additional systems are required to reconstruct muons above an energy of ~ 0.7 GeV/c as they are not reliably contained in the subdetector's fiducial volume. Measuring both the charge and momentum of muons will be critical to

reconstructing the full neutrino energy. To achieve this, the ND-GAr detector has to track, identify and momentum-analyse muons which exit the ND-LAr. Additionally, the muon system, while still in a very preliminary stage of design, will provide additional particle identification to separate muons and pions traversing the ECAL. This is crucial in determining the wrong-sign background in the neutrino and anti-neutrino beams.

8.3.3 SAND

SAND is novel amongst the ND complex as the only non-argon based detector system. The current design is based on a reuse of the magnet and calorimeter from the KLOE experiment, with an instrumented target and tracking system occupying the inner magnet and ECAL volume. An initial proposal design used an active target constructed from 1 cm^3 plastic scintillator cubes, each optically isolated, forming a central 3D scintillating tracker. This inner tracker would then be surrounded by low-density tracking chambers to measure the charge and momentum of outgoing particles. The tracking chambers would use TPCs, straw tube trackers (STT), or a mixture of both. This TPC design followed T2K's ND280 detector design, utilising a low-density gas tracker instrumented with Micromegas readout modules. The reference SAND design has since been updated to instead fill the magnetic volume with orthogonal XY planes of STT, interspersed with thin layers of additional target material such as carbon.

As the current DUNE-PRISM measurement program calls for $\sim 50\%$ of ND-LAr and ND-GAr beamtime to be spent taking measurements off-axis, SAND is instrumental in ensuring the beam remains stable, as variations or distortions of the beam are most easily seen and diagnosed from measurements on the beam axis.

8.4 Far Detector Complex

The far detector complex will be located in a large excavated underground cavern system at SURF, South Dakota. The complex will house four LArTPC detector modules, each contained in its own identically sized cryostat and each with a fiducial mass of at least 10 kt. The cavern system, shown in figure 8.5, will be situated approximately 1.5 km underground with two larger caverns housing the four FD modules and the smaller central cavern providing space for data acquisition (DAQ) and cryogenics systems. The first two detector modules will be installed by 2028, with the remaining two installed as rapidly as the funding allows, estimated to be within the following two years.

The current project proposal utilises two developing LArTPC technologies for the first three modules, single-phase (SP) and dual-phase (DP), with an




Figure 8.5: Layout of the underground caverns housing the DUNE FD and cryogenics at SURF, South Dakota. The arrow shows the direction of the neutrino beam. The long vertical structure on the right shows the refurbished Ross Shaft, providing surface access [77].

SP module planned first. In SP, all detector elements inside the cryostat are immersed in liquid Argon. Ionisation electrons, liberated by charged particles traversing the LAr volume, drift horizontally under a uniform \vec{E} field towards an instrumented anode plane where they are read out. The SP module design (DUNE-FD SP LArTPC) has a total mass of 17.5 kt, with an active fiducial mass of 10 kt. Alternating anode and cathode walls divide this mass into separate drift volumes, with three anode and two cathode planes providing a maximum drift length of 3.5 m. The 58.2 m length anode planes run the length of the module and will be constructed from 6 m high by 2.3 m wide anode plane assemblies (APAs) stacked two high and 25 wide for a total of 150 per module. Similarly, the cathode planes are each formed from 150 stacked cathode plane assemblies (CPAs) spanning the same total length and height.

Each APA consists of three layers of active wires on each side, with the outer two recording bipolar induction signals as electrons pass through them and the final layer collecting the drifting electrons. The collected ionization pattern on the anode wires provides event reconstruction in the coordinates perpendicular to the drift direction, as shown in figure 8.6, whilst the angle of each wire layer is chosen to reduce reconstruction ambiguities.

DP operation is very similar to SP, but ionization charges instead drift vertically upward and are extracted into a layer of argon gas above the liquid. A large electron multiplier amplifies the charge in the gaseous phase before readout

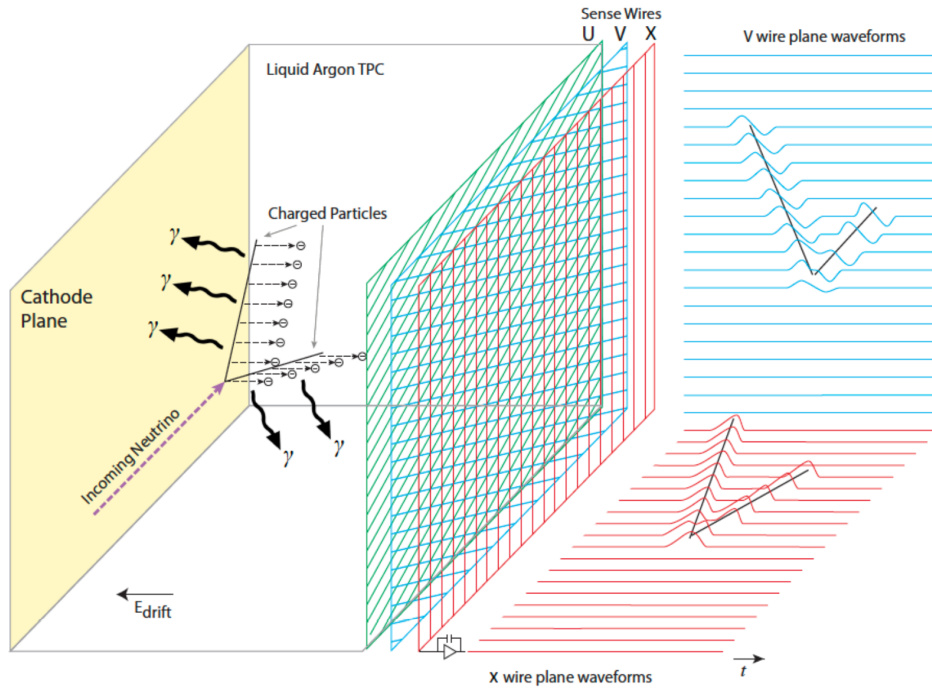


Figure 8.6: Operating principle of a single-phase LArTPC module, showing electron drift toward an APA on the left and the corresponding signal waveform reconstruction on the right [77].

on a horizontal 2D anode plane consisting of two sets of gold-plated copper strips. The gain achieved in the gas phase provides a better signal-to-noise ratio, offering a lower threshold on the charge required to form a reconstructed particle track and increasing the possible drift length. The detector design also brings new challenges, requiring over three times higher cathode voltages, with additional complexity managing space charge build-up and Argon boil-off at the liquid-gas boundary. Testing of the DP design is currently underway in ProtoDUNE-DP, located at the CERN Neutrino Platform Facility.

In both SP and DP, the drift volume is surrounded by a field cage which ensures electric field uniformity to the 1% level, nominally of 50 kV m^{-1} . Additionally, both technologies exploit the fast scintillation properties of argon, collecting UV scintillation light, shifted from 126.8 nm to visible, to provide precision timing for each event. This photon detection indicates when the electron drift began and provides reconstruction of the drift coordinate by comparison with the ionization signal measurement time. The photon collection systems will be installed on the frame of each SP anode plane and the bottom of the DP TPC, utilising X-Arapuca bars and PMTs respectively.

Whilst the installation timetable and order of FD modules is dependent on forthcoming results from prototype detectors, the assumption at the time of writing is that the first two modules will utilise SP technology, with a third

DP module installed shortly after. Several studies are currently underway for a more advanced detector design to occupy the fourth module.

Chapter 9

DUNE-PRISM

9.1 Off-Axis Neutrino Beams

The relativistic kinematics of a two-body meson decay define a well-known relationship between the meson parent energy, the neutrino energy, and its angle of emission with respect to the parent meson direction, as shown for pions in figure 9.1. Previous experiments such as T2K and NO ν A have employed this to produce more monochromatic neutrino beams, stationing their detectors at 44 mrad (2.52°) and 15 mrad (0.86°) off-axis, respectively.




Figure 9.1: Neutrino energy spectrum as a function of parent pion energy at different off-axis angles away from the pion momentum, θ .

To further utilise this effect beyond static off-axis detector systems, the construction of the nuPRISM intermediate water cherenkov detector in the J-PARC neutrino beamline was proposed in 2014 [78]. The proposed nuPRISM

physics program would take multiple measurements continuously spanning the off-axis range 1° to 4° with a movable detector, linearly combining measured spectra to produce Gaussian-like beams at specific energies. In doing so, precision measurements of interaction cross-sections could be made using pseudo-monochromatic beams, better informing other water cherenkov physics analyses. Linearly combining off-axis measurements to produce a new better-understood flux in this way underpins a PRISM analysis.

In experiments employing liquid argon as the active target material, difficulties remain in controlling the systematic uncertainties due to neutrino interaction modelling. Current models of the relationship between observable final state particles from neutrino interactions on argon and the incident neutrino energy remain imprecise, with significant deviations shown between some modern neutrino event generators in the low energy range. As a near-far analysis observes differing spectra at each detector due to oscillations and differences in detector geometries, missing energy from undetected or misidentified particles will not cancel in a near/far flux ratio. Contributors towards this missing energy include neutrons produced in neutrino interactions which induce variable secondary interactions where reconstruction of final states can poorly correlate to the initial kinetic energy of the primaries. Additionally, the relative energy carried by neutrons at the near- and far-detectors is expected to differ between neutrino and anti-neutrino interactions, further biasing any δ_{CP} measurement.

As any feed-down in reconstructed neutrino energy can bias measurements of oscillation parameters, addressing these uncertainties can significantly improve the sensitivity of oscillation parameter measurements. The DUNE Precision Reaction-Independent Spectrum Measurement (DUNE-PRISM) program exploits the off-axis kinematics by performing measurements continuously spanning the range of off-axis angles using a movable detector. As in nuPRISM, linear combinations of these measured fluxes can then be incorporated into the near-far analysis to control and reduce interaction model uncertainties using a data-driven methodology. By providing linear combinations which form pseudo-monochromatic beams *and* a ND flux which is very similar to the expected oscillated FD flux, oscillation parameters can be extracted largely without relying on an interaction model.

9.1.1 Off-Axis Flux at DUNE

A PRISM analysis requires measurement of a sufficiently large and distinct set of neutrino fluxes, each of which provides additional flexibility in any linear flux combination. To this end, DUNE-PRISM has considered the neutrino fluxes available by moving the near detector to off-axis distances up to 37.5 m;

Stop	Run Duration	ND-LAr				ND-GAr
		All Int.	Selected			All Int.
		$N_{\nu_{\mu}CC}$	N_{sel}	WSB	NC	$N_{\nu_{\mu}CC}$
On axis	14 wks.	21.6M	10.1M	0.5%	1.3%	580,000
4 m off axis	12 dys.	2.3M	1.2M	0.3%	1.0%	61,000
8 m off axis	12 dys.	1.3M	670,000	0.5%	0.9%	35,000
12 m off axis	12 dys.	650,000	330,000	0.8%	0.7%	17,000
16 m off axis	12 dys.	370,000	190,000	1.1%	0.7%	10,000
20 m off axis	12 dys.	230,000	120,000	1.3%	0.7%	6,200
24 m off axis	12 dys.	150,000	75,000	1.8%	0.7%	4,100
28 m off axis	12 dys.	110,000	50,000	2.1%	0.8%	2,900
30.5 m off axis	12 dys.	87,000	39,000	2.3%	0.7%	2,300

Table 9.1: A sample DUNE-PRISM run plan, outlining the yearly distribution of time spent at various on- and off-axis positions. The number of ν_{μ} charged current interactions ($N_{\nu_{\mu}CC}$) in both ND-LAr and ND-GAr detectors is shown, along with the number passing ND event selection (N_{sel}), the fraction of wrong-sign background (WSB), and fraction of neutral current (NC) events [76].

the off-axis muon neutrino flux between -1.5 m and 37.5 m is shown in figure 9.2. As a result of the relativistic kinematics, the highest energy peak of ν_{μ} flux is found when the detector is on-axis, with larger distances off-axis observing lower ν_{μ} energies. This restricts the flexibility of the set of muon fluxes at energies above the on-axis energy peak, limiting the ability to construct flux features at such energies. While a slight increase in muons in the high energy tail of the flux distribution is present around 5° off-axis, analysis has shown such tail populations provide limited additional utility. Furthermore, as the detector approaches further off-axis positions, the neutrino flux distributions there become closer in shape and structure, whilst providing lower flux in total. Consequently, ever-increasing the off-axis detector range returns ever-reducing useful information for the analysis. The impact this limit has on the analysis informs the maximum off-axis positions proposed for such a program.

The latest DUNE-PRISM measurement program plans to take off-axis detector measurements with a maximal excursion of 32.5 m off-axis, though this may be subject to change. The current run plan spans -1.5 m to 32.5 m over 9 detector stop positions, assuming a 4 m wide fiducial volume. The flux achieved at each off-axis detector stop is shown in figure 9.3. Whilst the nominal ND-LAr design operates with a 6 m wide, 3 m deep, 2 m high fiducial volume, this measurement plan assumes only a 4 m wide fiducial volume, providing room for geometrically large event reconstruction, such as during hadronisation events. Additional corrections for geometry-dependent event rejection are still required to ensure fair comparison of ND and FD fluxes, with studies on-going.




Figure 9.2: FHC muon neutrino flux through each off-axis near detector 50 cm flux window. The colour axis is normalised to neutrinos per GeV per cm^2 per POT.




Figure 9.3: FHC muon neutrino flux through each 4 m wide fiducial volume detector stop.

9.2 Flux Matching

The primary objective of DUNE-PRISM is to replicate useful spectral distributions of neutrino flux through the linear combination of ND events, prior to the onset of, potentially CP-violating, flavour mixing. With careful choice of target spectra, linearly combining ND measurements can effectively calibrate the relationship between observed and true neutrino energies. In addition, neutrino spectra expected at the FD in the presence of neutrino oscillations can be estimated by forming linear combinations of neutrino fluxes at the ND. This approach can yield significant de-coupling of flux and neutrino interaction uncertainties, whilst like-uncertainties between detectors, e.g. from missing energy, will cancel.

To do this, the range of ND flux measurements is separated into discrete flux windows, encompassing all neutrinos passing through a detector segment of a given width at the upstream face of the detector, divided by the number of off-axis bins this segment accounts for. This averaging is used to more easily equate regularisation parameters between solutions with differing flux-window widths and allows combinations of varying widths across the off-axis range, if required. Each window can be smaller than the fiducial width, and is selected to balance maximising the number of neutrino events per window, thereby reducing statistical fluctuations, against maximising the number of flux windows to provide additional flux matching capability. Rather than computationally expensive fitting procedures, a linear algebra approach can be used to find an exact solution which matches the resulting linear combination of ND flux windows with the oscillated FD flux. Defining the flux matrix describing the off-axis ND fluxes (shown in Fig. 9.2) as \mathbf{N} , and the vector of oscillated FD flux bins as \bar{F} , the simplest method to produce a set of coefficients is obtained by solving the linear least squares (LLS) expression:

$$\|\mathbf{N}\bar{c} - \bar{F}\| = 0 \quad (9.1)$$

where $\mathbf{N} \in \mathcal{R}^{i \times j}$, $\bar{F} \in \mathcal{R}^i$, $\bar{c} \in \mathcal{R}^j$, with \bar{c} the vector of coefficients to be identified, and $\|\cdot\|$ the ℓ_2 (Euclidean) norm. Here i represents the number of energy bins and j represents the number of bins in off-axis displacement. We can then solve for \bar{c} with

$$\bar{c} = \mathbf{N}^{-1}\bar{F}, \quad (9.2)$$

given \mathbf{N} is invertible. In the case where the number of energy bins, i , does not match the number of off-axis flux windows, j , then \mathbf{N} is a non-invertible non-square matrix and \mathbf{N}^{-1} is instead the left-inverse of \mathbf{N} .

This solution minimises the norm of the error residual, thereby providing a good match between fluxes, but generally produces coefficients with large

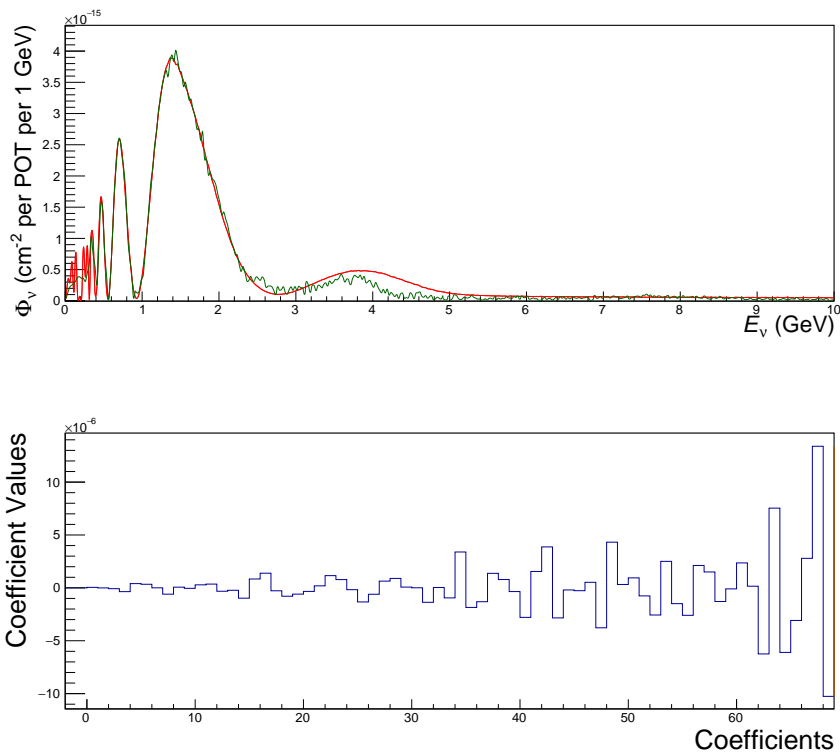


Figure 9.4: (Top) Flux matching between an oscillated FD flux (red) and the linear combination of ND fluxes (green) obtained by the linear least squares method shown in eqn. 9.2 and (bottom) their corresponding coefficients. The coefficients correspond to the weighting of each 50 cm flux window, with flux windows at larger distance off-axis having larger coefficient number.

variations between neighbouring values (fig 9.4), inherently susceptible to large errors with a reliance on large cancellations between similar flux windows. This is a known tendency of LLS solutions which solve such inverse problems, amplifying noise in the measurements¹.

In practice, $i \neq j$ in most cases (i.e. the solution is ill-posed), and so many possible solutions exist which provide good flux matching. Where the set of linear equations expressed by equation 9.2 describes an overdetermined system ($i > j$) with linearly independent constraints, no exact solution exists, and a perfect matching cannot be found. We can instead compute a regularised solution with a similar quality of matching, with the additional goal of preferring desirable properties of the solution. In the case where equation 9.2 is underdetermined (where $i < j$), the solution is not unique, and again we are well-placed to solve for a regularised solution.

¹The smallest singular values of the initial matrix become largest in the inverted matrix.

9.2.1 Tikhonov Regularisation

A regularisation term, Γ , can be included in the LLS minimisation which preferences solutions that minimise the change between neighbouring coefficients:

$$\|\mathbf{N}\bar{c} - \bar{F}\| + \|\Gamma\bar{c}\|. \quad (9.3)$$

This method is known as Tikhonov regularization. The preferred solution is achieved with the choice of a suitable Tikhonov matrix, Γ . In this case, Γ is chosen as a difference matrix \mathbf{A} , multiplied by a tunable scaling factor λ , such that the additional minimisation condition is

$$\|\Gamma\bar{c}\| = \|\lambda\mathbf{A}\bar{c}\| = 0 \quad (9.4)$$

with

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 & 0 \\ -1 & 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & -1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & -1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & 0 & \cdots & -1 & 0 \end{pmatrix}. \quad (9.5)$$

The expression to solve for the regularised coefficients is now [79]:

$$\bar{c} = (\mathbf{N}^T \mathbf{P} \mathbf{N} + \mathbf{\Gamma}^T \mathbf{\Gamma})^{-1} \mathbf{N}^T \mathbf{P} \bar{F}, \quad (9.6)$$

with \mathbf{P} an $i \times j$ inverse-covariance matrix. \mathbf{P} is currently employed as a diagonalised scaling matrix of the form

$$\begin{pmatrix} \overbrace{\alpha}^{OORF} & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & \alpha & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & \alpha & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & \alpha \end{pmatrix} \quad (9.7)$$

with scaling out-of-range factor (OORF) $\alpha < 1$ replacing the diagonal entries corresponding to low and/or high energy bins to be de-prioritised. This re-scaling of each energy bin's importance to the solution re-weights flux regions which otherwise reduce matching quality in areas of importance. Alternatives

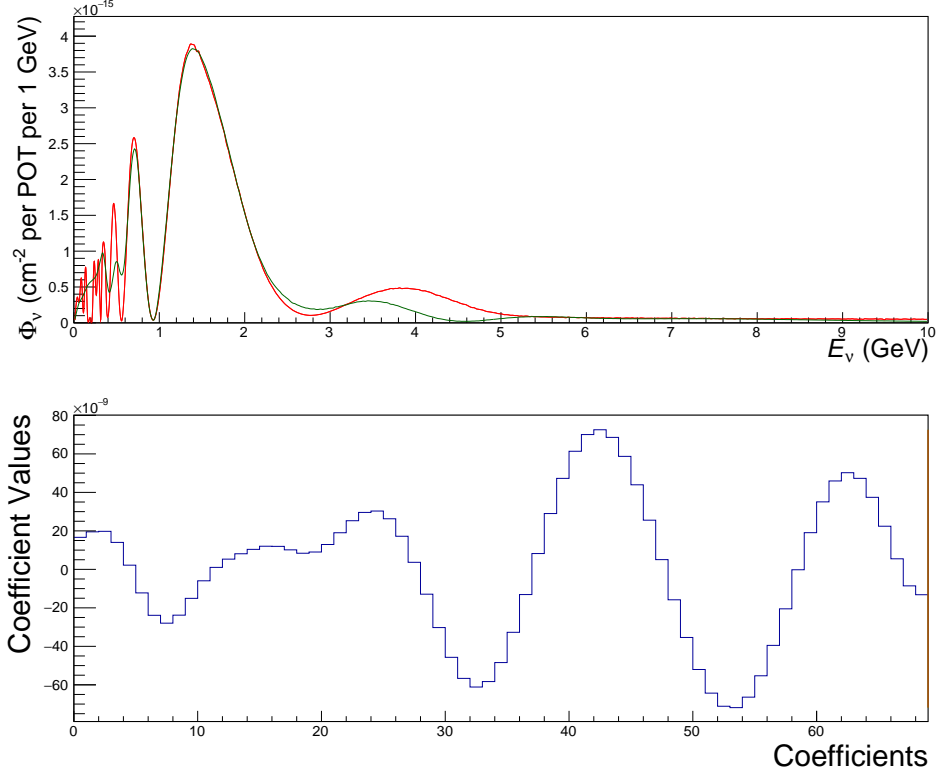


Figure 9.5: (Top) Flux matching between an oscillated FD flux (red) and the linear combination of ND fluxes (green) obtained using Tikhonov regularisation with $\lambda = 10^{-8}$, and (bottom) the corresponding regularised coefficients.

can be employed, such as reweighting fluxes by their energy bin values to prioritise better matching at larger neutrino energy. This can be important for managing feed-down energy effects.

To yield a solution, the strength of the regularisation parameter λ must first be chosen, with larger values enforcing greater smoothness in the resulting coefficients. A first pass with a value of λ equal to 10^{-8} returns a reasonable flux matching and significantly smoothed coefficients (Fig. 9.5), whilst choosing $\lambda = 10^{-9}$ returns an improved matching with more well-defined peaks at 0.5 GeV, but significantly rougher coefficients (Fig. 9.6). Clearly a more sophisticated method is required to select an optimal value. A robust approach for choosing a value for the scaling factor λ is the L-curve method [80, 81]. An L-curve is produced by comparison of the evolution of two terms related to the norms in equation 9.3,

$$\rho = \log (\|\mathbf{N}\bar{c} - \bar{F}\|) \quad (9.8)$$

and

$$\eta = \log (\|\mathbf{A}\bar{c}\|). \quad (9.9)$$

The optimal value of λ is found at the point where decreasing λ ceases to

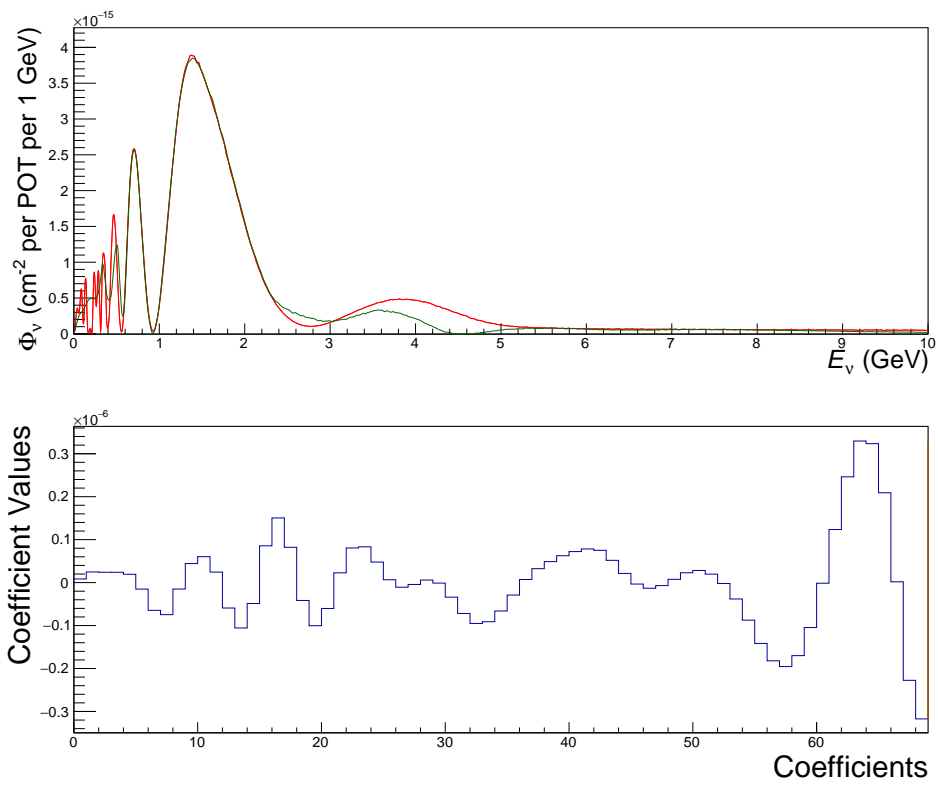


Figure 9.6: (Top) Flux matching between an oscillated FD flux (red) and the linear combination of ND fluxes (green) obtained using Tikhonov regularisation with $\lambda = 10^{-9}$, and (bottom) the corresponding regularised coefficients.

improve the matching, no longer reducing the residual norm, ρ , and instead increases the regularization norm term, η , corresponding to a sharp increase in L-curve curvature. This curvature is calculated as

$$2 \frac{\rho' \eta'' - \eta' \rho''}{((\rho')^2 + (\eta')^2)^{3/2}}. \quad (9.10)$$

An example of one such L-curve, its curvature, and the resulting nominal DUNE-PRISM solution is shown in figure 9.7. Here, the flux-matching solution uses an OORF to ignore agreement below 0.5 GeV and above 3.5 GeV, thereby allowing for greater agreement over this central interval. It should be noted that finding an optimal λ with this method significantly increases run-time, dependant on the sparsity of such a scan. The solution obtained provides a good match to the oscillated FD flux between the 0.8 GeV and 3.5 GeV, but struggles with both the rapidly oscillating flux features at low energies and with features above 3.5 GeV.

With a sufficiently large set of linearly independent states to form a basis over the space, unregularised linear combinations can model any such distribution. In DUNE-PRISM this is not the case due to both the regularisation term and the limited number of flux windows, such that an exact solution does not exist even without regularisation (in fact, a solution which provides an exact matching in this way may do so through large cancellations between similar fluxes which are inherently more susceptible to larger propagated statistical or systematic errors). To provide good flux-matching agreement, this constrains DUNE-PRISM to require energy spectra from off-axis flux-windows which more largely resemble the oscillation features being targeted. However, the peak energy of the flux distribution is maximised on-axis over an energy range of 2 to 2.5 GeV. As the detector moves further off-axis, each ν_μ flux window forms a progressively sharper, lower energy peak. For this reason, reproducing features above 3 GeV is difficult with DUNE-PRISM, with the poor matching of the FD flux oscillation feature between 3 and 5 GeV an example of this. To remedy this, additional ND flux information is required. Chapter 10 considers the introduction of additional ND fluxes and the improvements this offers to a FD flux matching solution.

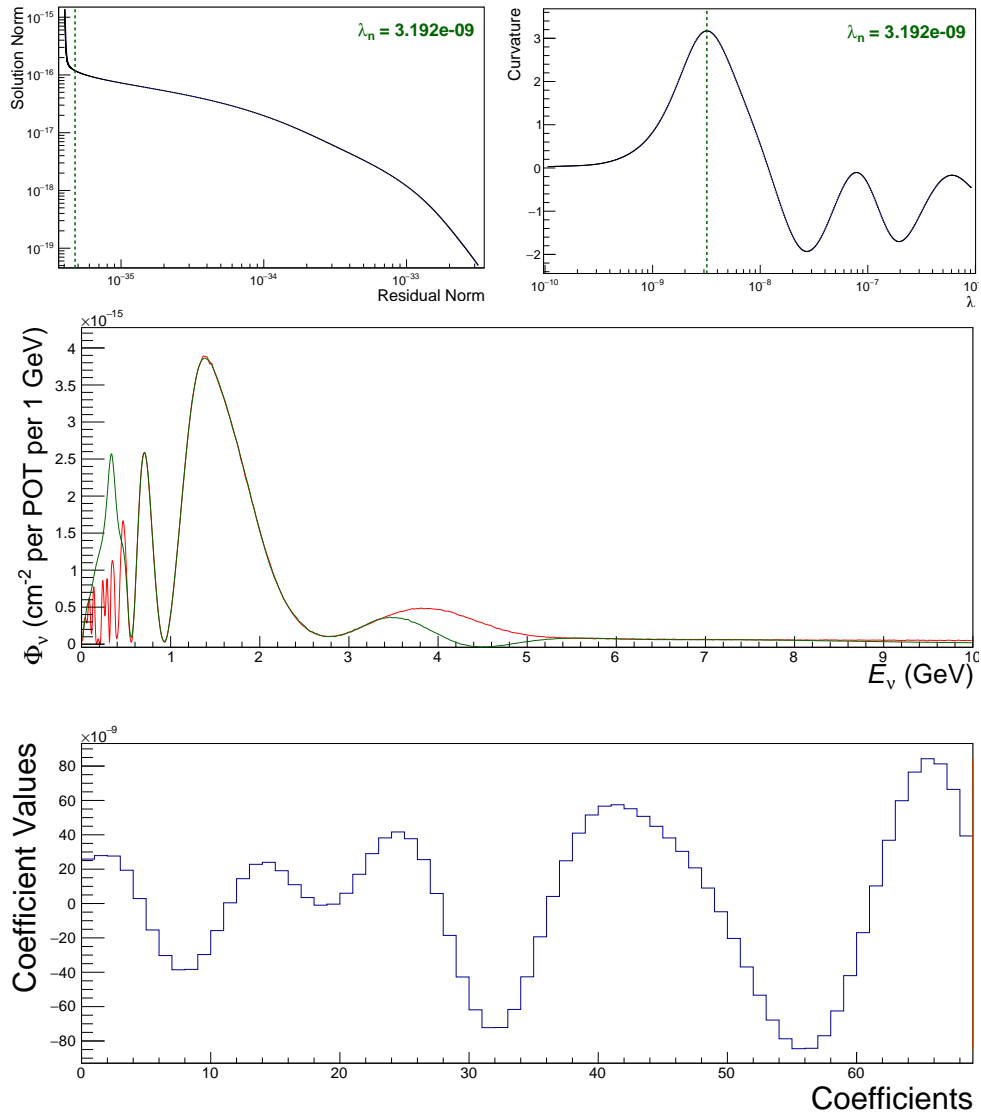


Figure 9.7: An L-curve (top left) constructed from the evolution of the residual norm and solution norm terms, ρ and η , and its curvature (top right), with the resulting value of $\lambda = 3.192 \times 10^{-9}$ chosen producing the flux-matching (middle) and regularised coefficients (bottom).

9.3 Summary

The DUNE-PRISM analysis strategy provides a data-driven approach to control systematic uncertainties due to neutrino interaction modelling. The uncertainty associated with reconstructing the incident neutrino energy from observable final state particles is too large to meet the precision requirements set by the DUNE physics goals. This is due to the potentially large missing energy contribution from neutrons (and other particles), as well as the misidentification of charged particles. With PRISM, the differences in energy spectra observed by the near- and far-detectors can be significantly reduced, with additional work on-going to reduce the impact of detector effects such as differences in detector geometry. This analysis strategy can then be used to extract the neutrino oscillation parameters with minimal interaction model dependence.

The DUNE-PRISM methodology shown here offers significant improvements on earlier iterations which used minimisation fitting algorithms to identify appropriate linear combinations of near-detector fluxes. In replacing this minimisation approach with a linear algebra solution, the computational time required to find a solution is several orders of magnitude lower. In addition, the inclusion of a regularisation term as a Tikhonov matrix allows for the enforcement of smoothly varying coefficients from adjacent flux windows. By smoothly varying coefficients, the reliance on large cancellations prone to large uncertainties is reduced. Additionally, the introduction of an inverse covariance matrix offers a high level of solution tuning to prioritise or de-prioritise flux-matching over different energy intervals. Better agreement can be found over the flux features which most significantly inform an oscillation analysis by prioritising those regions of the energy spectra over others.

This strategy is expected to make a large contribution to the eventual DUNE oscillation analysis, with studies showing its planned implementation already outlined in the ND CDR [76].

Chapter 10

Beam-PRISM

10.1 Beam Options

The underlying principle of a PRISM analysis does not require off-axis detector measurements, but rather well-known flux predictions with multiple linearly independent fluxes measured with the same detector. An alternative to moving off-axis to obtain such ND fluxes is to directly alter the beam. This is most easily achieved within the nominal design scope by operating the powered magnetic focusing horns under non-nominal current. In addition, if the horn design allows for disconnection of individual horns from the power-supply, subsets of the nominal three-horn design can be operated to produce and expand such a set of non-nominal ND fluxes.

The magnetic focusing horn design at time of this study specifies 3 horns, labelled US to DS as A, B, and C, connected in series to one power-supply. Further details regarding the target and horns system are discussed in section 8.2.1. Under the most recent design, all powered horns are required to operate on equal current. By manually disconnecting one or more magnetic horns, sub-combinations of the three horn design could be operated. This work considers combinations of horns A, AB, AC, ABC and BC at 14 currents between 0 and 350 kA: 0, 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA.¹ Additionally, further studies undertaken here merited the inclusion of the ABC horn combination (all horns powered) at currents of 260, 270, and 280 kA. Horn configurations running only B or C individually were not considered due to poor focusing resulting in prohibitively low event rates. Fluxes are simulated with 5E8 Protons On Target (POT) per configuration and current combination. Primary analysis has utilised “nu-mode” beams produced with FHC running.

Each horn configuration allows a continuum of off-axis ND fluxes, up to the same off-axis angle as PRISM. On-axis ND fluxes composed from the 4 m

¹As 0 kA designates all horns being operated with no current, this setting is equivalent for each sub-combination and so is included once, grouped under the A+B+C configuration.

fiducial detector width for each combination of powered horns are shown in figure 10.1, with their ratio with respect to the nominal flux shown in figure 10.2, highlighting the change in flux shape observed on-axis as current is varied in each configuration. The full OA flux spectra for each configuration and current combination are shown in figures 10.3, 10.4, 10.5, 10.6, and 10.7.

Further beam options which involve altering the proton beam energy delivered to the target have also been investigated, but these fluxes are not shown here as their utility has proven limited. Additionally, linear algebra solutions involving matrix inversion increase in computational complexity with order $\gtrsim \mathcal{O}(n^{2.5})$, depending on the implementation and algorithms used, hence including all such additional flux windows would be prohibitively computationally demanding. The set of off-axis flux windows incorporated in the standard DUNE-PRISM analysis are referred to hereafter as nominal fluxes. The non-nominal beam fluxes and their sub-component flux windows introduced here by operating alternative horn settings are referred to as horn configuration (HC) fluxes hereafter for brevity.

10.1.1 Non-Nominal Run-Time

The expanded set of ND fluxes introduced here allows for the inclusion of up to 68 additional beam settings, each spanning up to 37.5 m off-axis. With an off-axis flux window for each HC flux of 0.5 m width, concordant with the width utilised in the nominal PRISM flux-matching, over 5000 flux contributions are available for inclusion in a Beam-PRISM solution. Such a large set of fluxes will likely span the space, given a coarse enough energy binning, and hence a perfect matching solution utilising all fluxes is straightforward to compute. However, each additional flux incorporated must be sampled in sufficient number so as to provide statistically reliable distributions.

If we require a statistical error of 1% in the peak bin content of each sampled flux distribution (with 40 bins per GeV binning), then given a HC flux distribution with peak intensity of similar size to the nominal flux intensity, the required event rate corresponds to an approximate week of running where all fiducial events over the 4 m fiducial width at a detector stop are aggregated. Comparisons can be made with the sample DUNE-PRISM run plan shown in table 9.1. Similar measurements of every off-axis detector stop for each beam setting would necessitate years of non-nominal beam running. As such, solutions featuring many or all of the available flux-windows at once are infeasible. The problem is then instead how best to reduce this sample to an optimal subsample, providing the most improved flux-matching for the least time spent running in non-nominal beam conditions.

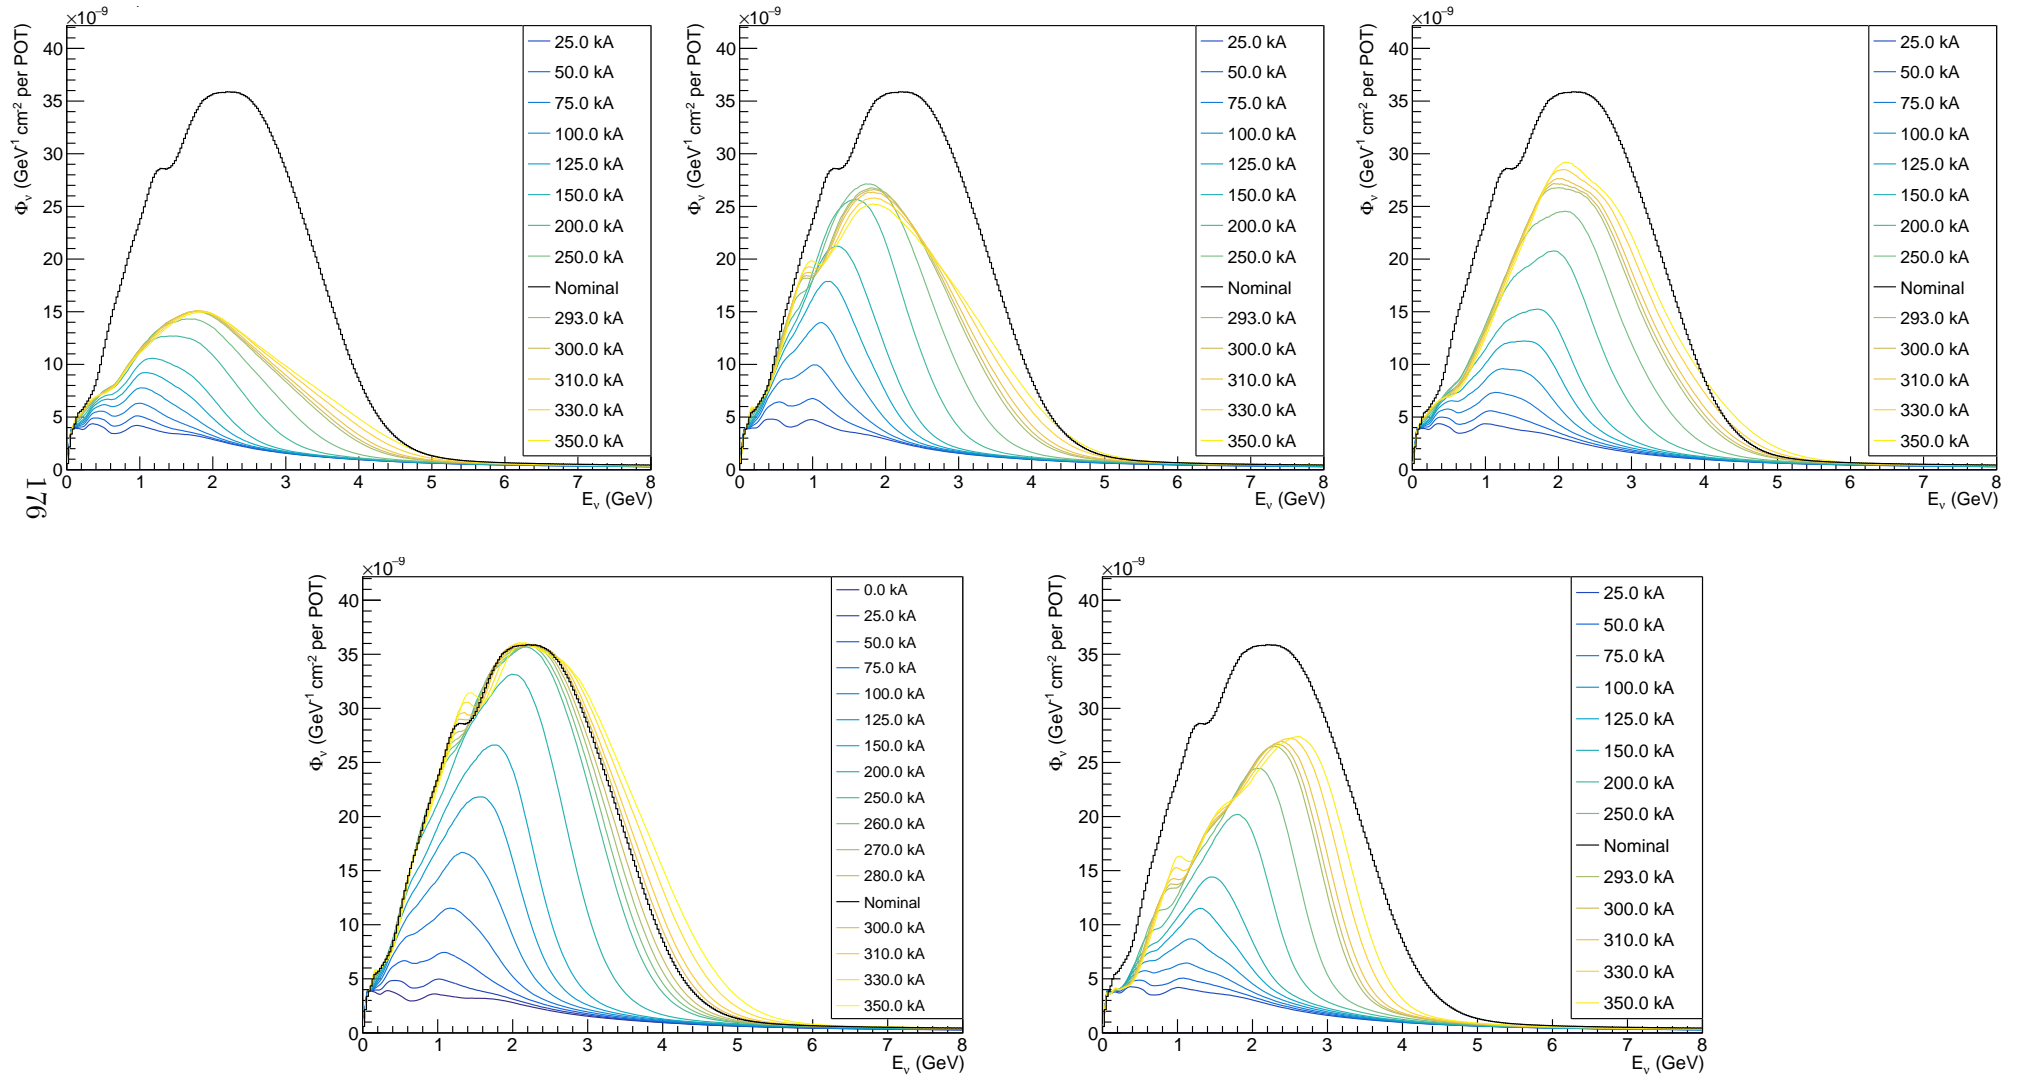


Figure 10.1: On-axis ν_μ flux distributions for a 3 m fiducial width ND 575m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper row first, for FHC horn combinations A, AB, AC, ABC and BC respectively.

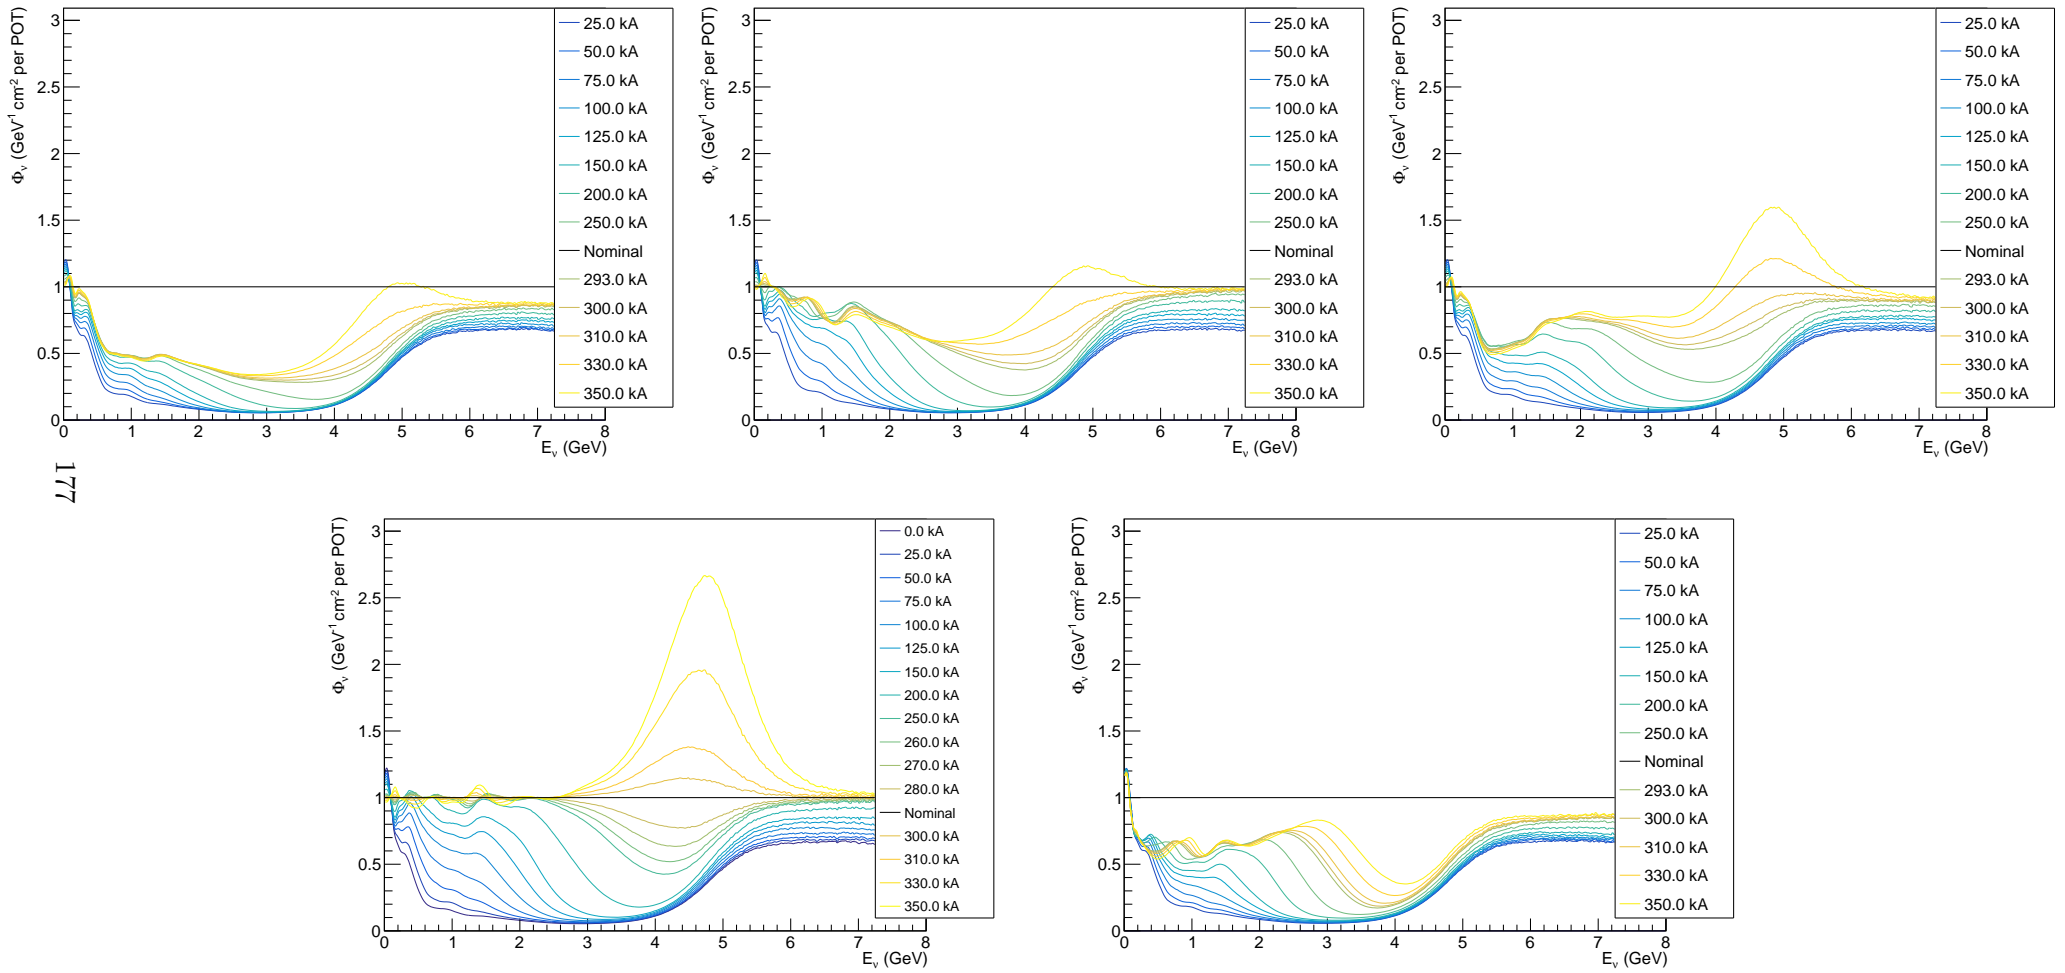


Figure 10.2: Normalised on-axis ν_μ flux distributions for a 3 m fiducial width ND 575m downstream, generated for 5×10^8 protons on target (POT) each. All fluxes are normalised with respect to the nominal operating flux. Fluxes are shown, left to right, upper row first, for horn combinations A, AB, AC, ABC and BC respectively.




Figure 10.3: Off-axis ν_μ flux spectra produced by varying horn A current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA respectively.




Figure 10.4: Off-axis ν_μ flux spectra produced by varying horn A and B current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA respectively.




Figure 10.5: Off-axis ν_μ flux spectra produced by varying horn A and C current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA respectively.




Figure 10.6: Off-axis ν_μ flux spectra produced by varying horn A, B and C current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 0, 25, 50, 75, 100, 125, 150, 200, 250, 260, 270, 280, 293, 300, 310, 330 and 350 kA respectively.




Figure 10.7: Off-axis ν_μ flux spectra produced by varying horn B and C current powerings for a ND 575 m downstream, generated for 5×10^8 protons on target (POT) each. Fluxes are shown, left to right, upper rows first, for currents of 25, 50, 75, 100, 125, 150, 200, 250, 293, 300, 310, 330 and 350 kA respectively.

10.2 Method

Due to the induced variation in the FD event rate, flux contributions which involve running the neutrino beam in a non-nominal state must be minimised. Including a limited number of additional flux-windows in the Tikhonov flux-matching shown in section 9.2.1 is straightforward, requiring only a small modification of the regularisation matrix with the addition of diagonal matrix elements, λ_{hc} , for each additional row and column corresponding to each additional flux. This yields a solution which incorporates all selected HC fluxes into the nominal PRISM flux-matching, with their contribution controlled by the regularisation parameter λ_{hc} . However, this procedure does not provide any insight into which HC fluxes should be selected. To include all additional fluxes in a DUNE-PRISM-style fit, a more sophisticated flux-matching method is required which can consider multiple beam options alongside the N_n nominal DUNE-PRISM off-axis fluxes, reducing the number of additional beams so as to identify the sparsest solution which achieves the required quality of matching.

10.2.1 Possible Approaches

Several different approaches to this problem have been investigated, beginning with the simplest case: a double application of the Tikhonov solution applied in section 9.2.1. This method first performs a Tikhonov solve using all available flux information but with an expanded regularisation matrix Γ employing a diagonalised regularisation parameter λ_{hc} for non-nominal fluxes. For these additional fluxes, diagonalised rows and columns are added to Γ , of the form $\Gamma_{ij} = \lambda_{hc}\delta_{ij} \forall i, j > N_n$ with a prohibitively high λ_{hc} value used to limit the contribution from additional fluxes. All but the N largest HC fluxes are removed for a second iteration of Tikhonov solve, with N chosen by the user, providing the final matching. A case-by-case evaluation of the flux matching improvement with increasing N would allow for the choice which best compromises between non-nominal beam run-time and matching quality.

More complex solutions involving orthogonalisation/orthonormalisation via the Gram-Schmidt procedure were also tested, based on [82]. These solutions iteratively construct an orthogonal or orthonormal flux set with respect to the nominal flux sample plus any HC fluxes included from prior iterations, introducing from the orthonormalised fluxes the sample with the largest projection onto the oscillated FD flux or FD flux-matching residual. At each step, the projection of a function n onto another function f is calculated as

$$proj_n = \frac{\langle n, f \rangle}{\langle f, f \rangle} f \quad (10.1)$$

where $\langle \circ, \circ \rangle$ denotes the inner product of the two functions,

$$\langle n, f \rangle = \int_i^j n \times f \, dE, \quad (10.2)$$

over the energy range (i, j) .

Each of these methods suffers from the same larger issue, in that they do not appropriately consider the impact of multiple HC fluxes used in combination with one-another. In the double Tikhonov method, the inclusion of flux-windows arises from the solution coefficients of each flux-window when considered in combination with the entire flux sample. This has the tendency to wash out the most useful flux-windows, particularly when solving with larger flux samples, and instead prioritise a larger number of atypical fluxes. Conversely, the orthonormalised solution assesses the incorporation of only one flux window at each stage. As a result, if any two (or more) flux windows would provide a significant improvement to the matching when combined but a poor contribution in isolation they would be introduced only after many higher ranking fluxes. Due to these limitations, an improved method is required.

10.2.2 Compressed Sensing

Compressed sensing is a modern signal processing technique designed for efficient signal reconstruction by identifying solutions to underdetermined linear systems [83]. The technique operates by enforcing sparsity on the solutions, thereby adding additional constraints which are exploited to recover a signal with far fewer samples than usually required. In Beam-PRISM, the goal is similarly to find a sparse solution which best recovers the oscillated FD flux. The problem of solving for the most sparse solution, minimisation of the ℓ_0 norm², is considered NP-hard and has until recently been very computationally expensive. However, recent developments related to compressed sensing have shown that a combinatorial ℓ_0 minimisation solution can be approximated with an iterative linear ℓ_1 minimisation approach. This can be extended to an iterative ℓ_2 minimisation, such as that obtained from Tikhonov regularised solutions [84–86].

To retrieve the sparse solution of coefficients which satisfies the flux-matching criteria, the approach mirrors the Tikhonov regularisation shown in section 9.2.1. The problem remains to minimise eqn. 9.3, but with an expanded regularisation matrix Γ which properly distinguishes between the N_n nominal and N_{HC} HC fluxes. To do so, an expanded rank $(N_n + N_{HC})$ regularisation

²The ℓ_0 norm counts the number of non-zero elements in a vector, while the ℓ_1 norm sums the magnitudes of the vector elements. These operations contrast the ℓ_2 norm, or Euclidean norm, which sums the vector elements in quadrature and is more commonly used in physics.

matrix is constructed of the form

$$\begin{pmatrix} \overbrace{\lambda_n & 0 & 0 & 0}^{N_n} & \overbrace{0 & \dots & 0}^{N_{HC}} \\ -\lambda_n & \ddots & 0 & 0 & 0 & \dots & 0 \\ 0 & \ddots & \lambda_n & 0 & 0 & \dots & 0 \\ \vdots & \ddots & -\lambda_n & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & \omega_1^t & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \omega_{N_{HC}}^t \end{pmatrix} \quad (10.3)$$

where λ_n is the nominal PRISM regularisation value, chosen using the L-curve method during nominal flux-matching, and ω_i^t are iteratively calculated HC flux weights. For each flux-window, N_i , a corresponding weight, ω_i , is initialised at step $t = 0$ with a value of $\omega_i^0 = \lambda_n$ (alternative initial values are valid). These $\omega_i^{t=0}$ are used to solve for the regularised coefficients c_i^t , as before. For each successive step $t + 1$, a new ω_i^{t+1} is calculated from the previous coefficient solution as

$$\omega_i^{t+1} = \left| \frac{\mathcal{N}^{-1}}{\mathcal{N}c_i^t + \epsilon} \right|, \quad (10.4)$$

where \mathcal{N} is a constant normalisation factor and ϵ is a stability factor to prevent weightings going to infinity [85]. The normalisation factor rescales flux coefficients to close to unity, and is typically set to 10^8 or 10^7 . Varying this parameter alters the resulting regularisation strength and can be selected to prefer sparser solutions. The choice of ϵ ensures that a zero-valued component in a coefficient c_i^t does not prohibit a non-zero estimate for c_i^{t+1} , and is selected to be smaller than the expected renormalised magnitudes of returned coefficients, typically chosen between 10^{-1} and 10^{-3} . As varying ϵ alters the way in which low-weighted flux coefficients go to zero, this can be used to preference a solution to incorporate more or fewer non-zero HC coefficients. The solution is once again computed from Eq. 9.6, with the evolution of coefficients converging to a sparse solution over $\mathcal{O}(10)$ iterations. The evolution of iteratively reweighted coefficients for one such example flux-matching is shown converging over 12 iterations in figure 10.8. Iterative reweightings are usually run for 30 steps to ensure convergence. As all fluxes are equally weighted for the first solution, most HC fluxes return non-zero coefficients. Subsequent steps increase the penalty weighting of the least contributing fluxes, whilst the most useful suffer a reduced penalty and subsequently make larger contributions. The nominal PRISM fluxes vary with each solution as reweighted HC fluxes make larger




Figure 10.8: Evolution of iteratively reweighted coefficients over 12 iterations for an oscillated FD flux-matching solution. Coefficients 270 to 318, shown between two black lines, show the variation of nominal PRISM fluxes with constant regularisation in response to the reweighting of contributing HC fluxes. Coefficient ranges for different horn configurations are shown in figure 10.9.




Figure 10.9: Coefficients resulting from the filtered solution of Fig. 10.8. Green lines separate coefficients corresponding to flux-windows from differing current powerings, red lines separate differing horn combinations.

or smaller contributions. Once the procedure has converged, a final filtering step is applied to ensure very low coefficient values are set to zero and a final Tikhonov solution returns the minimised solution of coefficients. In this example case, the filtered solution coefficients are displayed in figure 10.9, with green lines separating coefficients corresponding to flux-windows produced with differing horn currents, and red lines separating flux-windows from differing combinations of powered horns. As with the nominal PRISM solutions, each subset of HC coefficients is shown with lower number coefficients at lower axial displacement and higher number coefficients at greater displacement.

The process leads to only a small number of HC configurations being utilised in matching the oscillated FD flux, while still accommodating variations in the nominal DUNE-PRISM off-axis flux combination as they respond to additional

Horn Configurations				
A only	A+B	A+C	A+B+C	B+C
330 kA, 0.5 m	300 kA, 16.5 m		125 kA, 0.5 m	250 kA, 16.5 m
			330 kA, 0.5 m	310 kA, 20.5 m

Table 10.1: Horn configurations, their operating currents and the central off-axis positions of the detector utilised in the example iterative flux-matching solution with coefficients shown in figure 10.9. The solution allows for HC fluxes at currents up to 330 kA.

HC flux components. The HC fluxes which are utilised as a result are listed in table 10.1.

In order to reduce both statistical uncertainties on the utilised flux distributions and computing resource intensity requirements, solutions shown here consider HC fluxes aggregated over 4 m fiducial width flux windows. This reduces the maximum size of the involved matrices to $\sim 1000 \times 1000$.

10.3 Flux-Matching Objectives

The importance of flux-matching over different sections of the oscillated FD spectra should be noted, as not all sections contribute equally in a long baseline CP violation analysis. Oscillation parameters are primarily observed in a PRISM-style analysis by their impact on the first and second oscillation peaks. As highlighted by the choice of oscillation parameters in each ν_μ disappearance measurement in figure 10.10, Δm_{atm}^2 alters both the height of the highest energy peaks at 0.6 and 4 GeV and the positioning of the oscillation maxima/minima. Conversely, altering $\sin^2 \theta_{23}$ deforms the shape of the oscillations by varying the height of the oscillation maxima/minima. While matching these features

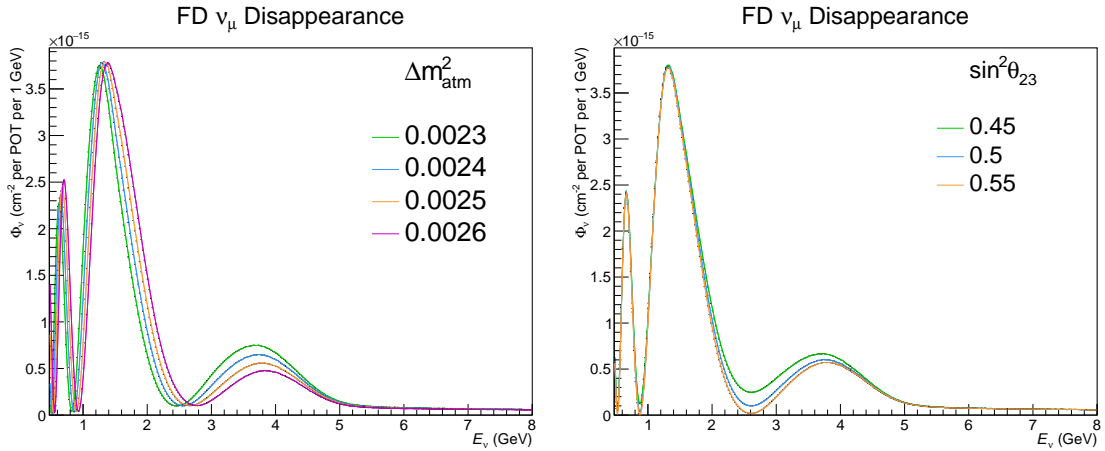


Figure 10.10: Oscillated ν_μ FD fluxes showing the impact of different predicted values for (left) Δm_{atm}^2 and (right) $\sin^2 \theta_{23}$.

contributes to measurements of the corresponding oscillation parameters, good flux-matching at higher energies reduces the impact of feed-down effects from higher-energy interactions which can mis-attribute reconstructed neutrino energies.

As such, it is important to provide good agreement between the oscillated FD spectrum and the flux-matching solution over not just the oscillation features used to extract parameters, but also over the extended high energy region. For this reason, solutions are constructed to prioritise higher energy agreement over lower energy agreement by reweighting fluxes by their energy bin values using the matrix \mathbf{P} in equation 9.6. In addition, this further de-prioritises matching the rapidly oscillating region below 0.5 GeV, an area of limited contribution in the near-far oscillation analysis. One such method to disregard the rapidly oscillating low energy region appends a gaussian decay to the target oscillated FD flux over this energy interval, replacing the low energy interval. This can be beneficial for the fit quality elsewhere.

10.4 Results

Having defined an appropriate solution method, a number of useful applications of a Beam-PRISM solution are explored.

10.4.1 PRISM On-Axis

With the inclusion of multiple on-axis beam options, an entirely on-axis PRISM-style flux-matching can be attempted which negates the need for the complex technological requirements involved in moving a cryogenically cooled, liquid-filled near detector to arbitrary off-axis positions.

Several variations are shown, with figure 10.11 showing the resulting flux-matching where we require horn current operation of no greater than the nominal operating current of 293 kA. Each HC flux is constructed from a single 4 m fiducial width on-axis flux-window spanning a distance of -1.5 m to 2.5 m transverse to the beam. The nominal flux spans the same distance but is composed of eight 0.5 m flux-windows which are allowed to vary in the same manner as the standard DUNE-PRISM analysis. As each flux-window is averaged against the number of off-axis bins, this means coefficient values shown for HC fluxes appear eight times larger than the underlying neutrino event rate would suggest. To produce a high quality flux-matching across the energy range above 1 GeV, 10 different non-nominal horn configurations are required, shown in table 10.2. Due to the limited off-axis distance, the rapidly varying lower energy flux is difficult to replicate. This is a consequence of both the loss of the very sharply defined flux information at large off-axis distances, and the reduced flexibility from combining fewer flux windows. An entirely on-axis approach with similar flux-matching success to the nominal PRISM solution may be possible but would likely require many additional low current HC fluxes to provide useful flux information in this low energy interval.

Horn Configurations				
A only	A+B	A+C	A+B+C	B+C
100 kA	250 kA	100 kA	125 kA	293 kA
200 kA		293 kA	150 kA	
			200 kA	
			270 kA	

Table 10.2: Horn configurations and their operating currents utilised in the on-axis solution for flux-matching, allowing HC fluxes at or below the nominal current of 293 kA.

By expanding the current constraints to include higher current horn fluxes, the number of additional flux configurations can be reduced to 9 HC fluxes, shown in figure 10.12 and table 10.3. However, the agreement with the target

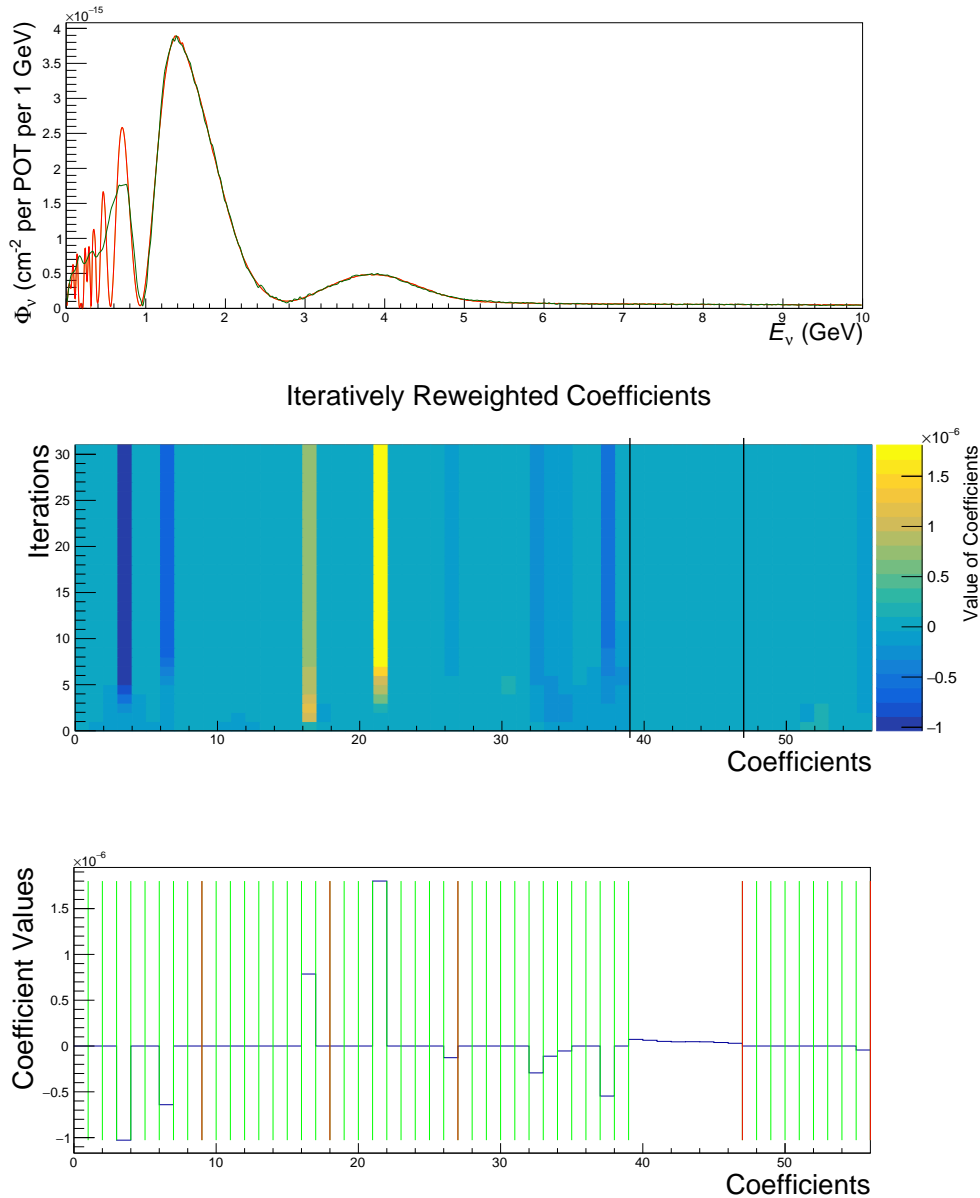


Figure 10.11: Energy-weighted iterative flux-matching over a targeted energy range of 0-10 GeV utilising eight on-axis nominal fluxes spanning -1.5 m to 2.5 m off-axis position, combined with 4 m wide on-axis HC fluxes produced from utilising up to 293 kA horn current. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.

oscillated FD spectrum at energies below 1 GeV is not improved even with higher allowed currents, as these beams do not provide significant additional flux information in this low energy region.

Horn Configurations				
A only	A+B	A+C	A+B+C	B+C
100 kA	100 kA	100 kA	100 kA	
	250 kA		150 kA	
			200 kA	
			250 kA	
			330 kA	

Table 10.3: Horn configurations and their operating currents utilised in the on-axis solution for flux-matching, allowing for HC fluxes at currents up to 350 kA.

Different solutions are obtained by varying the stability coefficient ϵ , the normalisation factor \mathcal{N} , and the initial regularisation parameter, λ_{HC} . Varying ϵ and \mathcal{N} alters the coefficient value at which each coefficient’s regularisation weighting will begin to increase over successive iterations, thereby rejecting those flux-windows from the solution, barring new-found utility at future iterations. This allows for the increase or reduction of the number of additional HC fluxes in a given solution, within some limit unique to the problem posed. In practice, even with such a restricted set of allowed fluxes, several equally well-matched solutions can be obtained with alternative combinations of HC fluxes, each of which returns a residual norm of similar size.

For each solution, the merits of incorporating an entirely on-axis flux-matching such as this must be balanced against the significant impact on FD event collection as any changes to the ND beam will similarly alter neutrino events at the FD site. In the case of an entirely on-axis PRISM-style analysis, this would avoid the high expense incurred from extension of the ND hall to an off-axis distance of 30 m, as well as the complications of designing and constructing the movement apparatus for the large cryo-cooled ND.

In exchange, requiring 9 or more HC fluxes would involve committing to a large period of data-taking with non-nominal flux conditions. With a well-understood neutrino beam, such changes could be accommodated in a long-baseline CP analysis to allow use of this data, though this may be non-trivial in some cases. The long-baseline ND+FD analysis is not currently well-enough understood to predict the impact of this type of data inclusion, and so if a minimal analysis would only utilise FD events using the nominal neutrino beam, extensive incorporation of HC fluxes would require a costly reduction in FD neutrino events. This would, in turn, delay reporting of the core DUNE oscillation measurements beyond the schedule currently being assumed.

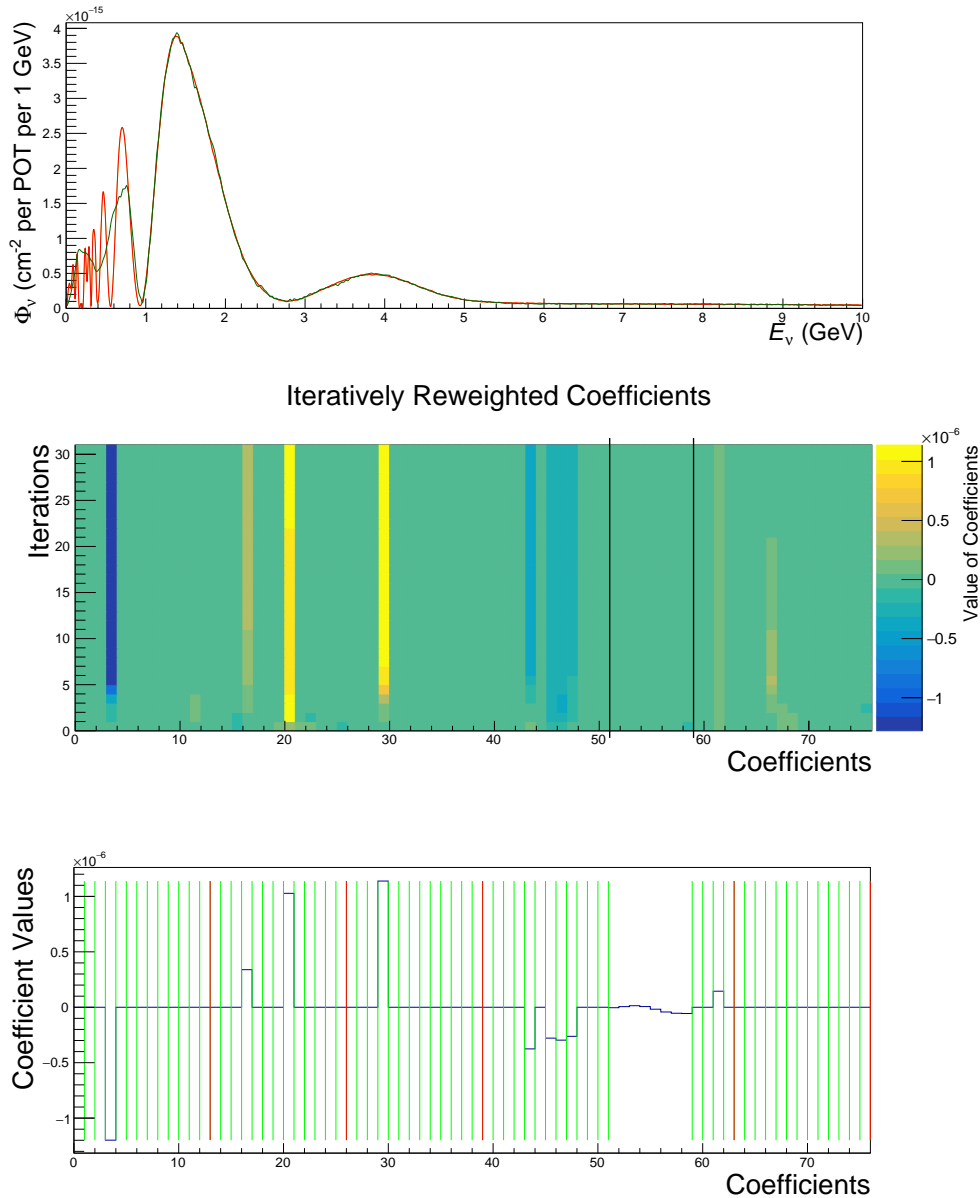


Figure 10.12: Energy-weighted iterative flux-matching over a targeted energy range of 0-10 GeV utilising eight on-axis nominal fluxes spanning -1.5 m to 2.5 m off-axis position, combined with 4 m wide on-axis HC fluxes produced from utilising up to 350 kA horn current. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.

Furthermore, varying power supply between different horn configurations may incur additional costs associated with any required adaptation of high-voltage systems used, whilst the procedure to remove one or more focusing horns from the power supply would likely require physical intervention, itself requiring additional periods without data-taking to allow for a decrease in target hall radioactivity and physical intervention time. Additionally, such physical interventions may introduce further systematic uncertainties in the horn positions over time which would need to be understood.

10.4.2 Expanding DUNE-PRISM

The addition of one or several non-nominal fluxes obtained by varying beam-line operating conditions can significantly expand the flux-basis spanned by the nominal-current off-axis fluxes. Two approaches to introduction of a complementary flux-set are considered:

- Minimal addition of HC fluxes, facilitating the maximal improvement with respect to the nominal PRISM flux-matching.
- Maximal useful addition of HC fluxes, to provide significant increase to flux-matching across all areas of the energy spectrum.

Unrestricted DUNE-PRISM and Beam Options

It is straightforward, either by scanning each flux individually or applying the Tikhonov iterative solver, to identify powerful contributors to flux-matching solutions. For each solution shown here, the HC fluxes are divided into 4 m fiducial width flux-windows, spanning off-axis distances of -1.5 m to 30.5 m, with nominal PRISM fluxes separated into 0.5 m fiducial width flux-windows between -1.5 m and 32.5 m off-axis.


For solutions which include only one additional HC flux, dramatic improvement to the matching agreement over the 4 GeV oscillation minimum is observed, replacing the nominal DUNE-PRISM's need for MC corrections in this region. Comparison of the available contributing flux-windows shows the most impactful additional fluxes are those on-axis fluxes at currents close to the nominal current. Initial studies using fewer fluxes produced at currents close to the nominal flux (prior to the inclusion of 260 kA, 270 kA, and 280 kA fluxes) found a significant improvement to the matching agreement over the 4 GeV oscillation minimum when allowing for horn currents up to 350 kA. The flux-matching shown in figure 10.13 allows currents up to 350 kA and includes a single on-axis flux at 330 kA, resulting in much better agreement over the 3 to 5 GeV oscillation feature. As this current significantly exceeds the nominal operating horn current, an alternative matching with a restricted current range

is made. Due to concerns with the maximum power safely transferable by the reference design for horn striplines, fluxes above 293 kA will ultimately not be utilised as the flux-matching is of similar quality without exceeding this value. Figure 10.14 shows the result when reducing the maximum allowed current to 293 kA. The single included HC coefficient corresponds to an on-axis flux at 250 kA. Here, an improvement over the 3 to 5 GeV interval is still observed, but with reduced agreement at and above 4 GeV. This arises due to the lower neutrino energies present in the 250 kA beam, contributing less new flux information in the region of interest. This motivated further investigation of horn currents below the nominal operating current, introducing fluxes produced from horns A, B, and C powered at 260, 270, and 280 kA.

With the introduction of this expanded flux-basis, another iterative solution using currents of 293 kA or lower is made. This flux-matching is shown in figure 10.15, where an additional horn flux of 260 kA observed on-axis is included, with events aggregated over the 4 m ND fiducial width. The iterative coefficient evolution during this procedure initially utilises on-axis fluxes from horns ABC at 250, 260, 270, and 280 kA before converging, suggesting each of these fluxes can contribute valuable flux information. A case-by-case analysis, separately introducing 260, 270, or 280 kA current fluxes returns the best result using a 280 kA current flux, with a larger albeit manageable coefficient value, shown in figure 10.16. This higher current flux returns better agreement over the 3 to 4 GeV interval. As the iterative solution must compromise between matching quality and the penalty term penalising larger non-nominal coefficients, the 260 kA flux which provides good matching for a smaller coefficient value is selected by the iterative solver.

This is straightforward to see when comparing the flux-matching deficit for nominal PRISM with the on-axis flux ratios at each current, normalised against the nominal flux, shown in figure 10.2. Those fluxes produced with greater current deviations from 293 kA show a larger relative signal, providing greater impact for a fixed coefficient value. When the allowed additional fluxes are increased progressively to accommodate currents up to 350 kA, the higher current fluxes are chosen at each stage for the same reason up to 330 kA, as this results in the largest signal in the region which previously showed poor matching. Due to the engineering concerns with utilising currents of this magnitude, the subsequent studies will not include fluxes above 293 kA.

Further improvement beyond this level by including significantly more HC fluxes has limited impact. The region with most error is at low energy, where rapidly fluctuating oscillation minima and maxima would require sharply defined low energy flux information to successfully match this region. Shape testing has shown the flux obtained from moving off-axis in each configuration begins to converge with the nominal off-axis flux at large distances, providing



Iteratively Reweighted Coefficients






Figure 10.13: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 32.5 m off-axis position, combined with 4 m wide HC fluxes up to 30.5 m off-axis produced from utilising horn currents up to 350 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.



Iteratively Reweighted Coefficients






Figure 10.14: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 32.5 m off-axis position, combined with 4 m wide HC fluxes up to 30.5 m off-axis produced from utilising horn currents up to 293 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.



Iteratively Reweighted Coefficients





Figure 10.15: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 32.5 m off-axis position, combined with 4 m wide HC fluxes up to 30.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the evolution of iteratively reweighted coefficients over 30 iterations, plus a filtering step, and (bottom) final coefficient values resulting from a filtered solution.

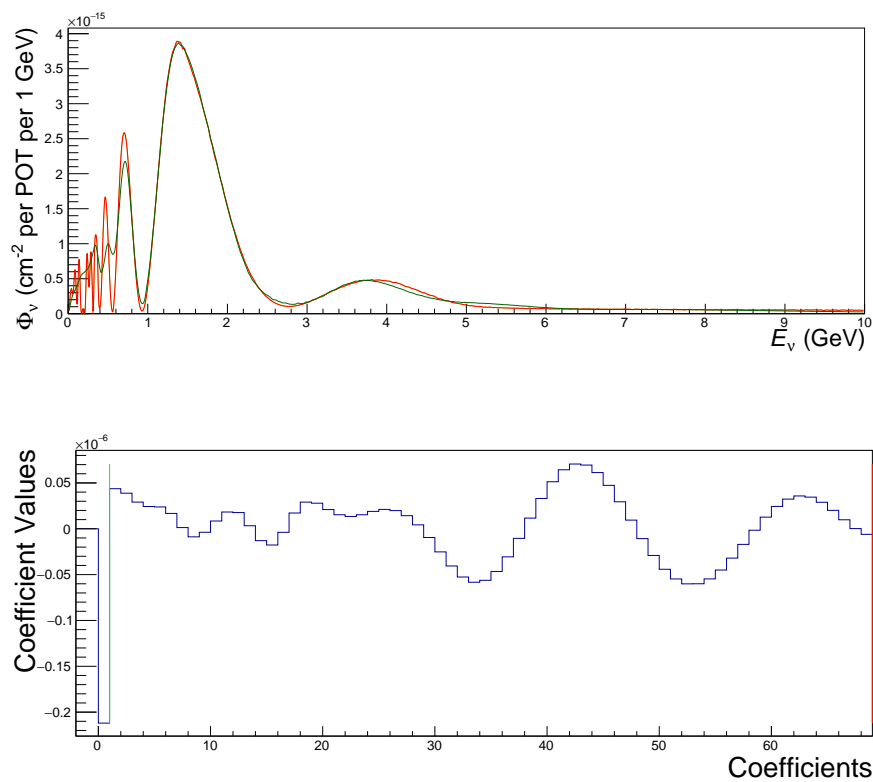


Figure 10.16: Energy-weighted flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 32.5 m off-axis position, combined with a single 4 m wide on-axis 280 kA flux-window. The two panels show (top) The resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (bottom) solution coefficient values.

little additional utility.

10.4.3 Restricted Off-Axis DUNE-PRISM

Alternative PRISM-style programs can be considered, where limitations on the off-axis freedom the near detector has can be compensated for with the inclusion of one or more HC fluxes. This is of interest particularly in evaluating the requirements for a full 32.5 m off-axis range, as the cost of extending the ND hall can incur many hundreds of thousands of dollars per metre. The following solutions consider the restriction of off-axis excursion to distances of 27.5 m, 20.5 m, and 14.5 m. Additional scenarios with maximum off-axis displacement limited to 25.5 m, 23.5 m, and 18.5 m are shown in appendix A. These solutions each show similar behaviour to a solution presented here.

Each BeamPRISM solution shown utilises a λ_n of 5.52×10^{-9} , with applied energy-weighting and a 20% de-preference for matching below 0.6 GeV to promote better agreement at higher energies. The comparative nominal PRISM solutions share these solution parameters. Each BeamPRISM solution allows for the inclusion of all HC fluxes at currents of 293 kA or below. To effectively vary the number of HC fluxes included, solutions utilise different values for the normalisation factor, \mathcal{N} , thereby returning different regularisation strengths for each solution and altering the preference for additional HC flux contributions. In addition, the scaling factor ϵ is varied between 10^{-2} for solutions with lower normalisation factor and 10^{-1} for solutions with large normalisation factor. This ensures coefficient regularisation remains small over a larger number of successive iterations and the procedure can more easily reintroduce HC fluxes at a later iteration in response to other included HC fluxes.

Further optimisations of both the nominal DUNE-PRISM solutions and BeamPRISM solutions could be made by the restriction of the target flux-matching interval to only consider subsections of the energy spectrum, thereby preferencing better matching over user-specified regions of interest at the cost of worse agreement overall, as used in the nominal DUNE-PRISM flux-matching shown in figure 9.7.

27.5m

The most generous of the restricted off-axis range solutions shown limits the maximum off-axis displacement to 27.5 m. Three solutions with increasing numbers of included HC fluxes are shown in figure 10.17, alongside a DUNE-PRISM only solution using the 27.5 m maximum off-axis range. This displacement off-axis still affords much of the low energy flux information utilised in the full 33 m solution, as shown in figure 9.3. As a result, the PRISM only solution over this range returns similar match quality to the 33 m solution, with losses only apparent below 0.6 GeV for the most part. Over the interval below this energy, the decay to zero flux at 0 GeV is poorly constrained, with large excesses in neutrino flux here. Additionally, as seen with the full 33 m extended solution, the lack of sharply peaked high energy flux information results in poor flux matching over the 3 to 5 GeV interval.

For each BeamPRISM solution, the inclusion of even a single additional HC flux affords greater performance over both of these poorly matched higher and lower energy flux intervals. The HC fluxes used for each case are shown in table 10.4, introducing one, two, or three additional HC fluxes as the normalisation factor is increased for each solution. Each utilised 4 m fiducial width flux spans the on-axis position where differing horn configurations return the largest differences in neutrino flux. The result is a more sharply defined oscillation peak at 0.7 GeV in each case, with increasingly better agreement over higher energies as additional flux information is included. However, agreement remains poor over the energy spectrum below 0.6 GeV where the rapidly oscillating features prove difficult to match. This is partially by construction as solutions are weighted to preference high-energy agreement, and partially due to the inherent difficulty of matching these features. Despite this, additional HC flux contributions allow these solutions to loosely follow the average position between these oscillation maxima and minima. Difficulties are noted in reducing the flux as closely to zero at the oscillation maximum around 0.95 GeV. Similar results are obtained when restricting the off-axis range to 25.5 or 23.5 m.

\mathcal{N}	Horn Configurations				
	A only	A+B	A+C	A+B+C	B+C
1×10^8		150 kA, 0.5 m	250 kA, 0.5 m	270 kA, 0.5 m	
5×10^7				200 kA, 0.5 m 260 kA, 0.5 m	
3×10^7				250 kA, 0.5 m	

Table 10.4: Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 27.5 m.

20.5m

With a reduced maximum displacement of 20.5 m off-axis, larger reductions in flux-matching at low energy emerge, with differences in preferred HC fluxes for each BeamPRISM solution. The HC fluxes used for each case are shown in table 10.5. Restricting the nominal PRISM solution in this way increases the discrepancy around the oscillation maximum at 0.95 GeV, with large features emerging at lower energies. For sparser solutions, the same non-nominal HC fluxes are introduced, but agreement over most of the 0.7 GeV peak feature is lost, even with the additional flexibility in flux information this includes. The least sparse solution shown now utilises 5 additional fluxes, with two detector stops centred at 4.5 m off-axis. The additional information this provides allows for significant improvement in matching the low energy features below 1 GeV in addition to the gains at higher energy. This solution has similar performance to the 18.5 m case, with superior agreement over the low energy region below 0.6 GeV.

\mathcal{N}	Horn Configurations				
	A only	A+B	A+C	A+B+C	B+C
1×10^8		250 kA, 4.5 m		200 kA, 0.5 m 200 kA, 4.5 m 270 kA, 0.5 m	200 kA, 0.5 m
5×10^7				200 kA, 0.5 m 260 kA, 0.5 m	
3×10^7				250 kA, 0.5 m	

Table 10.5: Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 20.5 m.

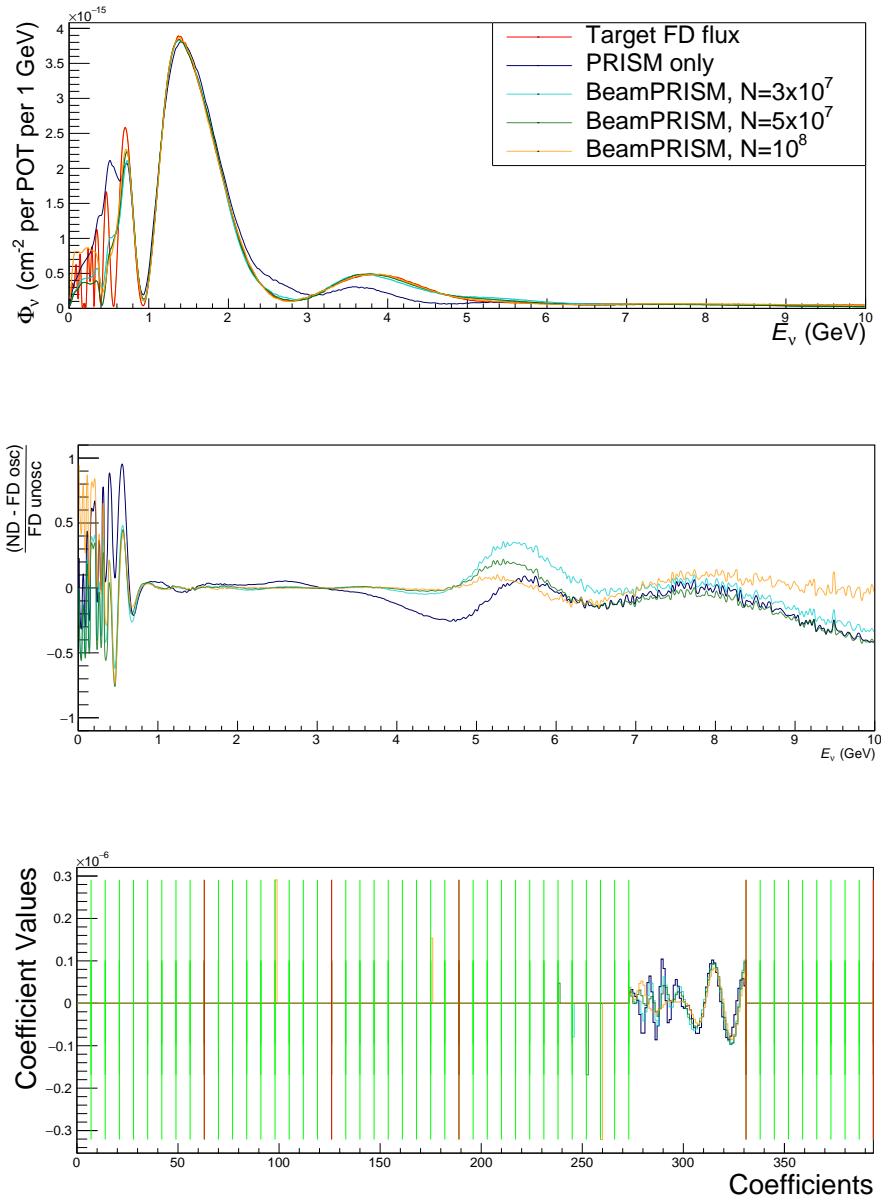


Figure 10.17: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 27.5 m off-axis position, combined with 4 m wide HC fluxes up to 26.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.

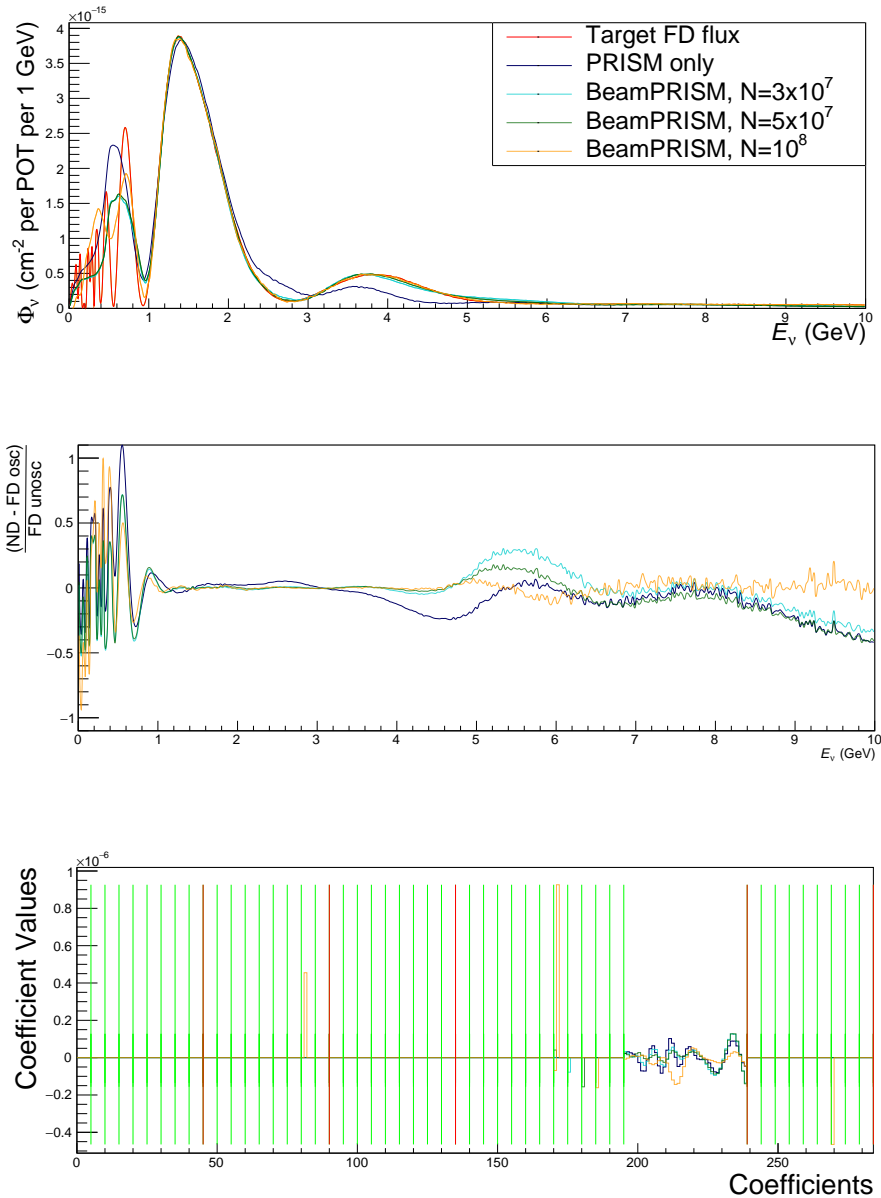


Figure 10.18: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 20.5 m off-axis position, combined with 4 m wide HC fluxes up to 18.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.

14.5m

Each of the utilised HC fluxes are highlighted in table 10.6, with the resulting flux-matching shown in figure 10.19. With the total off-axis distance restricted to 14.5 m, large discrepancies emerge in the PRISM only solution. The oscillation maxima and minima between 0.6 GeV and 1 GeV are largely ignored, with a displacement also induced in the energy position of the main oscillation feature at 1.5 GeV. The sparsest BeamPRISM solution focuses first on improving higher energy agreement between 3 and 5 GeV, with limited impact at lower energies. However solutions with $\mathcal{N} = 5 \times 10^7$ and 1×10^8 both introduce off-axis HC fluxes powering horns A and B at 250 kA, centred at 4.5 m off-axis. This affords better agreement over features below 1 GeV and at 1.5 GeV.

The least sparse solution shown now selects for 7 HC fluxes, with most displaced at varying off-axis distances. The inclusion of this additional flux information retains equivalent matching to fully extended solutions for features below 1 GeV.

\mathcal{N}	Horn Configurations				
	A only	A+B	A+C	A+B+C	B+C
1×10^8		100 kA, 12.5 m		200 kA, 4.5 m	293 kA, 4.5 m
		250 kA, 4.5 m		250 kA, 0.5 m	293 kA, 8.5 m
					293 kA, 12.5 m
5×10^7		250 kA, 4.5 m		150 kA, 0.5 m	
				260 kA, 0.5 m	
3×10^7				250 kA, 0.5 m	

Table 10.6: Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 14.5 m.

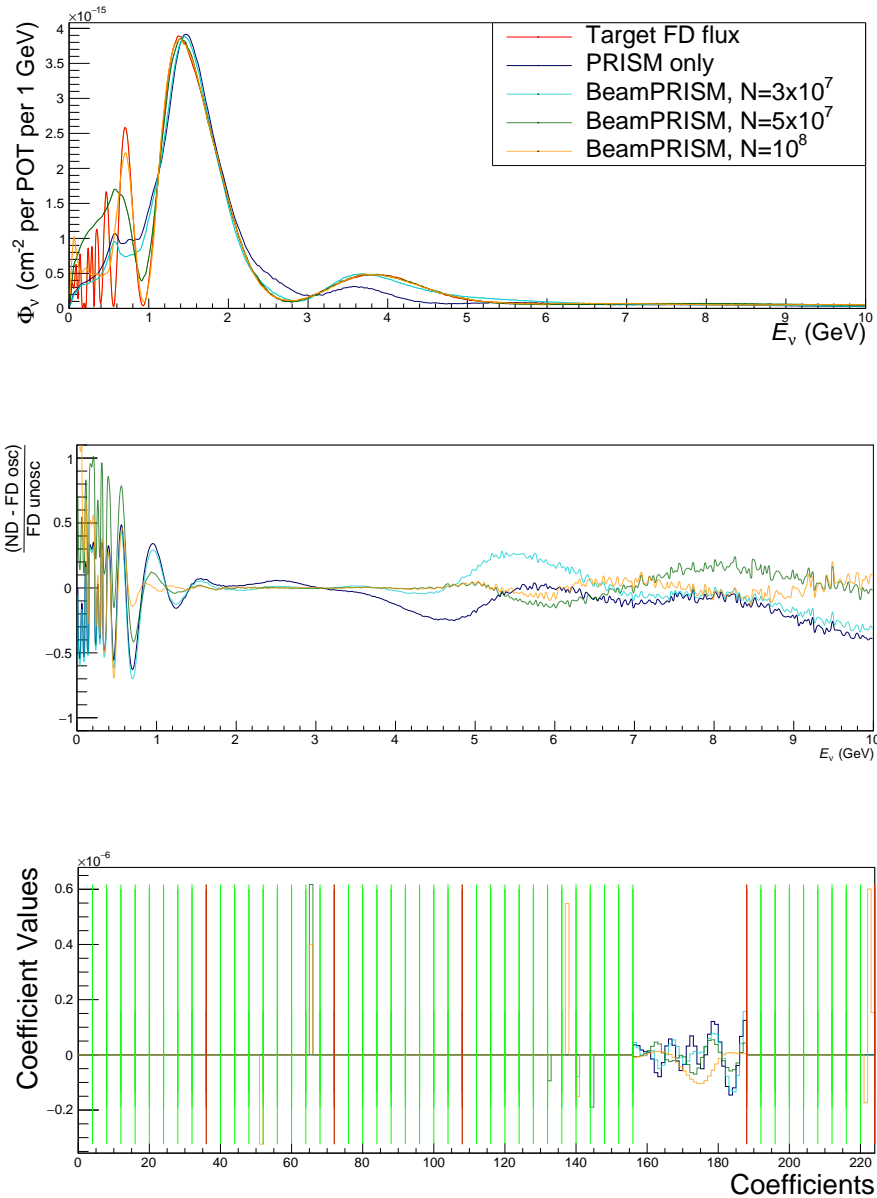


Figure 10.19: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 14.5 m off-axis position, combined with 4 m wide HC fluxes up to 12.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.

10.5 Summary

Solutions using the iteratively-reweighted Tikhonov solver show excellent performance in flux-matching to the FD oscillated flux, with equal or greater agreement than is obtained with the full off-axis range nominal DUNE-PRISM solution using any potential set of off-axis ranges. In each case, this is at the cost of the inclusion of one or more additional compensatory non-nominal fluxes, each with an associated FD data deficit.

Significant improvement to the nominal DUNE-PRISM flux-matching can be made, utilising as few as one additional HC flux in combination with nominal off-axis measurements up to 33 m. Alternatively, entirely on-axis solutions can be obtained with the inclusion of 10 or fewer non-nominal fluxes.

Whilst initial application of this method presents promising results, the imposition of such non-nominal running on FD statistics means the inclusion of these settings in an extended run-plan becomes prohibitively time-consuming. Despite this, if cross-section uncertainties remain challenging with MC corrections still required to provide a sufficiently well-matching ND-based flux, worthwhile applications in the later stages of the experiment may arise once a large enough number of FD events have been measured.

Chapter 11

Conclusions and Future Work

11.1 Summary

The work presented in chapter 10 builds on the method and tools established by the DUNE-PRISM analysis group which are themselves reviewed in chapter 9. The approach outlined here defines a tunable method with which additional useful flux information can be introduced to the DUNE-PRISM measurement program with varying degrees of sparsity, with the aim at each stage to provide the maximum increase in accurate modelling of the oscillated FD flux with the minimal number of additional flux components.

This procedure succeeds in providing equivalently performing solutions to the 33 m nominal DUNE-PRISM flux-matching using any restricted off-axis range, including an entirely on-axis solution which requires no moving near detector. In practice, these solutions are obtained at the expense of the additional run-time spent in non-nominal beam configurations, thereby impacting FD statistics, and so such an on-axis solution would be less desirable than the current DUNE-PRISM proposal.

Additionally, improvements to the baseline DUNE-PRISM program which will measure up to 33 m off-axis have been found, with superior agreement at energies above 3.5 GeV obtained from the inclusion of even a single additional flux measurement.

Following discussion with the DUNE-PRISM research group, it was clear that given the limited capability of multiple non-nominal beam flux matching to improve agreement beyond that seen with the addition of a single non-nominal flux, a solution relying on many beam-varying components would not be utilised. Whilst a combined approach succeeds in providing equal or greater agreement with the target oscillated FD flux than the nominal DUNE-PRISM for the examples shown, the reduction in overall FD event rate due to both lower luminosity and increased downtime during horn reconfiguration, as well as the additional complications induced on the long baseline analysis by such

FD flux variations would significantly impact the progress of the experiment. Instead, the nominal DUNE-PRISM method which avoids these issues has been accepted into the current design brief, with the inclusion of one additional HC flux at 280 kA to provide the large improvement shown in flux-matching over the 3 to 5 GeV energy interval [76]. However, incorporation of the 280 kA ABC horn flux into the nominal DUNE-PRISM analysis has already been undertaken, with this now posited as the baseline design due to the significant gains returned in matching the spectra across the 4 GeV oscillation minimum.

11.2 Future Work

To more deeply understand the utility of this technique in an oscillation analysis, systematic uncertainties would have to be considered. Work has already been done to understand some of these in the case of the nominal DUNE-PRISM analysis, with the currently understood sources of beam uncertainties shown in table 11.1, although focusing horn position or current uncertainties may differ due to the additional variations in running conditions such an approach would require.

To characterise the impact of these experimental uncertainties on a long-baseline analysis which incorporates a BeamPRISM ND flux combination, studies on the induced change in ND flux for each systematic uncertainty would need to be produced. These flux variations could be replicated in MC using systematically varied simulation geometries, as in section 4.5.2, assuming uncorrelated systematics. For each linearly combined solution corresponding to a given set of oscillation parameters, the same solution coefficients can then be applied to each systematically varied simulation as with the nominal simulation, returning a matched ND flux for each case. A strength of a PRISM-style analysis is that while each systematic uncertainty alters both the ND and FD flux, variations in both fluxes may cancel when comparing the two. This means the important characteristic to consider is the resulting variation in the difference between each linearly combined ND flux solution and the oscillated FD flux within the same systematically deviated simulation. The nominal DUNE-PRISM analysis has shown that large cancellations are observed when considering hadron production systematic uncertainties using the PPFX framework [76], where each nominal simulation also generates several hundred alternative simulations, each with different variations of hadron production cross sections. A similar investigation into the systematics highlighted for the BeamPRISM analysis would rely on simulating a significant number of additional MC variations to assess their impact for each included beam configuration, resulting in a prohibitively large computation time. For this reason, this study has not been conducted, though this would be the next step

Table 11.1: Sources of systematic flux uncertainty considered for the nominal-flux DUNE-PRISM analysis.

Systematic	1 Sigma Shift
Horn Current	+/- 2kA
Horn 1 X shift	+0.5mm
Horn 1 Y shift	+0.5mm
Beam Width X	+0.1mm
Decay Pipe Radius	+0.1m
Water Thickness	+0.5mm
Baffle Scraping	+0.25%
IC Skin Depth	$\infty \rightarrow 6.6\text{mm}$
Beam Shift X	+0.45mm
Target Density	+2%

in assessing the utility of this approach.

Furthermore, the preparation of such an analysis is constrained by its need to be resilient in its quality of flux-matching with a given combination of fluxes to perturbations of the neutrino oscillation parameters. Data-taking for the experiment will begin with several of these still ill-defined, and so taking non-nominal data which later proves to be unhelpful given the oscillation parameters should be avoided. To assess this, the development of an expanded method to select for the most useful flux information across many possible oscillation parameter combinations could be undertaken. With such a procedure in place, a rigorous study of the solution performance could be made given the variation of oscillation parameters through a wide parameter-space.

Further development of this analysis with less coarse discontinuities in the available current settings could increase the number of possible combinations and likely reduce the total required additional flux components for some solutions, but would lead to rapidly increased computation time as new flux information is included. This would advance an effort to find maximally sparse solutions.

Appendix A

Beam-PRISM Solutions

A.1 Restricted Off-Axis Range Solutions

A.1.1 25.5m

Restricting the off-axis interval to a maximum distance of 25.5 m shows similar performance as the 27.5 m case, with increasingly good agreement with the inclusion of more HC fluxes. Each solution introduces the same HC fluxes as for the 27.5 m case, with performance differing from the more extensive off-axis range only seen over the low energy interval below 0.6 GeV, where some linear combinations now extend into negative flux, with less smooth decay in flux as the energy falls to 0 GeV. Again, each successively less sparse solution improves the agreement at higher energy first due to the energy scaling preference included in the construction of the problem.

Comparatively, the PRISM only solution shown now fails to replicate the peak at 0.7 GeV, instead returning an extended feature which straddles this and the next lowest energy peak.

\mathcal{N}	Horn Configurations				
	A only	A+B	A+C	A+B+C	B+C
1×10^8		150 kA, 0.5 m	250 kA, 0.5 m	270 kA, 0.5 m	
5×10^7				200 kA, 0.5 m 260 kA, 0.5 m	
3×10^7				250 kA, 0.5 m	

Table A.1: Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 25.5 m.

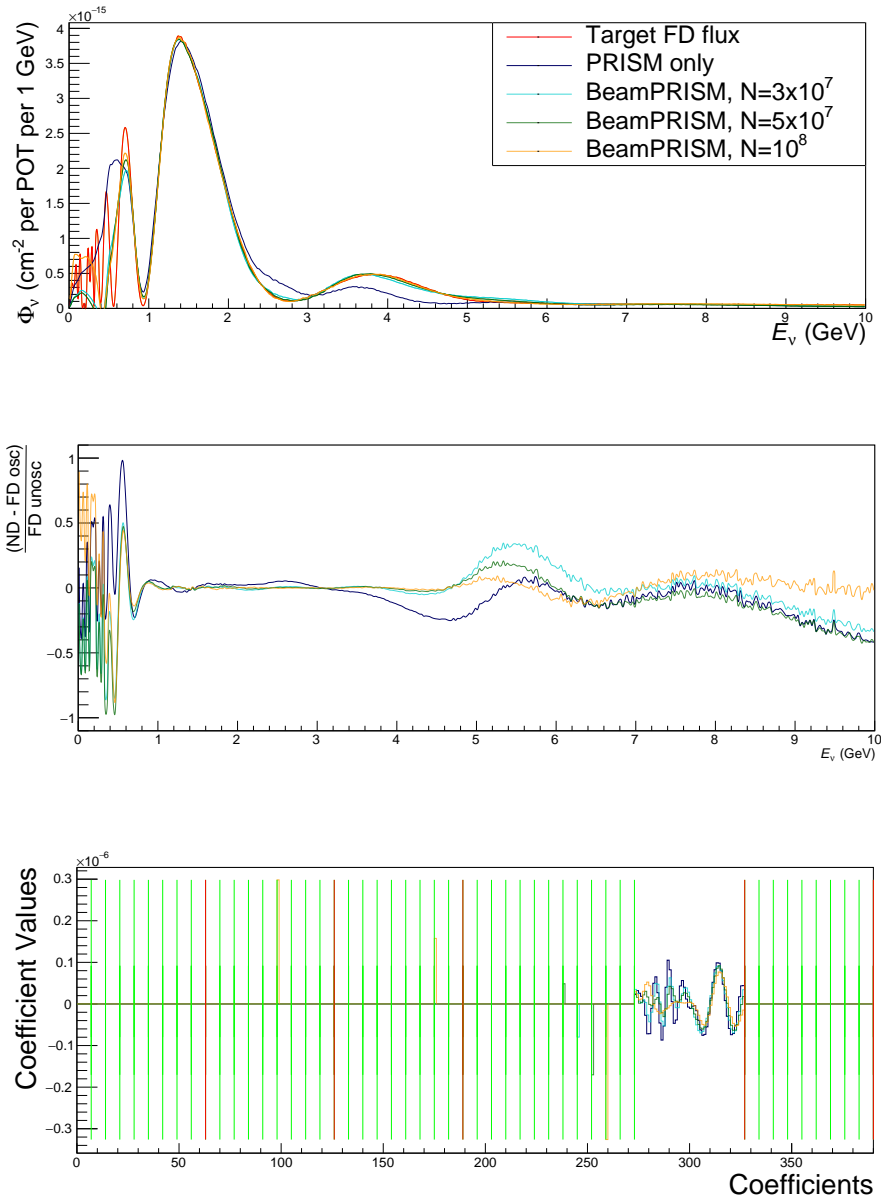


Figure A.1: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 25.5 m off-axis position, combined with 4 m wide HC fluxes up to 22.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.

23.5m

Restricting the off-axis range a further 2 m to a maximum displacement of 23.5 m shows similar results to the 25.5 m case, with the same additional HC fluxes introduced for each solution. Agreement at energies below 0.6 GeV remains inconsistent, with some negative neutrino fluxes constructed at very low energies. However the high energy matching remains consistent with more extensive off-axis range solutions, with successive improvement for each solution.

\mathcal{N}	Horn Configurations				
	A only	A+B	A+C	A+B+C	B+C
1×10^8		150 kA, 0.5 m	250 kA, 0.5 m	270 kA, 0.5 m	
5×10^7				200 kA, 0.5 m 260 kA, 0.5 m	
3×10^7				250 kA, 0.5 m	

Table A.2: Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 23.5 m.

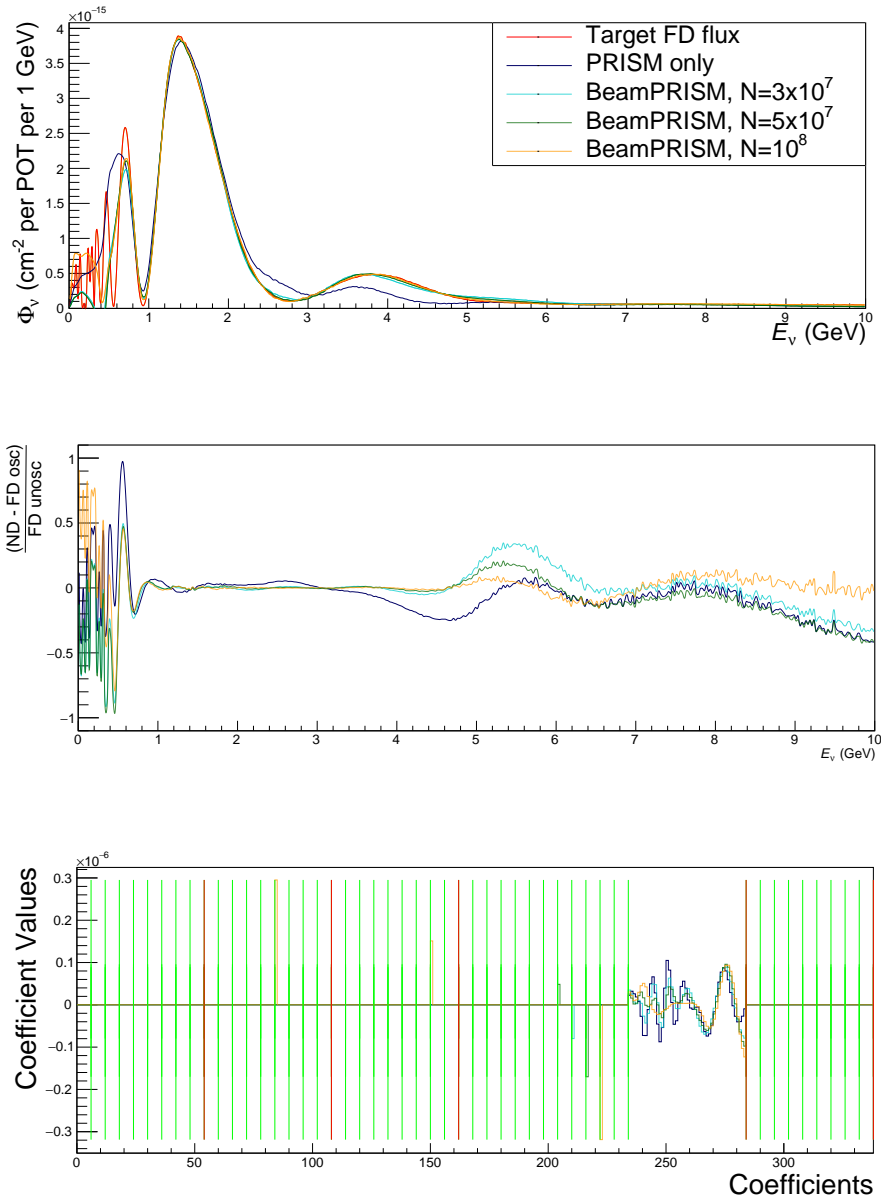


Figure A.2: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 23.5 m off-axis position, combined with 4 m wide HC fluxes up to 22.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.

18.5m

With a reduced maximum displacement of 18.5 m off-axis, a similar quality flux-matching is obtained as for the 20.5 m solution in each case. With this additional restriction, the low-energy region becomes poorly matched when including limited numbers of HC fluxes. This is corrected for in the less sparse solutions, with more HC contributions required than for the 20.5 m case.

\mathcal{N}	Horn Configurations				
	A only	A+B	A+C	A+B+C	B+C
1×10^8				200 kA, 4.5 m	293 kA, 4.5 m
				270 kA, 0.5 m	293 kA, 8.5 m
5×10^7				200 kA, 0.5 m	
				260 kA, 0.5 m	
3×10^7				250 kA, 0.5 m	

Table A.3: Horn configurations, their operating currents and the central off-axis positions of the detector utilised in each solution when restricted to a maximum off-axis distance of 18.5 m.

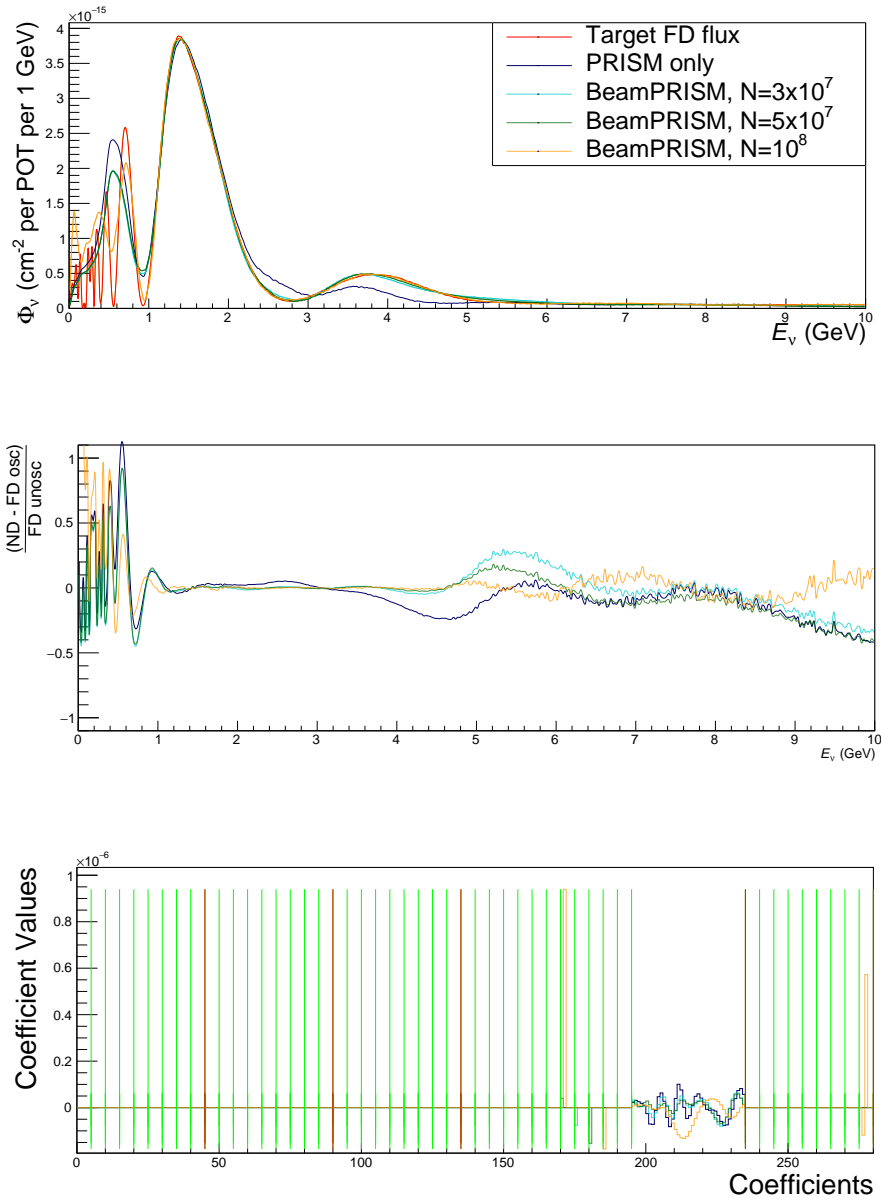


Figure A.3: Energy-weighted iterative flux-matching over a targeted energy range of 0.45-10 GeV utilising the nominal fluxes spanning -1.5 m to 18.5 m off-axis position, combined with 4 m wide HC fluxes up to 16.5 m off-axis produced from utilising horn currents up to 293 kA, including 260, 270, and 280 kA. The three panels show (top) the resulting flux-matching between (red) target oscillated FD spectrum and (green) linearly combined ND fluxes, (middle) the flux-matching residual divided by the unoscillated FD flux at each energy, and (bottom) final coefficient values resulting from a filtered solution.

Bibliography

- [1] D. Neuffer, P. Snopok, and Y. Alexahin. Front End for a neutrino factory or muon collider. *Journal of Instrumentation*, 12(11), 2017. ISSN 17480221. doi: 10.1088/1748-0221/12/11/T11007. URL <https://doi.org/10.1088/1748-0221/12/11/T11007>.
- [2] F. Foroughi, E. Morenzoni, T. Prokscha, et al. Upgrading the PSI muon facility. In *Hyperfine Interactions*, volume 138, pages 483–488. Springer, 2001. doi: 10.1023/A:1020830830050. URL <https://link.springer.com/article/10.1023/A:1020830830050>.
- [3] J. Grange, V. Guarino, P. Winter, et al. Muon (g-2) Technical Design Report. Technical report, 2015.
- [4] S Cook, R. D’Arcy, A Edmonds, et al. Delivering the world’s most intense muon beam. *Physical Review Accelerators and Beams*, 20(3), 2017. ISSN 24699888. doi: 10.1103/PhysRevAccelBeams.20.030101.
- [5] D Neuffer and V Shiltsev. On the feasibility of a pulsed 14 TeV c.m.e. muon collider in the LHC tunnel. *Journal of Instrumentation*, 13(10), 2018. ISSN 17480221. doi: 10.1088/1748-0221/13/10/T10003.
- [6] Jean Pierre Delahaye, Marcella Diemoz, Ken Long, et al. Input to the European Particle Physics Strategy Update Muon Colliders. Technical report, 2020. URL <https://muoncollider.web.cern.ch>.
- [7] IDS-NF Collaboration, R. J. Abrams, S. K. Agarwalla, et al. International Design Study for the Neutrino Factory, Interim Design Report. 2011. URL <https://arxiv.org/pdf/1112.2853.pdf><http://arxiv.org/abs/1112.2853>.
- [8] J. P. Delahaye, C M Ankenbrandt, S A Bogacz, et al. The NuMAX Long Baseline Neutrino Factory concept. *Journal of Instrumentation*, 13(6), 2018. ISSN 17480221. doi: 10.1088/1748-0221/13/06/T06003.
- [9] 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. ISBN 1507.08836v1.

- [10] M A Palmer. An Overview of the US Muon Accelerator Program. In *COOL13*, 2013. ISBN 9783954501403.
- [11] James B Rosenzweig. *Fundamentals of Beam Physics*. 2010. ISBN 9780198525547. doi: 10.1093/acprof:oso/9780198525547.001.0001. URL www.oup.com.
- [12] P A Zyla, R M Barnett, J Beringer, et al. Review of particle physics, aug 2020. ISSN 20503911. URL <https://academic.oup.com/ptep/article/2020/8/083C01/5891211>.
- [13] G Penn and J S Wurtele. Beam envelope equations for cooling of muons in solenoid fields. *Physical Review Letters*, 85(4):764–767, 2000. ISSN 00319007. doi: 10.1103/PhysRevLett.85.764.
- [14] Robert Phillips. Liouville’s theorem. *Pacific Journal of Mathematics*, 28(2):397–405, jan 1969. ISSN 0030-8730. doi: [pjm/1102983461](https://projecteuclid.org/journals/pacific-journal-of-mathematics/volume-28/issue-2/Liouvilles-theorem/pjm/1102983461). URL <https://projecteuclid.org/journals/pacific-journal-of-mathematics/volume-28/issue-2/Liouvilles-theorem/pjm/1102983461.full><https://projecteuclid.org/journals/pacific-journal-of-mathematics/volume-28/issue-2/Liouvilles-theorem/pjm/1102983461.short>.
- [15] A A Kolomenski and A N Lebedev. The effect of radiation on the motion of relativistic electrons in a synchrotron. In *Proc. CERN Symposium*, page 447, 1956.
- [16] F T Cole and F E Mills. Increasing the Phase-Space Density of High Energy Particle Beams. *Annual Review of Nuclear and Particle Science*, 31(1):295–335, 1981. ISSN 0163-8998. doi: 10.1146/annurev.ns.31.120181.001455. URL www.annualreviews.org.
- [17] G. I. Budker. An effective method of damping particle oscillations in proton and antiproton storage rings. *Soviet Atomic Energy*, 22(5):438–440, may 1967. ISSN 15738205. doi: 10.1007/BF01175204. URL <https://link.springer.com/article/10.1007/BF01175204>.
- [18] David Neuffer. Principles and applications of muon cooling. In *Los Alamos National Laboratory (Report) LA*, volume 1, pages 470–480, 1983.
- [19] David V. Neuffer and Robert B. Palmer. Progress toward a high-energy, high-luminosity $\mu^+\mu^-$ collider. In *AIP Conference Proceedings*, volume 356, pages 344–358. AIP Publishing, may 1996. doi: 10.1063/1.49615. URL <https://0-aip-scitation-org.pugwash.lib.warwick.ac.uk/doi/abs/10.1063/1.49615>.

- [20] H. Bethe. Zur Theorie des Durchgangs schneller Korpuskularstrahlen durch Materie. *Annalen der Physik*, 397(3):325–400, 1930. ISSN 15213889. doi: 10.1002/andp.19303970303.
- [21] F. Bloch. Zur Bremsung rasch bewegter Teilchen beim Durchgang durch Materie. *Annalen der Physik*, 408(3):285–320, jan 1933. ISSN 15213889. doi: 10.1002/andp.19334080303. URL <https://onlinelibrary.wiley.com/doi/full/10.1002/andp.19334080303><https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19334080303><https://onlinelibrary.wiley.com/doi/10.1002/andp.19334080303>.
- [22] Gerald R Lynch and Orin I Dahl. Approximations to multiple Coulomb scattering. *Nuclear Inst. and Methods in Physics Research, B*, 58(1):6–10, 1991. ISSN 0168583X. doi: 10.1016/0168-583X(91)95671-Y.
- [23] Diktys Stratakis, Juan C Gallardo, and Robert B Palmer. RF breakdown in magnetic fields: Previous work, recent theory, and future plans. In *AIP Conference Proceedings*, volume 1222, pages 303–307, 2010. ISBN 9780735407633. doi: 10.1063/1.3399327. URL <https://doi.org/10.1063/1.3399327>.
- [24] M. Mühlbauer, H. Daniel, F. J. Hartmann, et al. Frictional cooling: Experimental results. *Hyperfine Interactions*, 119(1-4):305–310, 1999. ISSN 03043843. doi: 10.1023/A:1012624501134. URL <https://link.springer.com/article/10.1023/A:1012624501134>.
- [25] D Taqqu. Compression and extraction of stopped muons. *Physical Review Letters*, 97(19), 2006. ISSN 00319007. doi: 10.1103/PhysRevLett.97.194801.
- [26] M Bogomilov, Y Karadzhov, D Kolev, et al. The MICE Muon Beam on ISIS and the beam-line instrumentation of the Muon Ionization Cooling Experiment. *Journal of Instrumentation*, 7(5), 2012. ISSN 17480221. doi: 10.1088/1748-0221/7/05/P05009.
- [27] D. Adams, D. Adey, R. Asfandiyarov, et al. First particle-by-particle measurement of emittance in the Muon Ionization Cooling Experiment. *European Physical Journal C*, 79(3):257–282, mar 2019. ISSN 14346052. doi: 10.1140/epjc/s10052-019-6674-y. URL <https://doi.org/10.1140/epjc/s10052-019-6674-y>.
- [28] Adam James Dobbs. *Particle Rate and Host Accelerator Beam Loss on the MICE Experiment*. PhD thesis, Imper-

- ial College London, 2011. URL http://mice.iit.edu/phd/adobbs_{_}thesis_{_}corrected_{_}2011-10-30.pdf.
- [29] Chris Booth and Sheffield Target Team. Proposal for Elimination of Dust from MICE Target Bearings. Technical report, 2007. URL <http://mice.iit.edu/mnp/MICE0180.pdf>.
- [30] C N Booth, P Hodgson, J Langlands, et al. The design and performance of an improved target for MICE. *Journal of Instrumentation*, 11(5), 2016. ISSN 17480221. doi: 10.1088/1748-0221/11/05/P05006.
- [31] H Witte, S Plate, J S Berg, et al. Partial return yoke for MICE Step IV and final Step. In *6th International Particle Accelerator Conference, IPAC 2015*, pages 2732–2734, 2015. ISBN 9783954501687.
- [32] A. Dobb, D Forrest, and F. J.P. Soler. The MICE luminosity monitor. In *Journal of Physics: Conference Series*, volume 408, 2013. doi: 10.1088/1742-6596/408/1/012084.
- [33] M Bonesini, R Bertoni, A De Bari, and G Cecchet. MICE-NOTE-DET-363 - The Refurbishing of MICE TOF0 and TOF1 detectors. Technical Report February, 2012.
- [34] François Drielsma. *Measurement of the increase in phase space density of a muon beam through ionization cooling*. PhD thesis, Université de Genève, 2018.
- [35] R Asfandiyarov, P Béné, R Bloch, et al. Electron-Muon Ranger (EMR) Construction and Tests. Technical report, 2011. URL <http://mice.iit.edu/tb/Meetings/2011-02-15-CM29/EMR.pdf>.
- [36] D. Adams, A. Alekou, M. Apollonio, et al. Electron-muon ranger: Performance in the MICE muon beam. *Journal of Instrumentation*, 10(12), 2015. ISSN 17480221. doi: 10.1088/1748-0221/10/12/P12012.
- [37] Malcolm Ellis. Mice software design and physics performance. In *Astroparticle, Particle and Space Physics, Detectors and Medical Physics Applications - Proceedings of the 10th Conference*, pages 763–770. World Scientific Publishing Co. Pte Ltd, 2008. ISBN 9812819088. doi: 10.1142/9789812819093_0126. URL <https://ui.adsabs.harvard.edu/abs/2008apsp.conf..763E/abstract>.
- [38] Rene Brun and Fons Rademakers. ROOT - An object oriented data analysis framework. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 389

- (1-2):81–86, apr 1997. ISSN 01689002. doi: 10.1016/S0168-9002(97)00048-X.
- [39] Lucien Cremaldi, David Sanders, and Don Summers. MICE-NOTE-DET-473 - Progress on Cherenkov Reconstruction in MICE. Technical report, 2015.
- [40] M. Botlo, M. Jagielski, and A. Romero. EPICS performance evaluation. Technical report, Superconducting Super Collider Project Office, sep 1993. URL <http://www.osti.gov/servlets/purl/10190661-aUjMlo/native/https://www.osti.gov/servlets/purl/10190661>.
- [41] Yordan Karadzhov. MICE Trigger System. Technical report, 2016. URL <https://launchpad.net/mice-trigger>.
- [42] A. Dobbs, C. Hunt, K. Long, et al. The reconstruction software for the MICE scintillating fibre trackers. *Journal of Instrumentation*, 11(12), 2016. ISSN 17480221. doi: 10.1088/1748-0221/11/12/T12001.
- [43] M. D. Petroff and M. G. Stapelbroek. Photon-counting solid-state photomultiplier. *IEEE Transactions on Nuclear Science*, 36(1):158–162, 1989. ISSN 15581578. doi: 10.1109/23.34424.
- [44] V. M. Abazov, B. Abbott, M. Abolins, et al. The upgraded DØ detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 565(2):463–537, sep 2006. ISSN 0168-9002. doi: 10.1016/J.NIMA.2006.05.248.
- [45] M. Ellis, P. R. Hobson, P. Kyberd, et al. The design, construction and performance of the MICE scintillating fibre trackers. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 659(1):136–153, 2011. ISSN 01689002. doi: 10.1016/j.nima.2011.04.041. URL <http://dx.doi.org/10.1016/j.nima.2011.04.041>.
- [46] M. Bogomilov, R. Tsenov, G. Vankova-Kirilova, et al. Performance of the MICE diagnostic system. *Journal of Instrumentation*, 16(08):P08046, 2021. ISSN 17480221. doi: 10.1088/1748-0221/16/08/p08046. URL <https://doi.org/10.1088/1748-0221/16/08/P08046>.
- [47] R Asfandiyarov, R Bayes, V Blackmore, et al. MAUS: The MICE analysis user software. *Journal of Instrumentation*, 14(4), 2019. ISSN 17480221. doi: 10.1088/1748-0221/14/04/T04005.
- [48] K. A. Olive, K. Agashe, C. Amsler, et al. Review of particle physics, aug 2014. ISSN 16741137. URL <https://iopscience.iop.org/>

article/10.1088/1674-1137/38/9/090001<https://iopscience.iop.org/article/10.1088/1674-1137/38/9/090001/meta>.

- [49] D. Adams, A. Alekou, M. Apollonio, et al. Pion contamination in the MICE muon beam. *Journal of Instrumentation*, 11(3), 2016. ISSN 17480221. doi: 10.1088/1748-0221/11/03/P03001.
- [50] David Neuffer. New approaches to final cooling. In *AIP Conference Proceedings*, volume 1777, page 100005, 2016. ISBN 9780735414396. doi: 10.1063/1.4965686. URL <https://doi.org/10.1063/1.4965686>.
- [51] Leon Brillouin. A theorem of larmor and its importance for electrons in magnetic fields. *Physical Review*, 67(7-8):260–266, 1945. ISSN 0031899X. doi: 10.1103/PhysRev.67.260.
- [52] J. Chadwick. The intensity distribution in the magnetic spectrum of beta particles from radium (B + C). *Verh. Phys. Gesell*, 16:383, 1914.
- [53] Wolfgang Pauli. Letter to the Gauverein meeting in Tübingen, with translation, 1930. URL <https://fermatslibrary.com/s/the-proposal-of-the-neutrino>.
- [54] Enrico Fermi. Tentativo di una Teoria Dei Raggi β . *Il Nuovo Cimento*, 11(1):1–19, sep 1934. ISSN 00296341. doi: 10.1007/BF02959820. URL <https://link.springer.com/article/10.1007/BF02959820>.
- [55] Kan Chang Wang. A suggestion on the detection of the neutrino. *Physical Review*, 61(1-2):97, jan 1942. ISSN 0031899X. doi: 10.1103/PhysRev.61.97. URL <https://0-journals-aps-org.pugwash.lib.warwick.ac.uk/pr/abstract/10.1103/PhysRev.61.97>.
- [56] C. L. Cowan, F. Reines, F. B. Harrison, et al. Detection of the free neutrino: A confirmation. *Science*, 124(3212):103–104, jul 1956. ISSN 00368075. doi: 10.1126/science.124.3212.103. URL <https://0-science-sciencemag-org.pugwash.lib.warwick.ac.uk/content/124/3212/103><https://0-science-sciencemag-org.pugwash.lib.warwick.ac.uk/content/124/3212/103.abstract>.
- [57] G. Danby, J. M. Gaillard, K. Goulianos, et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Physical Review Letters*, 9(1):36–44, 1962. ISSN 00319007. doi: 10.1103/PhysRevLett.9.36.
- [58] M. L. Perl, G. S. Abrams, A. M. Boyarski, et al. Evidence for anomalous lepton production in e^+e^- annihilation. *Physical Review Letters*, 35(22):1489–1492, 1975. ISSN 00319007. doi: 10.1103/PhysRevLett.35.1489.

- [59] J A Bagger, Martin Grünewald, S. Schael, et al. Precision electroweak measurements on the Z resonance, 2006. ISSN 03701573. URL www.elsevier.com/locate/physrepURL:<http://www.cern.ch/LEPEWWG>.
- [60] K Kodama, N Ushida, C Andreopoulos, et al. Observation of tau neutrino interactions. *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, 504(3):218–224, 2001. ISSN 03702693. doi: 10.1016/S0370-2693(01)00307-0.
- [61] Raymond Davis, Don S. Harmer, and Kenneth C Hoffman. Search for neutrinos from the sun. *Physical Review Letters*, 20(21):1205–1209, 1968. ISSN 00319007. doi: 10.1103/PhysRevLett.20.1205.
- [62] K S Hirata, T Kajita, T. Kifune, et al. Observation of B8 solar neutrinos in the Kamiokande-II detector. *Physical Review Letters*, 63(1):16–19, 1989. ISSN 00319007. doi: 10.1103/PhysRevLett.63.16.
- [63] John N. Bahcall. Solar neutrinos. I. theoretical. *Physical Review Letters*, 12(11):300–302, 1964. ISSN 00319007. doi: 10.1103/PhysRevLett.12.300.
- [64] B. Pontecorvo. Mesonium and antimesonium. *JETP Lett.*, 33:549–551, 1957.
- [65] B. Pontecorvo. Inverse β Processes and Nonconservation of Lepton Charge. Technical report, 1958. URL <https://www.osti.gov/biblio/4349231>.
- [66] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the Unified Model of Elementary Particles. *Progress of Theoretical Physics*, 28(5):870–880, 1962. ISSN 0033-068X. doi: 10.1143/ptp.28.870. URL <https://academic.oup.com/ptp/article/28/5/870/1858382>.
- [67] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Soviet Journal of Experimental and Theoretical Physics*, 26:984, 1968. ISSN 1063-7761.
- [68] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, sep 2013. ISBN 9781139525367. doi: 10.1017/cbo9781139525367.
- [69] K Abe, R Akutsu, A Ali, et al. Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations. *Nature*, 580(7803): 339–344, 2020. ISSN 14764687. doi: 10.1038/s41586-020-2177-0. URL <https://doi.org/10.1038/s41586-020-2177-0>.
- [70] M A Acero, P Adamson, L Aliaga, et al. An Improved Measurement of Neutrino Oscillation Parameters by the NOvA Experiment. 117312, 2021. URL <http://arxiv.org/abs/2108.08219>.

- [71] Peter B Denton, Julia Gehrlein, and Rebekah Pestes. CP -Violating Neutrino Nonstandard Interactions in Long-Baseline-Accelerator Data. *Physical Review Letters*, 126(5), 2021. ISSN 10797114. doi: 10.1103/PhysRevLett.126.051801. URL <http://arxiv.org/abs/2008.01110><http://dx.doi.org/10.1103/PhysRevLett.126.051801>.
- [72] B. Abi, R. Acciarri, Mario A. Acero, et al. Deep Underground Neutrino Experiment (DUNE) Far Detector Technical Design Report Volume I Introduction to DUNE. Technical report, DUNE Collaboration, 2020.
- [73] P. F. de Salas, D V Forero, S. Gariazzo, et al. 2020 global reassessment of the neutrino oscillation picture. *Journal of High Energy Physics*, 2021(2), 2021. ISSN 10298479. doi: 10.1007/JHEP02(2021)071.
- [74] Hyper-Kamiokande Proto-Collaboration, :, Ke. Abe, et al. Hyper-Kamiokande Design Report. 2018. URL <https://arxiv.org/pdf/1805.04163.pdf><http://arxiv.org/abs/1805.04163>.
- [75] DUNE Collaboration. DUNE Conceptual Design Report, Annex 3A, Optimized: Beamline at the Near Site Long-Baseline Neutrino Facility (LBNF), 2017. URL <http://docs.dunescience.org/cgi-bin/RetrieveFile?docid=1911&filename=Volume3Annex3ABeamlineattheNearSite25June2015.pdf&version=1>.
- [76] DUNE Collaboration. Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report. Technical report, 2021.
- [77] B. Abi, R. Acciarri, Mario A. Acero, et al. Deep Underground Neutrino Experiment (DUNE) Far Detector Technical Design Report Volume II DUNE Physics. Technical report, DUNE Collaboration, 2020.
- [78] S Bhadra, A Blondel, S Bordoni, et al. Letter of Intent to Construct a nuPRISM Detector in the J-PARC Neutrino Beamline. 2014. URL <http://arxiv.org/abs/1412.3086>.
- [79] Ralph A. Willoughby. Solutions of Ill-Posed Problems (A. N. Tikhonov and V. Y. Arsenin). *SIAM Review*, 21(2):266–267, apr 1979. ISSN 0036-1445. doi: 10.1137/1021044.
- [80] D. Calvetti, S. Morigi, L. Reichel, and F. Sgallari. Tikhonov regularization and the L-curve for large discrete ill-posed problems. *Journal of Computational and Applied Mathematics*, 123(1-2):423–446, nov 2000. ISSN 03770427. doi: 10.1016/S0377-0427(00)00414-3.

- [81] P C Hansen. The L-Curve and its Use in the Numerical Treatment of Inverse Problems. *in Computational Inverse Problems in Electrocardiology*, ed. P. Johnston, *Advances in Computational Bioengineering*, 4:119–142, 2000.
- [82] I Chernyshev and M Bishai. Reconstructing neutrino oscillation data using near detector results in a simulated Deep Underground Neutrino Experiment. Technical report, 2019. URL <https://docs.dunescience.org/cgi-bin/private/ShowDocument?docid=15538>.
- [83] David L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, apr 2006. ISSN 00189448. doi: 10.1109/TIT.2006.871582. URL <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp={&}arnumber=1614066>.
- [84] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006. ISSN 00103640. doi: 10.1002/cpa.20124.
- [85] Emmanuel J. Candès, Michael B. Wakin, and Stephen P. Boyd. Enhancing sparsity by reweighted l1 minimization. *Journal of Fourier Analysis and Applications*, 14(5-6):877–905, nov 2008. ISSN 10695869. doi: 10.1007/s00041-008-9045-x. URL <http://arxiv.org/abs/0711.1612>.
- [86] Ingrid Daubechies, Ronald Devore, Massimo Fornasier, and C. Sinan Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 63(1):1–38, 2010. ISSN 00103640. doi: 10.1002/cpa.20303. URL <http://www.dsp.ece.rice.edu/cs/>.