



University
of Glasgow

Forrest, David Alexander James (2011) *The Muon Ionisation Cooling Experiment*.
PhD thesis.

<http://theses.gla.ac.uk/2839/>

Copyright and moral rights for this thesis are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

The Muon Ionization Cooling Experiment

David Alexander James Forrest



University
of
Glasgow

September 2011

A thesis submitted for the degree of Doctor of Philosophy
at the College of Science & Engineering,
University of Glasgow

The Muon Ionisation Cooling Experiment

Abstract

Outstanding areas of ambiguity within our present understanding of the nature and behaviour of neutrinos warrant the construction of a dedicated future facility capable of investigating the likely parameter space for the θ_{13} mixing angle, the Dirac CP violating phase and clarifying the neutrino mass hierarchy. A number of potential discovery venues have been proposed including the beta beam, superbeam and neutrino factory accelerator facilities. Of these, the neutrino factory significantly outperforms the others.

A neutrino factory will deliver intense beams of 10^{21} neutrinos per year, produced from muons decaying in storage rings. This specification, coupled with the constraints of the short muon lifetime warrant the inclusion of a novel cooling channel to reduce the phase space volume of the beam to fall within the acceptance of the acceleration system.

Ionisation cooling is the only viable cooling technique with efficacy over the lifetime of the muon, however, it has yet to be demonstrated in practice. In a full cooling channel, a muon beam will traverse a periodic absorber and accelerator lattice consisting of low Z absorbers enclosed by focusing coils and accelerating radio-frequency cavities. Energy loss in the absorbers reduces both transverse and longitudinal momentum. The latter is restored by the accelerating cavities providing a net reduction in transverse momentum and consequently reducing the phase space volume of the muon beam.

The Muon Ionisation Cooling Experiment (MICE), under construction at the ISIS synchrotron at Rutherford Appleton Laboratory seeks to provide both a first measurement and systematic study of ionisation cooling, demonstrated within the context of a single cell prototype of a cooling channel. The experiment will evolve incrementally toward its final configuration, with construction and scientific data taking schedules proceeding in parallel. The stated goal of MICE is to measure a fractional change in emittance of order 10% to an error of 1%.

This thesis constitutes research into different aspects of MICE: design and implementa-

tion of the MICE configuration database, determination of the statistical errors and alignment tolerances associated with cooling measurements made using MICE, simulations and data analysis studying the performance of the luminosity monitor and a first analysis of MICE Step I data.

A sophisticated information management solution based on a bi-temporal relational database and web service suite has been designed, implemented and tested. This system will enable the experiment to record geometry, calibration and cabling information in addition to beam-line settings (including but not limited to magnet and target settings) and alarm handler limits. This information is essential both to provide an experimental context to the analysis user studying data at a later time and to experimenters seeking to reinstate previous settings. The database also allows corrections to be stored, for example to the geometry, whereby a later survey may clarify an incomplete description. The old and new geometries are both stored with reference to the same period of validity, indexed by the time they are added to the configuration database. This allows MICE users to recall both the best-known geometry of the experiment at a given time by default, as well as the history of what was known about the geometry as required. Such functionality is two dimensional in time, hence the choice of a bi-temporal database paradigm, enabling the collaboration to run new analyses with the most up to date knowledge of the experimental configuration and also repeat previous analyses which were based upon incomplete information.

From Step III of MICE onwards, the phase space volume, or emittance, of the beam will be measured by two scintillating fibre trackers placed before and after the cooling cell. Since the two emittance measurements are made upon a similar sample of muons, the measurement errors are influenced by correlations. This thesis will show through an empirical approach that correlations act to reduce the statistical error by an order of magnitude.

In order to meet its goals MICE must also quantify its systematic errors. A misalignment study is presented which investigates the sensitivity of the scintillating fibre trackers to translational and rotational misalignment. Tolerance limits of 1 mm and 0.3 mrad respectively allow MICE to meet the requirement that systematic errors due to misalignment of the

trackers contribute no more than 10% of the total error.

At present, MICE is in Step I of its development: building and commissioning a muon beamline which will be presented to a cooling channel in later stages of MICE. A luminosity monitor has been built and commissioned to provide a measurement of particle production from the target, normalise particle rate at all detectors and verify the physics models which will be used throughout the lifetime of MICE and onwards through to the development of a neutrino factory.

Particle identification detectors have already been installed and allow the species of particles to be distinguished according to their time of flight. This has enabled a study of particle identification, particle momenta and simulated and experimental beam profiles at each time of flight detector. The widths of the beam profiles are sensitive to multiple scattering and magnetic effects, providing an opportunity to quantify the success of the simulations in modelling these behaviours. Such a comparison was also used to detect offsets in the beam centre position which can be caused by misalignments of the detectors or relative misalignments in magnet positions causing asymmetrical skew in the magnetic axis. These effects were quantified in this analysis.

Particle identification combined with the earlier statistical analysis will be used to show that the number of muons required to meet the statistical requirements of MICE can be produced within a realistic time frame for each beam configuration considered.

Acknowledgments

This could take a while.

First of all I would like to thank Paul Soler, my supervisor, whose trust and grace throughout this adventure has been much appreciated. I thank him especially for being so generously and lavishly giving of his time for any student, no matter how busy he happens to be.

In retrospect I realise how fortunate I was to carry out my PhD research in a relatively small collaboration with limited resources but with numerous opportunities for new researchers to be granted meaningful areas of responsibility right from the start. I am grateful to MICE for every opportunity given and trust afforded to a new guy whose ‘accent’ they so clearly could not understand. Specifically Malcolm Ellis, both hyper-intelligent and hyper-kind; Adam Dobbs and Andri for making collaboration meetings far, far better; Chris Rogers for getting me going with G4MICE at the start of my PhD; JSG for fun chat; Marco for going to great lengths to understand me; Mark Rayner for promising he would show me how to measure emittance with only a toothpick and an elastic band and anyone who stayed awake for my talks the whole way through (which would appear to otherwise exclude some people who are already on this list!).

I would like to thank my parents (Hi Mum! Hi Dad!) for their great love and support. They should probably be given a PhD in something.

Throughout my time at university, it has been the most wonderful pleasure and privilege to have met and shared life with international students in Glasgow and who have helped me find answers to three of life’s fundamental questions: ‘What’s your name?’, ‘What do you study?’ and ‘Where do you come from?’. To name only a few: Marie Garde, Ilhan, Loik, Aurelie, Laurie, Max, Mario (Are we still on Map 6?), Andreaaaa, Jose, Bennet (from the republic of Scotland), Yannick, Loic, Illka & Suvi, Panu, Joseph, Simon, Wolfgang, Edu, Jeanette and the americanos, Marie, Weike, Helena, Yulia, Phillip, Matteo and the Italians (yes, the entire nation of Italy.), Daan, Tim, Badrig, Emmanuel, Hikmat, Calum (you count), Seven, Ziva, Andy, Luis and anyone I’ve ever lived with, even you Kamtoh. And to the unknown international student. You guys stayed for such a short time, but made such a

lasting impression.

There are plenty of UK folks who helped organise some of these adventures. Ruth and Dish especially are both seven different kinds of wonderful. Seven!

I'd like to thank Will Bell for some wise words at the beginning that made it all possible. Thanks Will!

Thanks are due to the Grid wizards at Glasgow (Sam, Graeme, Mike, David and Stuart), with their magical machines, whose technical expertise was invaluable in getting MICE on the Grid, achieving the number of events required for my statistics study and whose patience was appreciated when I accidentally stress tested the Grid instead of the configuration database.

The biggest thank you, more than any of these, to the God who provides, and made everything good.

Declaration

The results presented in this thesis are my own work unless otherwise stated. Chapters 1, 2 and 3 are reviews of the physics of neutrinos, a description of the baseline neutrino factory and of the Muon Ionisation Cooling Experiment (MICE) respectively. The reviews were carried out by the author, referencing from the established literature.

All work on the MICE configuration database described in chapter 4 is my own, including the requirements analysis, design, implementation, deployment, testing and user training, with clients provided by James Leaver. I performed all simulation and analysis for the statistical and misalignment studies presented in chapter 5, with a theoretical treatment provided later by John Cobb (Oxford) and included in section 5.3.1 of this thesis for completeness. The luminosity monitor described in chapter 6 was designed and built by Paul Soler (University of Glasgow) but I participated in the installation and commissioning of the device in MICE at the Rutherford Appleton Laboratory. I also performed the first analysis of the luminosity monitor data and subsequent analyses exploiting the luminosity monitor, the commissioning of which benefited from beamloss data provided by Adam Dobbs (Imperial College) which I merged manually with the luminosity monitor output. Marco Apollonio provided the G4Beamline description of the MICE geometry and the weighting algorithm for simulating particle production from the MICE target, used in chapter 7. I modified and used the geometry description to generate G4Beamline simulations comparing beam profiles at the MICE time of flight detectors to data and to normalise simulated detector readout to protons on target. The analysis in chapter 7 is completely my own work.

Appropriate references are provided when the results of others are mentioned. The research presented here was not submitted for another degree in any other department or university.

David Alexander James Forrest

Pleasure converges on wisdom

Like the sky onto the land.

Andrew Greig, *Western Swing*



'Are you finished that thesis yet?'

Various

Contents

1	Neutrino Physics	1
1.1	Neutrino History	1
1.2	Neutrino Interactions	3
1.3	Neutrino Oscillations	6
1.3.1	Neutrino Oscillations in Matter	12
1.4	Neutrino Mass	15
1.5	Review of Neutrino Experiments	16
1.5.1	Reactor Based	17
1.5.2	Solar Based	18
1.5.3	Atmospherically Based	21
1.5.4	Accelerator Based	23
1.6	Conclusions	31
2	Neutrino Factory	32
2.1	Introduction	32
2.2	The Case for a Neutrino Factory	32
2.3	Design Description	37
2.3.1	Proton Driver	37
2.3.2	Target	38
2.3.3	Bunching and Phase Rotation	39
2.3.4	Muon Cooling	41
2.3.5	Muon Acceleration	42
2.3.6	Muon Storage	44
2.3.7	Detectors	45
2.3.7.1	Near Detectors	45
2.3.7.2	Far Detectors	45
2.4	Future Work	47
2.5	Conclusions	49
3	The Muon Ionization Cooling Experiment (MICE)	50
3.1	Introduction	50
3.2	Emittance	51
3.3	Ionisation Cooling	54
3.4	Experimental Setup	57
3.5	Beamline	58
3.6	Target	60
3.6.1	Operation of the MICE Target	61

3.7	Absorbers	62
3.8	RF Cavities	64
3.9	Detectors	66
3.9.1	Luminosity Monitor	66
3.9.2	Particle Counters	67
3.9.3	Particle Identification	68
3.9.3.1	Time of Flight Detectors	68
3.9.3.2	Cherenkov Detectors (CKOV)	70
3.9.3.3	KLOE-Light Calorimeter (KL)	71
3.9.3.4	Electron Muon Ranger (EMR)	74
3.9.4	Particle Tracking	74
3.9.4.1	Scintillating Fibre Trackers	75
3.10	Proton Absorber	76
3.11	Software	77
3.11.1	Simulation & Analysis	77
3.11.1.1	G4MICE	78
3.11.1.2	G4Beamline	78
3.12	Outlook	79
3.13	Conclusions	80
4	The Configuration Database	81
4.1	Introduction	81
4.2	Relational Databases	82
4.3	Analysis and Use Cases	85
4.3.1	Beamline Settings	86
4.3.2	Geometry	87
4.3.3	Calibrations	87
4.3.4	Electronic Channel Mapping	88
4.3.5	Alarm Handler Limits	88
4.4	Design and Implementation	89
4.4.1	Bi Temporal Paradigm	89
4.4.2	Beamline Settings	91
4.4.3	Geometry	93
4.4.4	Calibrations	93
4.4.5	Electronic Channel Mapping	96
4.4.6	Alarm Handler Limits	97
4.4.7	Implementation	97
4.5	Remote Access	100
4.5.1	API	100
4.5.2	SOAP Connectivity for Remote Access	100
4.5.3	Web Interface	101
4.5.4	Security & Stability	102
4.6	Testing	103
4.7	Client-side Applications	105
4.8	Conclusions and Future Work	105

5	Statistical Errors and Alignment Tolerances of Emittance Change Measurements in MICE	108
5.1	Introduction	108
5.2	Emittance Change Through MICE	109
5.3	Statistical Errors of Fractional Change in Emittance in MICE	111
5.3.1	Theoretical Analysis of Statistical Errors and Correlations	111
5.3.2	Empirical Study of Statistical Errors in MICE	117
5.3.2.1	Use of the Grid	117
5.3.2.2	Ganga	118
5.3.2.3	G4MICE Simulations	119
5.3.2.4	Results	120
5.3.2.5	Discussion	120
5.4	MICE Tracker Alignment Tolerances	133
5.4.1	Motivation of Alignment Study	133
5.4.2	Simulation	134
5.4.3	Results	135
5.5	Conclusions	135
6	The MICE Luminosity Monitor	141
6.1	Purpose	141
6.2	Design	142
6.3	Beamloss	145
6.4	Commissioning	145
6.5	Simulations	147
6.5.1	Transmission	147
6.5.2	Geometry Factor	152
6.5.3	Particle Population	154
6.5.4	Yield	156
6.6	Saturation Limits	163
6.7	Conclusions	164
7	Step I Analysis of the MICE Beamline	166
7.1	Introduction	166
7.2	Step I Data Taking	167
7.3	Step 1 Simulations	168
7.3.1	MICE Beamline	168
7.3.2	Luminosity Monitor	170
7.4	Detector Efficiency	171
7.5	Particle Rate Per Detector	173
7.6	Particle Identification and Momentum Reconstruction	173
7.7	Predicted Performance of Step VI	178
7.8	Comparison with Simulation	179
7.8.1	Transmission Along Beamline	180
7.8.2	Simulated and Reconstructed TOF Profiles	181
7.8.2.1	Effective Misalignment	182
7.8.2.2	Profile Widths	187
7.9	Conclusions	192

8	Conclusions	193
A	Misalignment Tables for the Time of Flight Counters	197
B	Beam Profile Widths at the Time of Flight Counters	206

List of Tables

1.1	Best fit ($\pm 1\sigma$) values of the oscillation parameters. Results from both normal and inverted hierarchy assumptions are shown. Table from [67].	31
2.1	Experimental oscillation channels in a neutrino factory.	35
2.2	Proton driver parameters, taken from [78].	37
2.3	The current MIND baseline [70]	47
5.1	Estimated number of muons required with and without correlations for a range of beam conditions. Correlations in our measured sample of muons have the effect of reducing the total number of muons required by an order of magnitude.	133
5.2	Comparison between simulated and analytical results.	133
5.3	Percentage changes in emittance due to misalignment of the downstream scintillating fibre tracker.	139
6.1	Estimates of sensitivity at different points in the ISIS machine cycle. The MICE target dips into the beam at around 9 milliseconds. Table provided by Dean Adams (STFC).	147
6.2	Gradients of fitted lines from the graphs contained in figure 6.5, which provide conversion factors between luminosity counts and beamloss. The region of low beamloss < 0.5 Vms is excluded from the fit due to imprecision in the beamloss calculation. Data was taken during experimental runs 1447-1457.	147
6.3	Explanation of GEANT4[172] hadronic physics models used.	151
6.4	The percentage transmission between C12 and both C34 and C1234, calculated in simulation B. Note that the transmission may exceed 100% in C34 due to the larger area of the second set of scintillators (9 cm^2 in C34, compared with 4 cm^2 in C12).	151
6.5	Comparison of results for fin target geometry between this study and that of Walaron [169]. This study benefits from greater statistics, with 1.88×10^9 POT compared to 10^7	153
6.6	Particle populations in simulation A for the QGSP_BIC and QGSP_BERT hadronic models.	156
6.7	Particle populations in simulation A for the QGSC_CHIPS and QGSC_BERT hadronic models.	161
6.8	Particle populations in simulation A for the LHEP hadronic model.	161

6.9	The efficiency of neutron detection. For C12, this is the efficiency with which neutrons knocks off a proton in the first scintillator which then creates a coincidence in both scintillators. For C34, this is the efficiency with which a neutron knocks off a proton in either of the first two scintillators or the plastic shielding which subsequently creates a coincidence in C34. In both cases we consider only those events which cross all four scintillators.	162
6.10	Response of the luminosity monitor as a function of hadronic model. Each simulation has an identical geometry description and is given the same set of functional parameters, including 2×10^9 protons on target. Errors in data are dominated by the 30% error attributed to the beamloss measurements whilst errors in the simulation are statistical.	162
7.1	Numerical detector efficiencies taken as an average over each of the results in figure 7.4.	172
7.2	Detector rate per luminosity monitor count (C1234).	173
7.3	Number of muons per 1000 C1234 counts, for runs considered in each element of the EM matrix. The positive polarity (6,200) magnetic optics were the most efficient in transmitting muons to TOF1.	178
7.4	An estimate of the experimental running time required, at a beamloss of 2 Vms, to meet the statistical requirements of the MICE Experiment for positive polarity 200 MeV/c beams, rounded to the highest day. Though no formal calculation has been made of the error of these estimates, an error of 20% is suggested, to account for an unknown muon transmission and variations in muons per spill.	179
7.5	Effective misalignments which should be applied to MICE negative polarity MICE beamline simulations.	186
7.6	Effective misalignments which should be applied to MICE positive polarity MICE beamline simulations.	186
7.7	Weighted ratio of the width of the Monte Carlo distribution with respect to the data at each TOF plane, for a negative polarity beam.	191
7.8	Weighted ratio of the width of the Monte Carlo distribution with respect to the data at each TOF plane, for a positive polarity beam.	191
A.1	Mean values for positive polarity beam position profiles at TOF0, with no alignment corrections applied.	198
A.2	Mean values for positive polarity beam position profiles at TOF1, with no alignment corrections applied.	199
A.3	Mean values for negative polarity beam position profiles at TOF0, with no alignment corrections applied.	200
A.4	Mean values for negative polarity beam position profiles at TOF1, with no alignment corrections applied.	201
A.5	Mean values for positive polarity beam position profiles at TOF0, with corrections applied for effective misalignments.	202
A.6	Mean values for positive polarity beam position profiles at TOF1, with corrections applied for effective misalignments.	203
A.7	Mean values for negative polarity beam position profiles at TOF0, with corrections applied for effective misalignments.	204

A.8	Mean values for negative polarity beam position profiles at TOF1, with corrections applied for effective misalignments.	205
B.1	Values, in mm, of simulated and experimental standard deviations for negative polarity beam profiles at TOF0. The ratios of simulated and experimental values have been listed.	207
B.2	Values, in mm, of simulated and experimental standard deviations for negative polarity beam profiles at TOF1. The ratios of simulated and experimental values have been listed.	208
B.3	Values, in mm, of simulated and experimental standard deviations for positive polarity beam profiles at TOF0. The ratios of simulation and data have been listed.	209
B.4	Values, in mm, of simulated and experimental standard deviations for positive polarity beam profiles at TOF1. The ratios of simulation and data have been listed.	210

List of Figures

1.1	The experimentally observed energy spectrum for beta decay from a Radium-E source [1].	1
1.2	The thick arrows indicate the direction of spin whilst the thin arrows indicate the direction of the electron momentum. Figure taken from [6].	4
1.3	Rate of beta decay counts for ^{60}Co cooled to 0.01 K and normalised to a warm state. Electrons are preferentially emitted in the backward hemisphere of the ^{60}Co . Figure taken from Wu [5]	5
1.4	The Homestake experiment observed a third of the expected flux of electron neutrinos. Here 1 Solar Neutrino Unit=1 interaction per 10^{36} atoms per second. 8 SNU is predicted by the SSM. Figure taken from [14].	6
1.5	Geometrical illustration of the relationship between flavour and mass eigenstates, parametrised by the mixing angles which mediate neutrino oscillation. Figure taken from [17].	7
1.6	As yet the nature of the mass hierarchy is unknown due to ambiguities in our existing absolute mass scale measurements. Figure from [26], symbols changed to maintain consistency.	16
1.7	Schematic of the CHOOZ detector target housing with holes for PMTs. Figure from [29].	17
1.8	The standard solar model [33] describes expected terrestrial neutrino flux from nuclear reactions in the sun, assuming no neutrino oscillations.	19
1.9	Expected flux of solar neutrinos according to the SSM [33].	19
1.10	The Super-Kamiokande solar neutrino detector located 1 km underground at Mount Kamioka. Figure from [36].	20
1.11	The angular distributions of solar neutrinos detected at Super-K [37].	21
1.12	SNO results showing ν_{μ}, ν_{τ} flux plotted against ν_e flux. Dashed lines indicate the BP2000 SSM prediction for total ^8B flux [39] which compares favourably with the flux reconstructed from NC channel events (blue band). The y-axis intercepts of both bands represent $\pm 1\sigma$ uncertainties. The green band signifies the ES channel and the red band the CC channel. The intersection of the bands with the fit values for $\theta_{\mu\tau}$ and θ_e are consistent with the oscillatory neutrino behaviour. Figure from [38]	22
1.13	Zenith angle distributions for fully contained atmospheric neutrino candidates detected at Super-K. Data points are signified with black circles and Monte-Carlo simulations assuming no oscillations are provided in blue. The left graphs relate to electron-like events and the right graphs muon-like events. Figure taken from [44].	24

1.14	Parameter space showing oscillation parameter regions consistent with K2K results. Figure from [45].	26
1.15	Parameter space of $\sin^2(2\theta_{13})$ and δ_{CP} consistent with event rates seen at MINOS assuming normal mass hierarchy (top) and inverted mass hierarchy (bottom). Figure from [41].	27
1.16	The first ν_τ candidate event in OPERA. Top right and top left: views transverse to neutrino direction. Bottom: view longitudinal to neutrino direction [56].	29
2.1	A schematic of a future neutrino factory facility [70].	33
2.2	The physics reach of a range of possible future facilities. IDS-NF refers to a baseline neutrino factory described in this chapter, with a 100 kTon magnetised iron neutrino detector (to be described in section 2.3.7.2) at 4000 km from the storage rings and a similar 50 kTon detector positioned at a distance of 7500 km. BB100 refers to a Beta-Beam with $\gamma = 100$. MIND LE refers to a single baseline neutrino factory optimised for large $\sin^2 2\theta_{13} > 0.01$. The curve 2025 refers to the generation of experiments presently under construction. The remaining acronyms are consistent with those used in this thesis. Figure from [70].	34
2.3	Predicted discovery sensitivities. The left plot illustrates the $\sin^2 2\theta_{13}$ sensitivity as a function of energy and baseline, with the diamond indicating optimal sensitivity of $\sin^2 2\theta_{13} \simeq 5.0 \times 10^{-4}$ with contours showing factors of 0.5, 1, 2, 5 and 10 of optimal sensitivity. The right plot illustrates the potential for discovery of maximal CP violation with respect to $\sin^2 2\theta_{13}$. Note from the plots that some flexibility exists of order a few hundred kilometres within the optimal region around the diamond point. [72]	35
2.4	A possible design of a future beta beam facility, a rival proposal to the neutrino factory.[74]	36
2.5	The proposed updated baseline of the neutrino factory target based on a free flowing mercury jet placed within a solenoidal magnet. The acronym SC here stands for superconducting magnet. Image taken from [87].	39
2.6	The muon yield of various target materials, taking into account the neutrino factory acceptance. Plot from [88] with discussion in [77].	39
2.7	The buncher and phase rotation sections of the neutrino factory act to reduce the energy spread and increase the time spread of the muon beam [89, 90]. Figure from [70].	40
2.8	A diagram of a single cell of the phase rotator. Figure from [70].	40
2.9	Above: Relative cost with and without cooling. Below: Muons per proton within the accelerator acceptance as a function of cooling channel length for a variety of acceptance profiles [92]. The cooling channel begins at 215 m, which is the distance from the target.	43
2.10	A view of the accelerating path within a neutrino factory, taken from [70]. Radio-frequency (cavity) acceleration straights are shown in red.	44
2.11	The racetrack design of the neutrino factory muon storage rings. The top decay straight produces a neutrino beam from the decay of muons. The return straight performs acceleration, tuning and collimation. Each decay ring may store either polarity of muon. Figure from [70].	45

2.12	A single plate in the Magnetised Iron Neutrino Detector (MIND).	46
2.13	Figure by Laing [94] showing efficiency of signal identification of ν_μ (left) and $\bar{\nu}_\mu$ (right) charged current interactions as a function of energy.	47
2.14	Diagrams of (a) the neutrino factory and (b) muon Collider which illustrate their similarity upstream of the muon storage rings. The present baseline of the neutrino factory front-end is identical to that required for a muon Collider. Image sourced from [96].	48
3.1	A geometrical representation of emittance in phase space [100].	52
3.2	These graphs illustrate the principle of ionisation cooling, which MICE aims to demonstrate experimentally.	55
3.3	The staged progression of MICE is illustrated above. At time of writing, MICE is taking data in a step I configuration [104].	57
3.4	A detailed schematic of the final stage of MICE with the full cooling channel consisting of three absorber modules and two RF-cavities, braced by particle identification and tracking detectors. Flags indicate areas of responsibility [104].	58
3.5	Sketch of the MICE beamline running parasitically to the ISIS synchrotron. Beamline counters and particle identification detectors (including the luminosity monitor not shown) were used to commission the MICE muon beam.	59
3.6	The first dipole is tuned to a given pion momentum. The second dipole is tuned to the distinctive momentum of muons which decay in the backwards hemisphere of the pion centre of mass frame. Note that both forward and backward muons travel in the forward direction in the laboratory reference frame. Figure taken from [108]	60
3.7	A schematic of the MICE target body. Magnets in the shaft interact with the stator coils, causing actuation [109].	61
3.8	The target intersects the beam during the last 2 ms of the ISIS cycle.	62
3.9	MICE liquid absorber cross-section. The Beryllium safety windows are curved to withstand pressure. The absorber windows are made from 6061-T6 aluminium [111].	63
3.10	The MICE Absorber Focus Coil (AFC) module. [113].	64
3.11	Cross section showing two RF cavity modules, with four cavities per module, braced by three AFC modules [114].	64
3.12	A single prototype cavity with attached water cooling pipes. [114]	65
3.13	The design of the luminosity monitor features four low noise photo multiplier tubes in sets of two. We measure counts from coincidences of each pair of PMTs, and all four together. A block of polyethylene provides a filter for protons below 500 MeV/c and pions below 150 MeV/c.	66
3.14	Engineering diagram of the BPMs, taken from [116]. Measurements enclosed by square brackets are in units of millimetres, other measurements are in units of inches.	67
3.15	The TOF2 detector features a horizontal and vertical plane of 10 scintillator slabs. Each slab has a dual PMT readout. [117]	68
3.16	The MICE time of flight system is able to differentiate between muons, pions and electrons. Histogram by Rayner [122].	69

3.17	Left: Single CKOV detector, picture taken in direction of beam trajectory. Right: Two CKOV detectors in position between the second quadrupole triplet and TOF0.[123]	70
3.18	Cross-section of a single CKOV unit. An additional two PMTs are positioned horizontally to the beam axis.	71
3.19	The response of the two aerogels to pion, muon and electron events. A muon event in momentum region I will cause a signal in CKOV2 only. A muon event in region II will cause a signal in both CKOV detectors. Region III is not of practical interest in MICE.[124]	72
3.20	Lateral and rear view of the KL detector [125]	73
3.21	The KL under test in the MICE Hall at Rutherford Appleton Laboratory.	73
3.22	Left: Schematic of the EMR calorimeter, featuring an acceptance of 1 m^2 . Right: The triangular segmented scintillator bars are shown. These bars are arranged in 50 planes inside the EMR.	74
3.23	Above: Each tracker is comprised of five stations which each support three pairs of scintillating fibre planes. Below: Fibre optic readout is sent through a patch panel to a VLPC system.	76
3.24	Results of tracker performance study using cosmic rays showing the track residual distribution [127].	77
3.25	A visual representation of part of the MICE beamline geometry in G4Beamline [117]. Here solenoids are shown in yellow, quadrupoles in green and blue and a dipole in red. Smaller boxes and planes denote particle counters and detectors useful for tuning the simulation.	79
4.1	An analysis of the data flow of the MICE Experiment[139]. The area we are concerned with here relates to the configuration database and interface only. These are highlighted in purple. This diagram is included to illustrate that the database is integral to data flow in MICE, and that other systems exist to handle non-configuration data.	82
4.2	An example of a relational database containing four relations: Student, Lecturer, Course and StudentsStudyingCourse. A one to many relationship exists between Course and Lecturer and a many to many relationship exists between Student and Course. Red keys indicate the primary key for that relation. A green key indicates that the attribute is a foreign key from another relation. A blue key, not featured here, would indicate that a foreign key is also part of the primary key for that relation.	84
4.3	Here separate database records are represented by different coloured lines which have the same valid time period. Each record is added to the database at a different time, resulting in a different transaction time. These two time dimensions are orthogonal to one another. In the case of geometry information, valid times for different records representing the same module perfectly overlap in valid time, whereas calibrations may partially intersect one another in valid time.	90

4.4	Run meta data is tied to set values for target, magnet, ISIS and DAQ settings. Each run has one row of the Target, ISIS and DAQ tables associated with it and N rows of the magnet table where N is equal to the number of magnet modules deployed in MICE, nominally 12 (9 quadrupoles, 2 dipoles, 1 decay solenoid from the beamline) in step 1, increasing when magnets are added to the cooling channel.	92
4.5	The settings associated with a tagged run are a reduced set of those available for a real run. We do not wish to allow the user to automatically apply all the settings which are normally recorded with the set values associated with a real run. For example, the target is a critical piece of MICE equipment on which experimental running depends and target settings are the key driver of radioactive activity in the hall. Therefore target settings are an example of values which must remain under human expert rather than automated control.	92
4.6	The geometry of the experiment is modelled as a hierarchy of modules with parameters described above. Each module has knowledge of its parent module through the link to its unique identifier. A parent id of 0 defines the mother volume. The module structure was developed to mimic the geometry structure in G4MICE rather than develop a new database specific model. . .	94
4.7	Support is given for calibrations comprising arbitrary length and potentially inhomogeneous lists of parameters. The relationship of parameters with parameter sets follows a "many to one" correspondence, and the description of the type of parameter is provided by the paramType attribute.	95
4.8	An approach taken without the benefit of the flat file channel mapping strategy. Detectors are considered roots of cabling. Here the KL, TOF, Geneva and Tracker cabling are illustrated. The diagram quickly becomes more and more complicated and difficult to query and maintain with more resolution and greater number of detectors.	98
4.9	The flat file strategy. Here we have united the representation of the electronic channel mapping for all detectors, increasing maintainability and reducing complexity. The flat file is stored as a byte array and may be particular to some detector without introducing constraints on the rest of the table which remains sufficiently general as to be useful for all detectors.	99
4.10	Every alarm is tied to a process variable with a unique name, 'pvname' in the Alarm table shown. An alarm structure includes low and high limits and very low and very high limits indicating the severity of the alarm. The AlarmHandler table allows us to recall the set of alarms which were in force at a given time.	99
4.11	Web interface for the Configuration Database, allowing lookup of previous run conditions, review of tagged presets, a summary of all runs previously undertaken, and specialist queries from developers and expert users.	102
4.12	50 processors sent constant synchronous requests to the CDS for a sustained period of one minute. Response times are distributed here. These response times refer to read requests which formed the basis of our background traffic. The success or failure of the test would be determined on the time taken for a write request from a separate CPU simulating control room processes. This was found to be 0.110s on average in the presence of traffic and 0.109s without.	104

4.13	Manual interface for entering set values into the CDS. The majority of parameters are automated. The remainder require manual user entry in the current version of the client-side application. The number of parameters available for manual user entry should reduce as a function of time.	106
4.14	An EPICS application to reinstate settings from a previous run or associated with a preset string tag.	106
5.1	Emittance as a function of z -position in MICE Step VI, where z is along the axis of the beam. The central absorber is positioned at $z = 0$. Although deviations exist partly due to multiple scattering and limitations in the simulation, it can be shown that 4D transverse emittance reduces in steps along the MICE channel.	110
5.2	Fractional change in emittance as a function of input beam emittance. . . .	110
5.3	The structure of a job in the Ganga Job Submission software. Abstraction over application and backend afford the system generality over many different applications and systems (including Grid and batch systems). Splitters and mergers allow for the ready definition of composite jobs. Image taken from the Ganga project page [165].	119
5.4	Distributions of fractional change in emittance for a 0.2 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.	121
5.5	Distributions of fractional change in emittance for a 1.5 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.	122
5.6	Distributions of fractional change in emittance for a 2.5 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.	123
5.7	Distributions of fractional change in emittance for a 3.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.	124
5.8	Distributions of fractional change in emittance for a 4.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.	125
5.9	Distributions of fractional change in emittance for a 6.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.	126
5.10	Distributions of fractional change in emittance for a 8.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.	127

5.11	Distributions of fractional change in emittance for a 10.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.	128
5.12	For each beam, we plot σ_f (taken as σ in the plots) against $\frac{1}{\sqrt{N}}$ to find a constant of proportionality, K.	129
5.13	The constant of proportionality K between σ and $\frac{1}{\sqrt{N}}$ as a function of input beam size. The region of $K > 1$, where we make gains due to the effect of correlations, is highlighted in blue.	130
5.14	The correlation factor k_c due to measuring a similar sample of muons in two places. This correlation factor is included in the constant of proportionality K	130
5.15	The constant of proportionality relating the simulated error in fractional change in emittance to the inverse square root of the number of events is plotted and compared to an analytical solution provided by Cobb[160]. Agreement is good except in the case of the very small 0.2 mm rad beam.	132
5.16	Plots shown apply to misalignments using the 2.5 mm rad beam. The input beam is sufficiently small that multiple scattering dominates and the beam experiences fractional increase in emittance, or ‘heating’. Green markers indicate positive misalignment, blue markers indicate negative misalignment and black markers indicate nominal emittance with no misalignment.	136
5.17	Plots shown apply to misalignments using the 6 mm rad beam. Green markers indicate positive misalignment, blue markers indicate negative misalignment and black markers indicate nominal emittance with no misalignment. In this beam we see significant cooling.	137
5.18	Plots shown apply to misalignments using the 8 mm rad beam. Green markers indicate positive misalignment, blue markers indicate negative misalignment and black markers indicate nominal emittance with no misalignment. In this beam we see significant cooling.	138
6.1	The design of the luminosity monitor features four low noise photo multiplier tubes in sets of two. We measure counts from coincidences of each pair of PMTs, and all four together. A block of polyethylene provides a filter for protons below 500 MeV/c and pions below 150 MeV/c.	143
6.2	Scaler readout channels available to the luminosity monitor. The data from each run includes scaler counts of the coincidences of PMT sets 1 and 2 (C12), 3 and 4 (C34) and the fourfold coincidence of PMTs 1,2,3 and 4 (C1234).	144
6.3	The luminosity monitor is positioned ten metres from the target within the ISIS vault. The red line follows the path of the ISIS ring. Original figure provided courtesy of Geoff Barber (RAL).	144
6.4	Top: Taken from [171]. The target intersects the ISIS beam at the last millisecond of its acceleration cycle. Bottom: Beamloss as a function of time due to target oscillations, measured by ionisation chambers in ISIS. The large peak is caused by the dipping of the MICE Target.	146

6.5	First data from the luminosity monitor showing a direct relationship between luminosity and beamloss on all three scaler output channels. Graphs show counts for coincidences of signals for PMTs (a) 1,2 (b) 3,4 and (c) 1,2,3 and 4. The straight lines indicate a linear fit to the data between 0.5 Vms and 3 Vms, since the method for calculating beamloss is unreliable below 0.5 Vms.	148
6.6	Above: A top-down view of the active simulation geometry with the LM model positioned 10 m and 25 degrees from the target. Below: A magnification of the LM model.	149
6.7	The x positions of particles at the first scintillator, which is at an angle of 25 degrees to the x-axis. The area of the first scintillator has been enlarged by a factor of 100 however the density of particle hits remains constant.	150
6.8	A proton beam containing 10^7 particles inbound upon a model of the previous (left) and current (right) MICE target. The previous, thicker target produces a greater number of proton interactions than the new target. Here protons are coloured blue, neutrons red and pions green. Figure produced using G4Beamline.	153
6.9	The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the QGSP_BIC hadronic model.	155
6.10	The momentum distributions at the first set of scintillators for those events which create hits in all four scintillators (Channel C1234). Particle species is recorded at the first set of scintillators. Note the effect of the polyethylene in filtering the protons below approximately 570 MeV.	156
6.11	The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the QGSP_BERT hadronic model.	157
6.12	The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the LHEP hadronic model. The momentum distributions provided by this parametrised model appear to handle higher momentum protons particularly poorly and predict an overall reduced rate inconsistent with data.	158
6.13	The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the QGSC_BERT hadronic model.	159
6.14	The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the QGSC_CHIPS hadronic model.	160
6.15	Luminosity counts against beamloss for a set of high beamloss runs with DAQ gate width of 0.5 ms. Saturation is clearly evident from around 4 Vms.	165
7.1	The emittance-momentum (ϵ,p) matrix. For each emittance and momentum, a set of beamline optics has been defined. α and β are the familiar Twiss parameters and t defines the thickness of lead diffuser, used to inflate the emittance. P_{diff} is the momentum immediately before the diffuser [174].	167
7.2	Schematic of the beamline showing the positions of each of the magnets and particle detectors relative to ISIS.	168
7.3	Top: The momentum spectrum of positively and negatively charged pions leaving the target box. Bottom: The same plot, excluding particles which do not create events in TOF1.	170

7.4	Detector efficiency as a function of average C1234 counts. For each run, an average luminosity (C1234) is calculated from each spill. The graph shows efficiency remains flat as a function of luminosity for each detector considered.	171
7.5	The efficiencies of individual BPM planes for BPM1 (top) and BPM2 (bottom). The vertical plane (bottom left) of BPM2 operates with a greater efficiency than the horizontal plane (bottom right).	172
7.6	Particle rate for each detector. Detector rates are calculated by taking the average over all spills. Each data point corresponds to one run.	174
7.7	The reconstructed time of flight for particles travelling between TOF0 and TOF1.	175
7.8	Momentum distributions for muons and pions at TOF0 (top) and TOF1 (bottom). The number of events, mean and RMS figures for the muon distribution are shown. The pion contamination, coloured in red, is very small and spread over a large range of momenta. As such it is very difficult to observe from the plots.	176
7.9	Momentum distribution between TOF0 and TOF1 for particles identified as muons in a (6,200) positive polarity beam.	177
7.10	Rate through the MICE Beamline normalised to C12 LM counts, for experiment and simulation, for the positive polarity (6,200) beam.	180
7.11	Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF0 for the negative polarity (6,200) matrix element.	182
7.12	Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF1 for the negative polarity (6,200) matrix element.	183
7.13	Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF0 for the positive polarity (6,200) matrix element.	184
7.14	Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF1 for the positive polarity (6,200) matrix element.	185
7.15	Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF0 for the negative polarity (6,200) matrix element with effective misalignments applied. The mechanical misalignment is from reference [177].	187
7.16	Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF1 for the negative polarity (6,200) matrix element with effective misalignments applied. The mechanical misalignment is from reference [177].	188
7.17	Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF0 for the positive polarity (6,200) matrix element with effective misalignments applied. The mechanical misalignment is from reference [177].	189
7.18	Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF1 for the positive polarity (6,200) matrix element with effective misalignments applied. The mechanical misalignment is from reference [177].	190

Chapter 1

Neutrino Physics

1.1 Neutrino History

By 1930 results from an increasing number of studies of beta decay for various radioactive elements appeared to place the fundamental physical law of conservation of energy under threat. At the time, theoretical studies of beta decay implied the production of a two body decay characterised by discrete electron energy levels. Nevertheless a continuous electron energy spectrum was observed (figure 1.1).

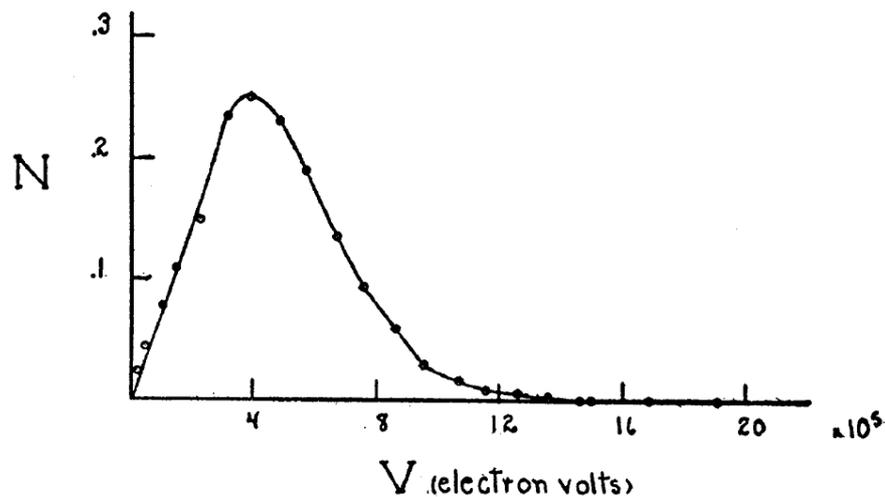


Figure 1.1: The experimentally observed energy spectrum for beta decay from a Radium-E source [1].

Pauli, in a now infamous letter to a conference in Tübingen, ventured to attribute the

phenomenon to a hitherto unobserved third particle [2]:

Dear Radioactive Ladies and Gentlemen,

As the bearer of these lines, to whom I graciously ask you to listen, will explain to you in more detail, because of the "wrong" statistics of the N and Li6 nuclei and the continuous beta spectrum, I have hit upon a desperate remedy to save the "exchange theorem" of statistics and the law of conservation of energy. Namely, the possibility that there could exist in the nuclei electrically neutral particles, that I wish to call neutrons, which have spin $\frac{1}{2}$ and obey the exclusion principle and which further differ from light quanta in that they do not travel with the velocity of light. The mass of the neutrons should be of the same order of magnitude as the electron mass and in any event not larger than 0.01 proton masses. The continuous beta spectrum would then become understandable by the assumption that in beta decay a neutron is emitted in addition to the electron such that the sum of the energies of the neutron and the electron is constant...

I agree that my remedy could seem incredible because one should have seen these neutrons much earlier if they really exist. But only the one who dares can win and the difficult situation, due to the continuous structure of the beta spectrum, is highlighted by a remark of my honoured predecessor, Mr Debye, who told me recently in Bruxelles: "Oh, It's better not to think about this at all, like new taxes". From now on, every solution to the issue must be discussed. Thus, dear radioactive people, look and judge.

Unfortunately, I cannot appear in Tübingen personally since I am indispensable here in Zürich because of a ball on the night of 6/7 December. With my best regards to you, and also to Mr Back.

Your humble servant,

W. Pauli

Although entangled here with the concept of the neutron, which was discovered by Chadwick in 1932 thus resolving the nuclear spin problem, Pauli's explanation of the continuous beta spectrum was later validated and augmented by the work of Enrico Fermi on weak forces. Fermi correctly described beta decay as the decay of a neutron producing a proton, electron and a third particle possessing $\frac{1}{2}$ spin and very slight mass much less than that of the neutron. Fermi's work clarified Pauli's explanation of a three body decay and Pauli's particle was christened the 'neutrino' by Fermi.

In 1934 Bethe and Peierls [3] first calculated the cross section of the inverse beta decay, $\bar{\nu}_e + p \rightarrow n + e^+$, to be $\sigma = 10^{-44} \text{ cm}^2$. This result suggested that the antineutrino was beyond the reach of contemporary sources at the time. The antineutrino was finally observed in 1956 by Frederick Reines and Clyde Cowan at an experiment sited near the Savannah River nuclear reactor in California [4]. Reines received the Nobel prize on behalf of both pioneers in 1995.

1.2 Neutrino Interactions

Also in 1956, Chien-Shiung Wu demonstrated that weak interactions maximally violated parity conservation [5]. Until that time, parity conservation was thought to be universal, having been applied successfully to interactions governed by both the strong and electromagnetic forces.

Parity is a spatial symmetry transformation about the origin whereby $x \rightarrow -x$. A parity value of -1 indicates a change of sign when the parity operation is imposed, for example, velocity (v) has a parity value of -1 since a transformation $(x,y,z) \rightarrow (-x,-y,-z)$ requires that $v \rightarrow -v$. Angular velocity however has a parity of $+1$ since it is the product of both a gradient operator and the velocity, both of which change sign under parity. Parity conservation holds that a process must possess the same probability as its parity counterpart or 'mirror' process. Wu performed a parity transformation on polarised ^{60}Co , placed inside a solenoid, through inversion of the magnetic field. The ^{60}Co was cooled to 0.01 K and decayed to

nickel-60 through the familiar beta-decay process:



Parity inversion reversed the electron momenta whilst leaving the spin unchanged, as shown in figure 1.2. If weak interactions were to exhibit symmetry with respect to parity then electrons should be emitted equally in the backward and forward nuclear hemispheres. However as shown in figure 1.3 this was not the case. Wu demonstrated maximal parity violation, meaning that the mirror process cannot occur. Parity violation was now established as a real and significant effect.

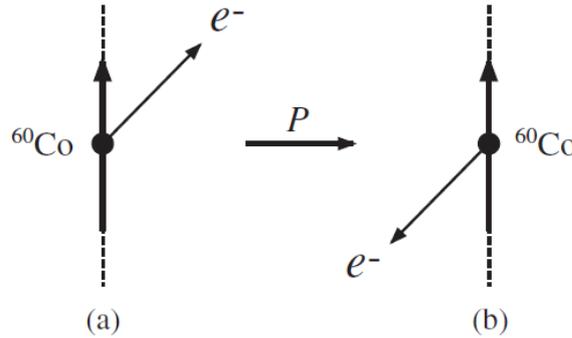
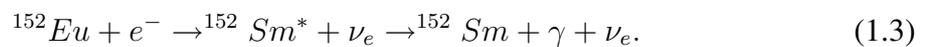


Figure 1.2: The thick arrows indicate the direction of spin whilst the thin arrows indicate the direction of the electron momentum. Figure taken from [6].

Studies of the decay by orbital electron capture of ${}^{152}\text{Eu}$ by Goldhaber et al [7] resulted in a measure of the neutrino helicity. We define helicity as the projection of spin in the direction of momentum as follows:

$$\mathcal{H} = \frac{\sigma \cdot \mathbf{p}}{|\mathbf{p}|}. \quad (1.2)$$

Where \mathcal{H} is the helicity, \mathbf{p} the momentum vector and σ the spin vector. The decay process is:



Momentum conservation requires that the ${}^{152}\text{Sm}^*$ and ν_e momenta are anti-parallel. γ

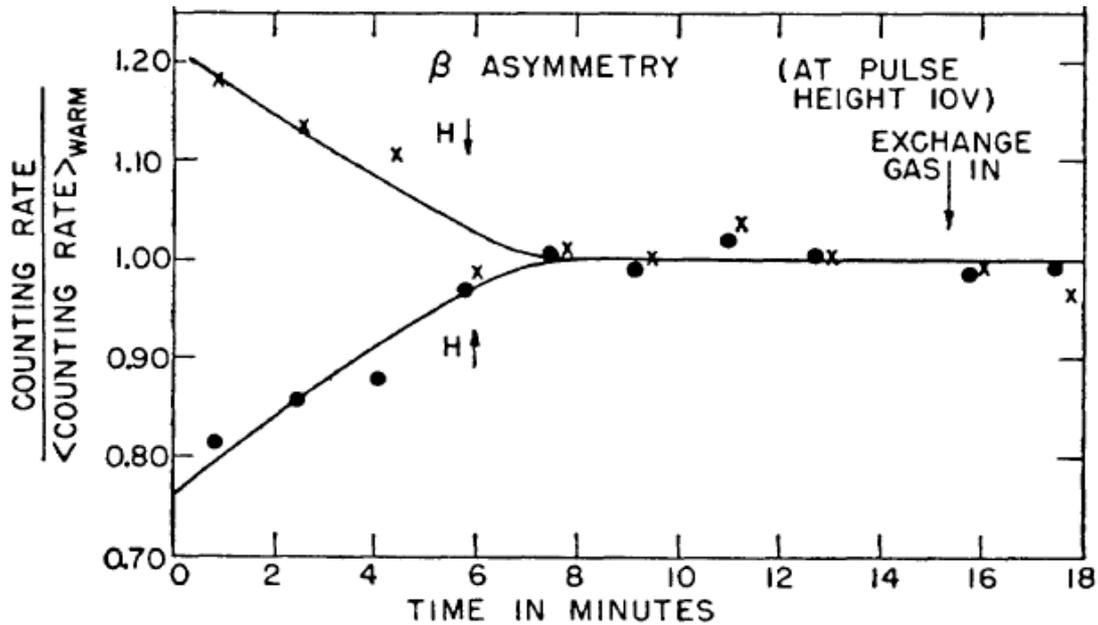


Figure 1.3: Rate of beta decay counts for ^{60}Co cooled to 0.01 K and normalised to a warm state. Electrons are preferentially emitted in the backward hemisphere of the ^{60}Co . Figure taken from Wu [5]

emission arrests the motion of the ^{152}Sm such that $\mathbf{p}_\gamma = -\mathbf{p}_{\nu_e}$, allowing indirect measurement of the neutrino momentum. The electron spin places a requirement that the sum of the photon and neutrino spin equate to $\pm\frac{1}{2}$. Given that the photon spin is ± 1 and the neutrino spin $\pm\frac{1}{2}$, their spins must be anti-parallel. This suggests that $\mathcal{H}(\nu) = \mathcal{H}(\gamma) = -1$. Experimentally only left handed (meaning, anti-parallel momentum and spin such that $\mathcal{H} = -1$) neutrinos and right handed anti-neutrinos (with parallel momentum and spin) have been observed.

These findings are explained by the V-A (vector minus axial) theory published in 1958 and developed by Feynman and Gell-Mann [8] which proposes that weak interactions operate only on left handed particles, thus explaining the maximal parity violation observed by Wu and others.

By 1959, neutrinos had been observed both in beta decay (ν_e) and pion decay (ν_μ). Pontecorvo was the first to question whether the neutrinos emitted in each process were indeed the same particle, noting that lepton flavour violating processes were not observed [9]. This was confirmed by experimental work at the Brookhaven Alternating Gradient Synchrotron

(AGS) in which a high purity ν_μ beam was shown to produce muons but not electrons [10], demonstrating that ν_μ and ν_e are distinct states and earning a Nobel prize for Lederman, Schwartz and Steinberger in 1988. In 1975 the τ lepton was discovered at the Stanford Linear Collider [11], and twenty five years later a third neutrino flavour, the ν_τ , was observed in the Direct Observation of Nu-Tau (DONUT) experiment at the Fermilab Tevatron facility [12].

1.3 Neutrino Oscillations

Since neutrinos participate only in weak interactions their cross-sections are vanishingly small. Despite their great flux on the earth's surface of 10^{11} per square centimetre per second (solar neutrinos), they remain extremely difficult to detect. It is known from observations of the invisible decay width of the Z^0 boson at the Large Electron Positron collider that the number of light neutrino flavour eigenstates (where a neutrino is classified as light if $m_\nu < \frac{m_{Z^0}}{2}$) is 2.9841 ± 0.0083 [13].

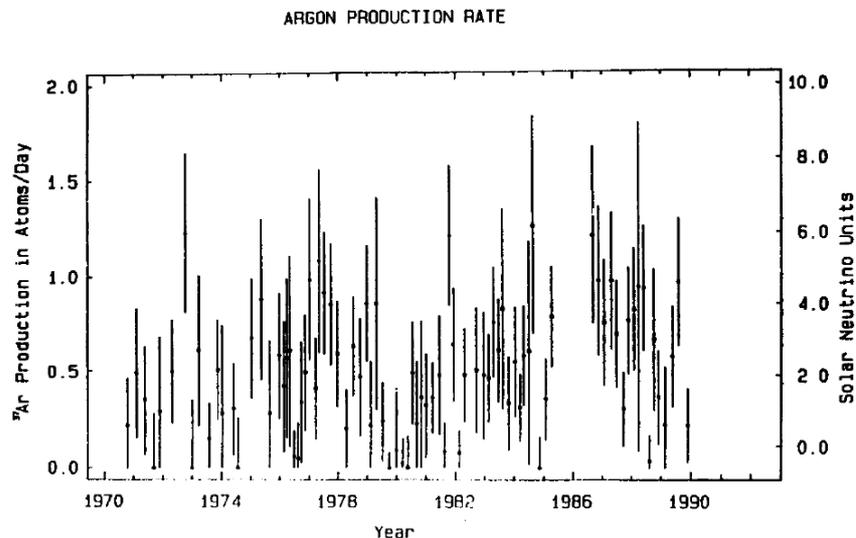


Figure 1.4: The Homestake experiment observed a third of the expected flux of electron neutrinos. Here 1 Solar Neutrino Unit=1 interaction per 10^{36} atoms per second. 8 SNU is predicted by the SSM. Figure taken from [14].

In 1968 Ray Davis devised the first solar neutrino experiment, a 100,000 gallon tank of perchloroethylene built in a mine in Lead, South Dakota. Electron neutrino collisions

convert chlorine atoms to argon through the inverse beta process, with a threshold energy of 0.814 MeV, allowing Davis to measure neutrino collisions by bubbling helium through the experiment and studying the amount of collected argon. Its observation of a third of the neutrino flux predicted by the successful Standard Solar Model (SSM) (figure 1.4) was initially controversial [14]. Later measurements by the Kamiokande and SuperKamiokande experiments confirmed the result [15]. The Sudbury Neutrino Observatory (SNO), described in detail in section 1.5.2, was sensitive to all three flavours of neutrinos and demonstrated that the total number of neutrinos predicted by the SSM was correct, however the fractions apportioned to each flavour were modified [16]. This provided clear evidence that neutrinos were able to change flavour, or 'oscillate' over solar distances.

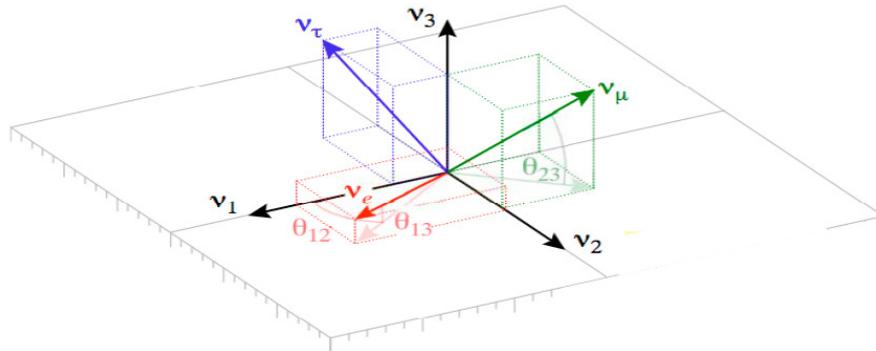


Figure 1.5: Geometrical illustration of the relationship between flavour and mass eigenstates, parametrised by the mixing angles which mediate neutrino oscillation. Figure taken from [17].

The quantum mechanical relationship between flavour states is shown in figure 1.5 and expressed:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle, \quad (1.4)$$

where ν_α is a designated neutrino flavour state ($\alpha = e, \mu, \tau$) and ν_i denotes a neutrino mass state ($i = 1, 2, 3$). $U_{\alpha i}$ is the Pontecorvo-Maki-Nagawa-Sakata (PMNS) matrix [18, 19], which must be unitary and is similar in its parametrised form to the Cabibbo-Kobayashi-Maskawa (CKM) quark mixing matrix.

$$U = \begin{pmatrix} c_{13}c_{12} & c_{13}s_{12} & s_{13}e^{-i\delta} \\ -c_{23}s_{12} - s_{13}s_{23}c_{12}e^{i\delta} & c_{23}c_{12} - s_{13}s_{23}s_{12}e^{i\delta} & c_{13}s_{23} \\ s_{23}s_{12} - s_{13}c_{23}c_{12}e^{i\delta} & -s_{23}c_{12} - s_{13}c_{23}s_{12}e^{i\delta} & c_{13}c_{23} \end{pmatrix}. \quad (1.5)$$

Subscripts indicate mixing angles, for example $s_{23} = \sin(\theta_{23})$, $c_{13} = \cos(\theta_{13})$. The phase δ is responsible for CP violation. δ must be non-zero (or not equal to π) and θ_{13} must be non-zero if CP violation effects are to be measured. Any oscillation of neutrino flavour states invokes a rotation with respect to mass eigenstates, expressed diagrammatically in figure 1.5.

The stationary mass states propagate in time according to:

$$|\nu_i(x, t)\rangle = e^{-iE_i t} |\nu_i(x, 0)\rangle \quad (1.6)$$

$$|\nu_i(x, 0)\rangle = e^{ipx} |\nu_i\rangle. \quad (1.7)$$

Here p and E denote momentum and energy in the laboratory reference frame. The weak eigenstates (ν_α) correspond to flavour states and not mass states. These are written:

$$|\nu_\alpha(x, t)\rangle = \sum_i U_{\alpha i} e^{iE_i t} |\nu_i\rangle = \sum_{i, \beta} U_{\alpha i} U_{\beta i}^* e^{ipx} e^{-iE_i t} |\nu_\beta\rangle \quad (1.8)$$

Assuming neutrino mass states are different, this then requires variation in the phase of equation 1.8 and consequently observations of flavour oscillations over large distances. The transition amplitude of such an oscillation is given by:

$$A(\alpha \rightarrow \beta)(t) = \langle \nu_\beta | \nu_\alpha(x, t) \rangle = \sum_i U_{\alpha i} U_{\beta i}^* e^{ipx} e^{-iE_i t}. \quad (1.9)$$

In the relativistic limit we can set the speed of light $c = 1$ and approximate $L \approx t$, $E \approx p$ where L is the distance between source and detector and E is the observed neutrino energy. Assuming light neutrinos, $p \gg m_i$, therefore $E_i = \sqrt{m_i^2 + p^2} \approx p + \frac{m_i^2}{2p} \approx E + \frac{m_i^2}{2E}$.

We can now write:

$$A(\alpha \rightarrow \beta)(L) = \sum_i U_{\alpha i} U_{\beta i}^* e^{\frac{-im_i^2 L}{2E}} \quad (1.10)$$

In the case of antineutrinos,

$$|\bar{\nu}_\alpha\rangle = \sum_i U_{\alpha i}^* |\bar{\nu}_i\rangle \quad (1.11)$$

$$A(\bar{\alpha} \rightarrow \bar{\beta})(t) = \sum_i U_{\beta i} U_{\alpha i}^* e^{ipx} e^{-iE_i t} \quad (1.12)$$

$$= \sum_i U_{\beta i} U_{\alpha i}^* e^{\frac{-im_i^2 t}{2E}}. \quad (1.13)$$

The oscillation probability is given by equation 1.15.

$$P(\alpha \rightarrow \beta)(t) = |A(\alpha \rightarrow \beta)|^2 = \sum_i \sum_j U_{\alpha i} U_{\alpha j}^* U_{\beta i}^* U_{\beta j} e^{-i(E_i - E_j)t} \quad (1.14)$$

$$= \sum_i |U_{\alpha i} U_{\beta i}^*|^2 + 2\text{Re} \sum_{j>i} U_{\alpha i} U_{\alpha j}^* U_{\beta i}^* U_{\beta j} e^{\frac{-i\Delta m_{ij}^2 L}{2E}} \quad (1.15)$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$. We see from equation 1.15 that if at least one mass state is non-zero then oscillations will occur. The second term of equation 1.15 accounts for the time or space dependency of neutrino oscillations. The first term describes an average transition probability:

$$\langle P_{\alpha \rightarrow \beta} \rangle = \langle P_{\beta \rightarrow \alpha} \rangle = \sum_i |U_{\alpha i} U_{\beta i}^*|^2 \quad (1.16)$$

Through applying CP invariance we may simplify equation 1.15:

$$P(\alpha \rightarrow \beta) = \sum_i U_{\alpha i}^2 U_{\beta i}^2 + 2 \sum_{j>i} U_{\alpha i} U_{\alpha j} U_{\beta i} U_{\beta j} \cos\left(\frac{\Delta m_{ij}^2 L}{2E}\right) \quad (1.17)$$

$$= \delta_{\alpha\beta} - 4 \sum_{j>i} U_{\alpha i} U_{\alpha j} U_{\beta i} U_{\beta j} \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \quad (1.18)$$

The probability of no oscillation is given by:

$$P(\alpha \rightarrow \alpha) = 1 - \sum_{\alpha \neq \beta} P(\alpha \rightarrow \beta). \quad (1.19)$$

We note that observation of neutrino oscillations does not allow a direct means of probing the absolute values of neutrino mass states but only affords study of the mass squared differences between mass states. Considering only three neutrino flavours, e , μ and τ , propagating in a vacuum, we have:

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_\beta) &= \delta_{\alpha\beta} - 4 \sum_{i>j=1}^3 \text{Re}(K_{\alpha,\beta,i,j}) \sin^2\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \\ &+ 4 \sum_{i>j=1}^3 \text{Im}(K_{\alpha,\beta,i,j}) \sin\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \cos\left(\frac{\Delta m_{ij}^2 L}{4E}\right) \end{aligned} \quad (1.20)$$

where $K_{\alpha,\beta,i,j} = U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}$. The possible mass splittings are referred to as solar (m_{sol}) and atmospheric (m_{atm}) mass splittings and relate to our existing mass terms as follows:

$$\Delta m_{sol}^2 = \Delta m_{21}^2 \quad (1.21)$$

$$\Delta m_{atm}^2 = \Delta m_{31}^2 \approx \Delta m_{32}^2 \quad (1.22)$$

$$\Delta m_{atm}^2 \gg \Delta m_{sol}^2 \quad (1.23)$$

For simplicity it is assumed that in the general case only $\Delta m_{atm}^2 \approx \text{few} \times 10^{-3} \text{eV}^2$ is

important when considering neutrino oscillation experiments on earth.

$$P(\nu_e \rightarrow \nu_\mu) = 4|U_{13}|^2|U_{23}|^2 \sin^2\left(\frac{\Delta m_{atm}^2 L}{4E}\right) \quad (1.24)$$

$$= \sin^2(2\theta_{13}) \sin^2(\theta_{23}) \sin^2\left(\frac{\Delta m_{atm}^2 L}{4E}\right) \quad (1.25)$$

$$P(\nu_e \rightarrow \nu_\tau) = 4|U_{33}|^2|U_{13}|^2 \sin^2\left(\frac{\Delta m_{atm}^2 L}{4E}\right) \quad (1.26)$$

$$= \sin^2(2\theta_{13}) \cos^2(\theta_{23}) \sin^2\left(\frac{\Delta m_{atm}^2 L}{4E}\right). \quad (1.27)$$

$$P(\nu_\mu \rightarrow \nu_\tau) = 4|U_{33}|^2|U_{23}|^2 \sin^2\left(\frac{\Delta m_{atm}^2 L}{4E}\right) \quad (1.28)$$

$$= \sin^2(2\theta_{23}) \cos^4(\theta_{13}) \sin^2\left(\frac{\Delta m_{atm}^2 L}{4E}\right) \quad (1.29)$$

CP violation occurs when a process has a different probability to its ‘mirror’ process with antiparticles. In our present example, oscillation between two differing weak flavour eigenstates must have the same probability for both neutrinos and anti-neutrinos for CP invariance to hold. In the event that the PMNS matrix (equation 1.5) is complex, a CP-violation signature may be found in the observable:

$$\Delta P_{\alpha\beta}^{CP} = P(\nu_\alpha \rightarrow \nu_\beta) - P(\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta) \neq 0, \quad (1.30)$$

for $\alpha \neq \beta$.

In the case of high intensity beams where the effect of $|\Delta m_{12}^2|$ is not negligible, considering electron and muon flavour oscillations, we may expand our oscillation formulae (equations 1.25-1.29) as follows [20]:

$$\begin{aligned}
 P_{\nu_e \rightarrow \nu_\mu (\bar{\nu}_e \rightarrow \bar{\nu}_\mu)} &= s_{23}^2 \sin^2 2\theta_{13} \sin^2 \left[\frac{\Delta m_{13}^2 L}{4E} \right] \\
 &+ c_{23}^2 \sin^2 2\theta_{12} \sin^2 \left[\frac{\Delta m_{12}^2 L}{4E} \right] \\
 &+ \tilde{J} \cos \left[\pm \delta - \frac{\Delta m_{13}^2 L}{4E} \right] \left(\frac{\Delta m_{12}^2 L}{4E} \right) \sin \left[\frac{\Delta m_{13}^2 L}{4E} \right] \quad (1.31)
 \end{aligned}$$

where \pm corresponds to the ν and $\bar{\nu}$ states and \tilde{J} , the Jarlskog co-efficient for CP violation, is given as:

$$\tilde{J} = c_{13} \sin 2\theta_{12} \sin 2\theta_{23} \sin 2\theta_{13} \quad (1.32)$$

These definitions are consistent with our previous assertion that CP violation is possible if θ_{13} is non-zero and $\delta \neq 0$ or π ; in effect if the PMNS matrix is complex.

1.3.1 Neutrino Oscillations in Matter

Previously we have only considered neutrino oscillations in vacuum. Following [21], we now expand our discussion to neutrino oscillations in matter. The introduction of a medium causes elastic, coherent, forward scattering from constituent electrons which modifies the neutrino effective masses and mixing angles. However, interactions between neutrinos and other standard model particles do not result in flavour change and therefore do not produce a background to an experimental neutrino oscillation signal.

The propagation of a neutrino in the laboratory frame may be described by the Schrödinger equation:

$$i \frac{\partial}{\partial t} \Psi(t) = \mathcal{H} \Psi(t), \quad (1.33)$$

where t is the time, $\Psi(t)$ the neutrino wavefunction and \mathcal{H} the hamiltonian represented by an $N \times N$ matrix, where N is the number of active neutrino flavours. One may write the flavour state (α) amplitude as $\Psi_\alpha(t)$. All known flavours of neutrino interact with nucleons

through exchange of the Z^0 boson, or neutral current interaction. We have shown that the oscillation amplitude depends upon the relative phases of the neutrino eigenstates. Since the neutral current effect is equally weighted to all neutrino flavours, the relative phases of the mass states and therefore, by extension, the oscillation parameters are unaltered. Z^0 exchange causes the Hamiltonian to include a term proportional to the identity matrix representing the interaction energy which may be omitted when considering neutrino oscillations.

However, in the case of charged current interactions the electrons in matter will couple only to the electron flavour state through exchange of the W^+ boson. This discrimination adds an interaction energy term to the $\nu_e - \nu_e$ -th element of the Hamiltonian:

$$V = \sqrt{2}G_F N_e, \quad (1.34)$$

where G_F represents the familiar Fermi coupling constant ($G_F = 1.16637 \pm 0.00001 \times 10^{-5} \text{ GeV}^{-2}$) and N_e the number of electrons per unit volume of material. In the case of antineutrinos we take the negative of equation 1.34. Since the effect of the charge current interaction does affect the relative phase of the neutrino mass states it does alter the parameters governing neutrino oscillation.

For the purpose of this explanation let us limit our discussion to two flavour oscillation between ν_e and ν_μ . The Hamiltonian as defined previously is therefore a 2×2 matrix in flavour space. Omitting terms which do not contribute:

$$\mathcal{H} = \frac{\Delta m_M^2}{4E} \begin{bmatrix} -\cos 2\theta_M & \sin 2\theta_M \\ \sin 2\theta_M & \cos 2\theta_M \end{bmatrix}, \quad (1.35)$$

where Δm_M^2 is the effective mass squared difference in matter, E the neutrino energy and θ_M the effective mixing angle in the material. One may express these quantities in terms of vacuum equivalents:

$$\Delta m_M^2 = \Delta m^2 \sqrt{\sin^2 2\theta + (\cos 2\theta - x_\nu)^2}, \quad (1.36)$$

$$\sin^2 2\theta_M = \frac{\sin^2 2\theta}{\sin^2 2\theta + (\cos 2\theta - x_\nu)^2}, \quad (1.37)$$

where

$$x_\nu \equiv \frac{2\sqrt{2}G_F N_e E}{\Delta m^2} \quad (1.38)$$

The component of the oscillation probability due to material is:

$$P_{MSW}(\nu_e \rightarrow \nu_\mu) = \sin^2 2\theta_M \sin^2 \left(\Delta m_M^2 \frac{L}{4E} \right) \quad (1.39)$$

The term ‘MSW’ refers to three scientists: Mikheyev, Smirnov and Wolfenstein whose work [22, 23] informed our current view of neutrino oscillations in matter. The effect of matter upon neutrino oscillations is named the MSW effect in their honour.

In a full three flavour analysis, the probability of $\nu_e(\bar{\nu}_e) \rightarrow \nu_\mu(\bar{\nu}_\mu)$ flavour oscillation in matter is adapted as follows:

$$\begin{aligned} P_{\nu_e \rightarrow \nu_\mu(\bar{\nu}_e \rightarrow \bar{\nu}_\mu)} &= s_{23}^2 \sin^2 2\theta_{13} \left(\frac{\Delta m_{13}^2}{2EB_\mp} \right) \sin^2 \left[\frac{B_\mp L}{2} \right] \\ &+ c_{23}^2 \sin^2 2\theta_{12} \left(\frac{\Delta m_{12}^2}{2EA} \right) \sin^2 \left[\frac{AL}{2} \right] \\ &+ \tilde{J} \cos \left[\pm\delta - \frac{\Delta m_{13}^2 L}{4E} \right] \left(\frac{\Delta m_{12}^2}{2EA} \frac{\Delta m_{13}^2}{2EB_\mp} \right) \sin \left[\frac{AL}{2} \right] \sin \left[\frac{B_\mp L}{2} \right] \end{aligned} \quad (1.40)$$

where $B_\mp = \frac{1}{2E} \sqrt{(\Delta m_{13}^2 \cos 2\theta_{13} \mp A)^2 + \Delta m_{13}^4 \sin^2 2\theta_{13}}$ and \pm is conditional on whether we have a neutrino or anti-neutrino.

Experimental observation of CP violation requires measurement of the difference between $P_{(\nu_e \rightarrow \nu_\mu)}$ and $P_{(\bar{\nu}_e \rightarrow \bar{\nu}_\mu)}$ and disentangling of the asymmetry from matter effects due to B_\mp and the asymmetry due to the CP phase $\pm\delta$ from the last term in equation 1.40.

1.4 Neutrino Mass

Until the observation of neutrino oscillations, it was assumed that neutrinos were massless. However, experimental evidence of oscillations requires that mass states are not degenerate and that at least one neutrino mass-state is non-zero. The mass states cannot themselves be directly probed, since only flavour states are eigenstates of the weak interaction. Oscillation experiments are sensitive to the mass squared differences between mass states but cannot clarify absolute masses.

Direct mass experiments reconstruct neutrino masses from the kinematics of weak decays, inferring neutrino mass from precision measurements of charged decay products. Tritium β decay experiments provide an opportunity to probe the mass of the ν_e in particular and have to date placed an experimental limit of $m_{\nu_e} < 2$ eV [24] (95% confidence level). Study of pion decay has yielded a limit on the mass of the muon neutrino of $m_{\nu_\mu} < 190$ keV (90% confidence level). The limit on the mass of the ν_τ is calculated from study of pairs of τ leptons produced in electron-positron colliders and decaying into multiple pions, yielding a limit $m_{\nu_\tau} < 18.2$ MeV. However, results from oscillation experiments suggest that the difference between the ν_μ and ν_τ masses is small [25].

Since the absolute values of neutrino masses are unknown so too is the mass hierarchy. The case of $m_{\nu_3} > m_{\nu_2} > m_{\nu_1}$ (we assume no degeneracies) is referred to as the normal hierarchy, and the case where $m_{\nu_2} > m_{\nu_1} > m_{\nu_3}$ is referred to as the inverted hierarchy (figure 1.6).

In the Standard Model, leptons are described as Dirac fermions whose mass depends upon a coupling of left and right handed chiral components. If a neutrino can be described as a Dirac particle then the antineutrino, the right handed neutrino state, must be distinct from the left handed state. However there exists a competing description put forward by Majorana [27] which holds that the neutrino is a two component object identical to its antiparticle. If so, we can expect to observe a violation of lepton number symmetry of ± 2 under CPT transformation. Double beta decay experiments provide a promising venue for the study of the possible Majorana nature of the neutrino [28].

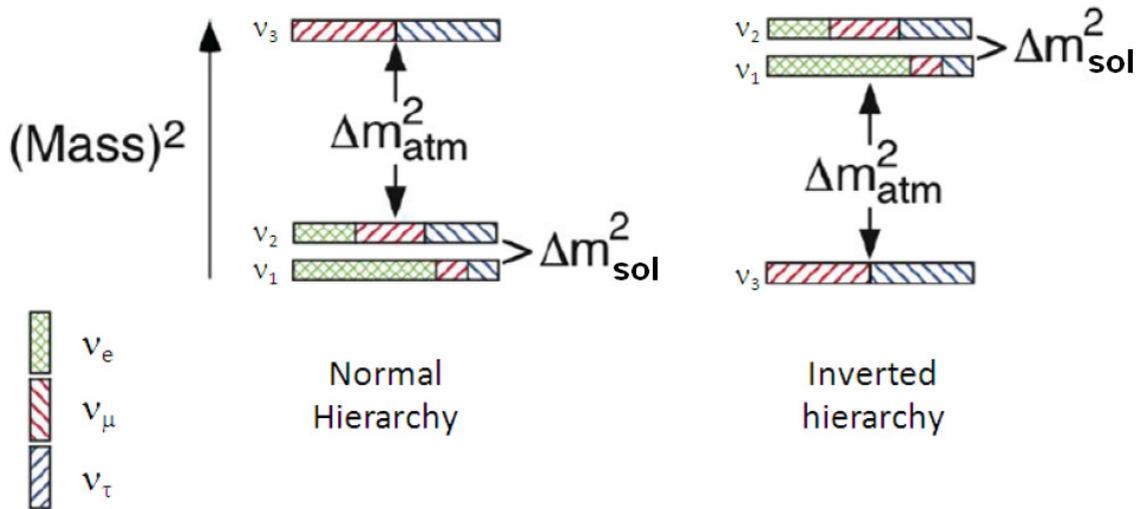


Figure 1.6: As yet the nature of the mass hierarchy is unknown due to ambiguities in our existing absolute mass scale measurements. Figure from [26], symbols changed to maintain consistency.

1.5 Review of Neutrino Experiments

In over eighty years since the existence of the neutrino was first postulated and despite improvements both in theory and detection, the number of questions regarding the nature of the neutrino has increased in step with our experimental advances. Both the nature and behaviour of the neutrino remain matters of considerable interest. Is it a Dirac or Majorana particle? What are the parameters governing neutrino oscillations? Questions remain regarding the number of neutrino flavours and possible existence of sterile neutrinos, whether or not the neutrino might provide a venue for observing CP violation and contribute to the matter-antimatter asymmetry present in the observable universe or what the mass of the neutrino is.

In this section we restrict ourselves to providing an overview of a selection of significant experiments in the field of neutrino oscillations. Within this field, considerable diversity exists with both artificial (nuclear reactors, accelerators) and natural neutrino sources (the sun, atmospheric production from cosmic rays) and a wide range of detection paradigms.

1.5.1 Reactor Based

Reactor experiments have played a crucial role in developing our understanding of the neutrino, from their discovery in 1956 at the Savannah River nuclear reactor experiment, to more recent oscillation experiments including KamLAND and CHOOZ¹. Nuclear reactors provide near monochromatic sources of $\bar{\nu}_e$ with a well defined mean energy (≈ 3.6 MeV) and a well understood flux. Measurement of mixing angles at reactor facilities mitigates degeneracies and correlations between oscillation parameters which may occur in accelerator experiments.

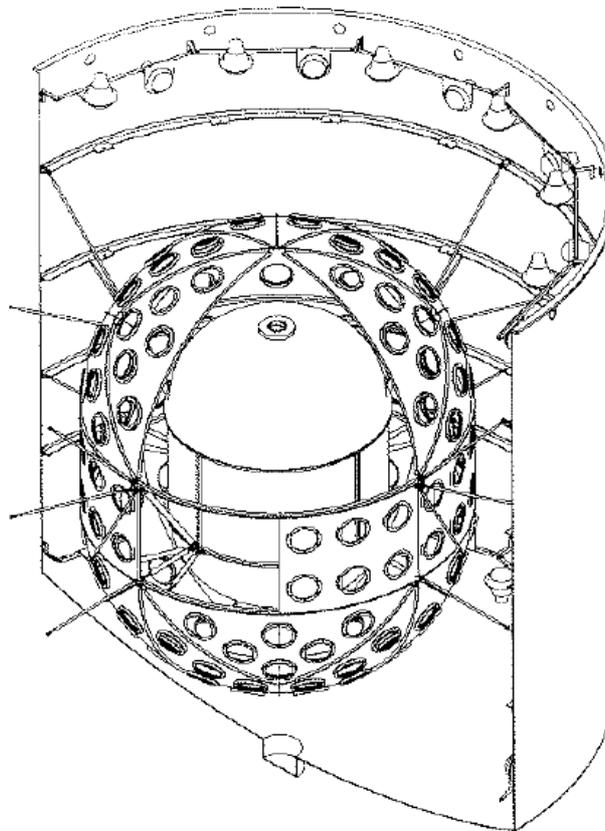


Figure 1.7: Schematic of the CHOOZ detector target housing with holes for PMTs. Figure from [29].

CHOOZ was a reactor experiment with a detector target in Ardennes, France receiving $\bar{\nu}_e$ flux from the Chooz nuclear power station, consisting of two 4.25 GWth twin-pressurised water reactors. The detector target, pictured in figure 1.7, follows a common liquid scintillator design with three concentric regions comprising a 5 ton central target with 0.09% Gd-

¹Not an acronym.

loaded scintillator in a Plexiglas container, a 70 cm thick, 17 ton region comprising 192 8 inch PMTs and a final 90 ton outer region with 24 8 inch PMTs. CHOOZ featured an average $\frac{L}{E}$ value of $300 \frac{\text{m}}{\text{MeV}}$ where $L \approx 1 \text{ km}$ and $E \approx 3 \text{ MeV}$. Neutrino signatures were produced through the inverse beta decay reaction:

$$\bar{\nu}_e + p \rightarrow e^+ + n, \quad (1.41)$$

leading to a time correlated positron and neutron signal distinctive from background. CHOOZ did not observe $\bar{\nu}_e$ disappearance, to a 90% confidence level, excluding $\sin^2 2\theta_{13} > 0.18$ and $\Delta m_{23}^2 > 0.9 \times 10^{-3} \text{ eV}^2$ [29]. An upgrade, Double CHOOZ, is planned which will use a second detector to reduce systematic uncertainties [30].

KamLAND (KAMioka Liquid scintillator ANti-neutrino Detector) comprises a 13 m diameter spherical steel inner volume holding 1 kiloton of ultra-pure liquid scintillator, with 1,879 PMTs lining the inner volume. This is surrounded by a 3.2 kiloton water-Cherenkov detector with 225 PMTs which provides shielding from natural sources of radiation. KamLAND detected $\bar{\nu}_e$ from all nuclear power plants in Japan (53 reactors), however 80% of its flux was received from reactors 140-210 km distant. KamLAND made the first measurement of reactor antineutrino disappearance and, by combining its data with SNO contributed to best fit values of $\tan^2 \theta_{12} = 0.47_{-0.05}^{+0.06}$ and $\Delta m_{21}^2 = 7.59_{-0.21}^{+0.21} \times 10^{-5} \text{ eV}^2$ [31]. KamLAND also performed a detailed geophysical study of the decay of ^{238}U and ^{232}Th in the earth [32].

1.5.2 Solar Based

We have described the contribution of the Homestake experiment in providing the first experimental evidence of a deficit in the terrestrial flux of ν_e when measured against the predictions of the Standard Solar Model (SSM) that models the sun according to the reaction chain in figure 1.8 and that produces the neutrino spectra illustrated in figure 1.9 [33]. The deficit of solar neutrinos was later confirmed by the SAGE [34] and GALLEX [35] experiments.

These experiments contributed to a picture of neutrino behaviour which is now known to be consistent with neutrino oscillations. However, they acted only as neutrino counters

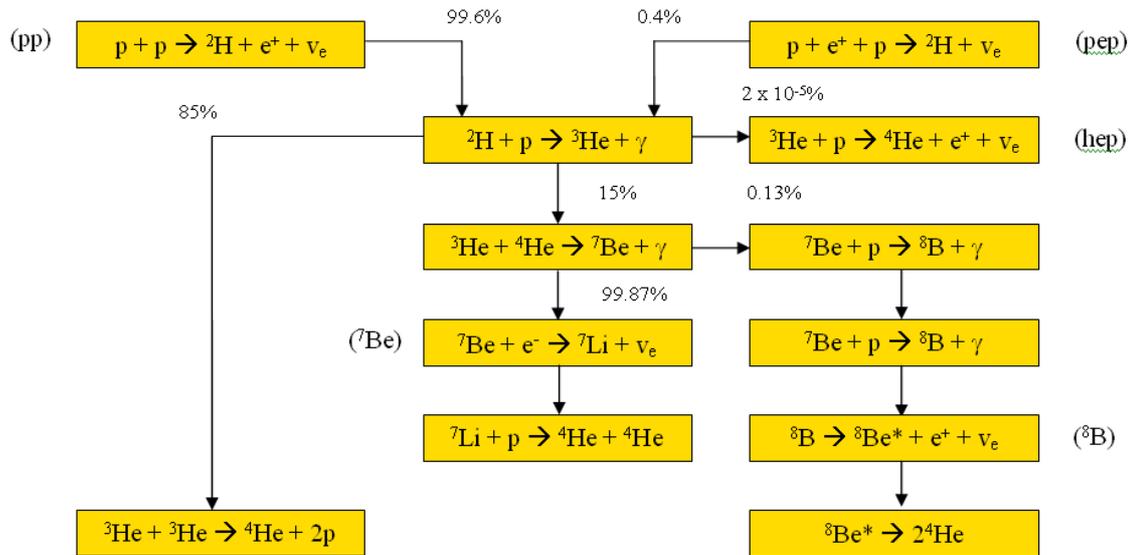


Figure 1.8: The standard solar model [33] describes expected terrestrial neutrino flux from nuclear reactions in the sun, assuming no neutrino oscillations.

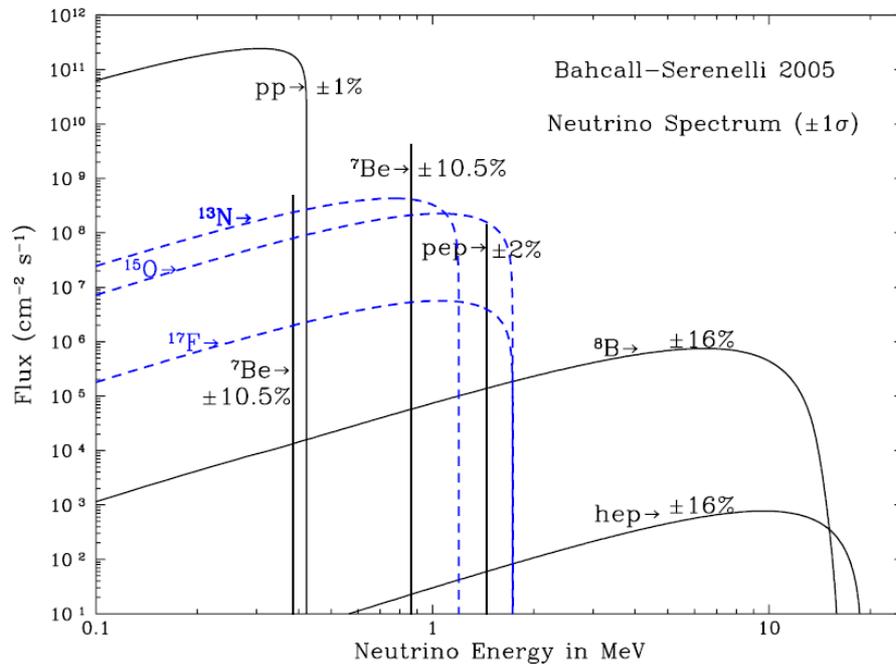


Figure 1.9: Expected flux of solar neutrinos according to the SSM [33].

and were unable to discern the direction of the neutrinos producing reactions in their volumes. Final confirmation would require angular and timing information to exclude competing theories such as neutrino decay and decoherence. The Kamiokande (KAMIOKANDE: KAMIOKA Nucleon Decay Experiment) and upgraded Super-Kamiokande (Super-K) experiments exploit $\nu_e + e^- \rightarrow \nu_e + e^-$ scattering which produces Cherenkov light, providing

energy and directional information on inbound solar electron neutrinos. Super-K, shown in figure 1.10, was a 50 kTon pure water Cherenkov detector [36] located 1 km underground with an inner volume lined with 1,885 8-inch photomultiplier tubes (PMTs) and 11,146 PMTs surrounding the outer volume of the detector. A Super-K run of 1496 days produced angular distributions (figure 1.11) confirming the inbound neutrinos as solar neutrinos. The flux recorded at Super-K was $2.35 \pm 0.02 \pm 0.08 \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$, suggesting a 64% deficit when compared with the SSM prediction [37].

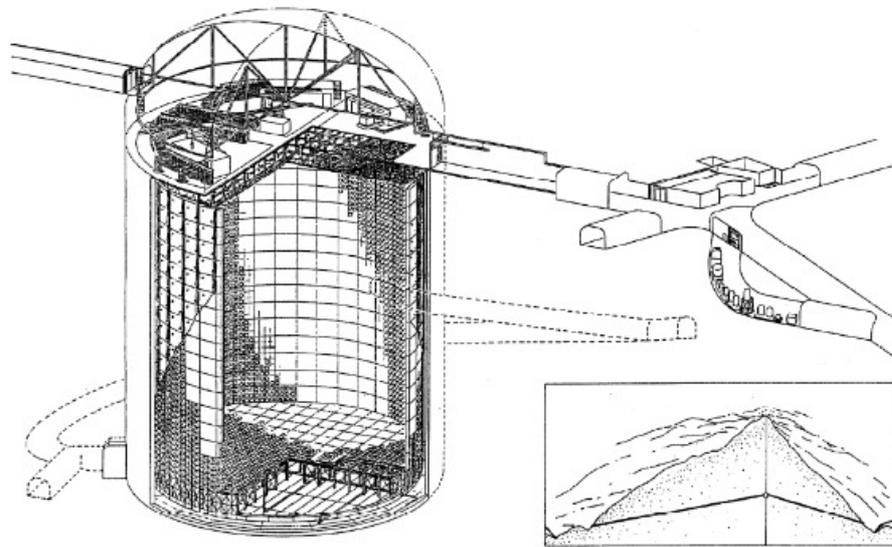


Figure 1.10: The Super-Kamiokande solar neutrino detector located 1 km underground at Mount Kamioka. Figure from [36].

For both the SSM and the theory of neutrino oscillations to be simultaneously correct, sensitivity to all known neutrino flavours is required. The Sudbury Neutrino Observatory (SNO), a 1 kTon heavy water (D_2O , (with subsequent addition of salt, NaCl , for enhanced neutron detection) Cherenkov detector positioned ~ 2 km underground in Ontario, was sensitive to ^8B solar neutrinos of all three flavours through elastic scattering (ES), charged current (CC) and neutral current (NC) interaction channels summarised below:

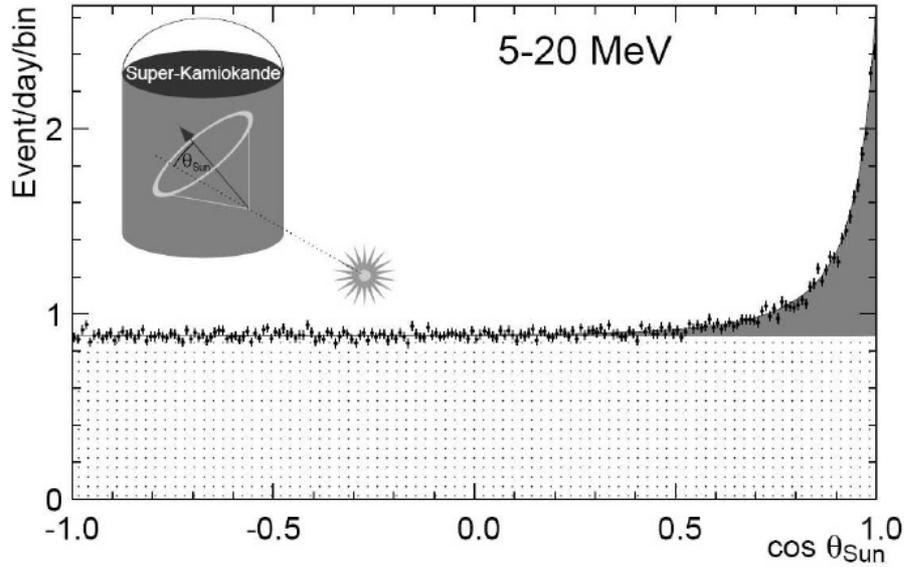


Figure 1.11: The angular distributions of solar neutrinos detected at Super-K [37].

$$\text{ES } \nu_x + e^- \rightarrow \nu_x + e^- \quad (1.42)$$

$$\text{CC } \nu_e + d \rightarrow e^- + p + p \quad (1.43)$$

$$\text{NC } \nu_x + d \rightarrow \nu_x + p + n \quad (1.44)$$

where ν_x denotes any active flavour neutrino. Elastic scattering provides sensitivity to all neutrino flavours but reduced sensitivity for ν_μ and ν_τ . The CC channel confirmed the deficits observed in previous experiments, measuring roughly $\frac{1}{3}$ of the SSM predicted flux. Neutrino oscillations would require that the ES flux exceed the CC flux. The NC channel performed identically for all known neutrino flavours, enabling a measurement of total neutrino flux which was consistent with the SSM (Figure 1.12), providing evidence for neutrino oscillations and solving the solar neutrino problem [38].

1.5.3 Atmospherically Based

Neutrinos are produced in the Earth's atmosphere from cosmic rays, where a pion decays into a muon ($\pi^\pm \rightarrow \mu^\pm + \nu_\mu(\bar{\nu}_\mu)$) which subsequently produces an electron- or muon-flavoured

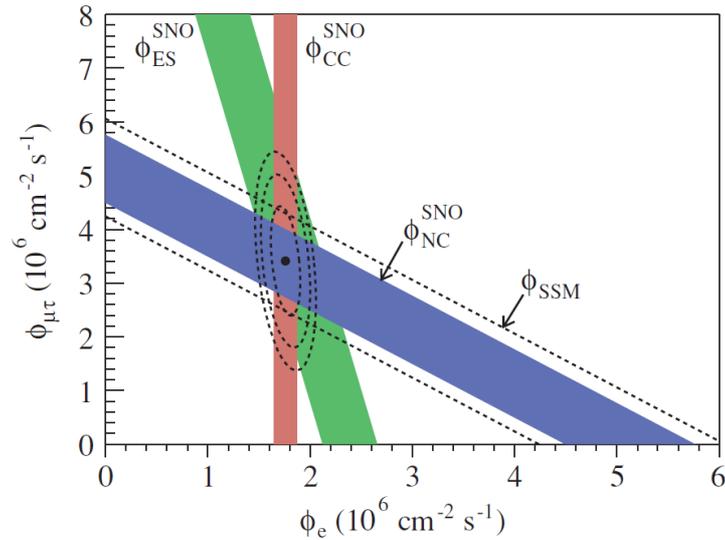


Figure 1.12: SNO results showing ν_μ, ν_τ flux plotted against ν_e flux. Dashed lines indicate the BP2000 SSM prediction for total ${}^8\text{B}$ flux [39] which compares favourably with the flux reconstructed from NC channel events (blue band). The y-axis intercepts of both bands represent $\pm 1\sigma$ uncertainties. The green band signifies the ES channel and the red band the CC channel. The intersection of the bands with the fit values for $\theta_{\mu\tau}$ and θ_e are consistent with the oscillatory neutrino behaviour. Figure from [38]

neutrino ($\mu^\pm \rightarrow e^\pm + (\bar{\nu}_\mu)\nu_\mu + (\nu_e)\bar{\nu}_e$). To date there exists two means of studying atmospheric neutrinos; through water Cherenkov and fine resolution tracking detectors (FRTDs). A more rigorous overview is provided in [40].

Water Cherenkov detectors record signals of Cherenkov light from relativistic charged particles produced in the neutrino interaction. The internal surface is covered by a two dimensional array of PMTs which record the number of photo-electrons received, allowing reconstruction of particle energy, as well as time and pulse height. Data from the PMT array can be used to infer vertex position.

FRTDs provide the advantage of sensitivity to particles whose momentum is below the Cherenkov threshold in water. Various configurations have been built and proposed including magnetised detectors (MINOS [41], SOUDAN2 [42]), and liquid scintillator time of flight layers (first used in MACRO [43]) which resolve the direction of muons from their time of flight.

The largest atmospheric neutrino detector to date is Super-K, an example of a water Cherenkov design. We have already discussed the sensitivity of Super-K to directional in-

formation through the electron scattering channel. This is of particular use in studying atmospheric neutrinos. As a consequence, Super-K is able to discern whether a neutrino candidate event was incident on the detector in the downwards direction suggesting a distance of travel ~ 15 km ($\cos\theta = 1$) or on the upwards direction ($\cos\theta = -1$) suggesting a distance of ~ 13000 km having been produced in the atmosphere on the other side of the earth. This provides the experiment with a continuous distribution of baselines through which to study oscillation behaviour in atmospheric neutrinos.

Figure 1.13 shows the zenith angle distributions observed at Super-K with clear evidence of oscillatory behaviour in the muon data and no evidence of disappearance in the electron data. The Super-K Collaboration have performed a full three flavour neutrino oscillation analysis and, assuming a normal mass hierarchy, provided best fit results for the mixing angles and mass squared difference of $\Delta m_{23}^2 = 2.5 \times 10^{-3}$, $\sin^2\theta_{23}=0.5$ and $\sin^2\theta_{13} = 0.0$ [44] which are consistent with existing experimental bounds imposed by CHOOZ [29], K2K [45] and Palo Verde [46].

1.5.4 Accelerator Based

Accelerators produce neutrino beams from decays of muons, which themselves decay from pions created by proton interactions in a nuclear target. In the future accelerators promise a range of advantages to neutrino oscillation studies with the next generation of super-beams promising high flux and betabeam facilities providing an opportunity to study monochromatic beams. A neutrino factory facility, described in the following chapter, will provide unprecedented precision in measuring oscillation parameters and unmatched discovery potential for CP violation in the neutrino sector. Mixing angle measurements from atmospheric experiments suggests a high probability of muon flavour eigenstates oscillating to tau flavour eigenstates (muon disappearance), a promising venue for studying the mass difference between the two flavours of neutrino through comparison of oscillation probability as a function of energy/distance at fixed distance/energy.

The electron neutrino appearance channel also provides excellent opportunities. Indeed,

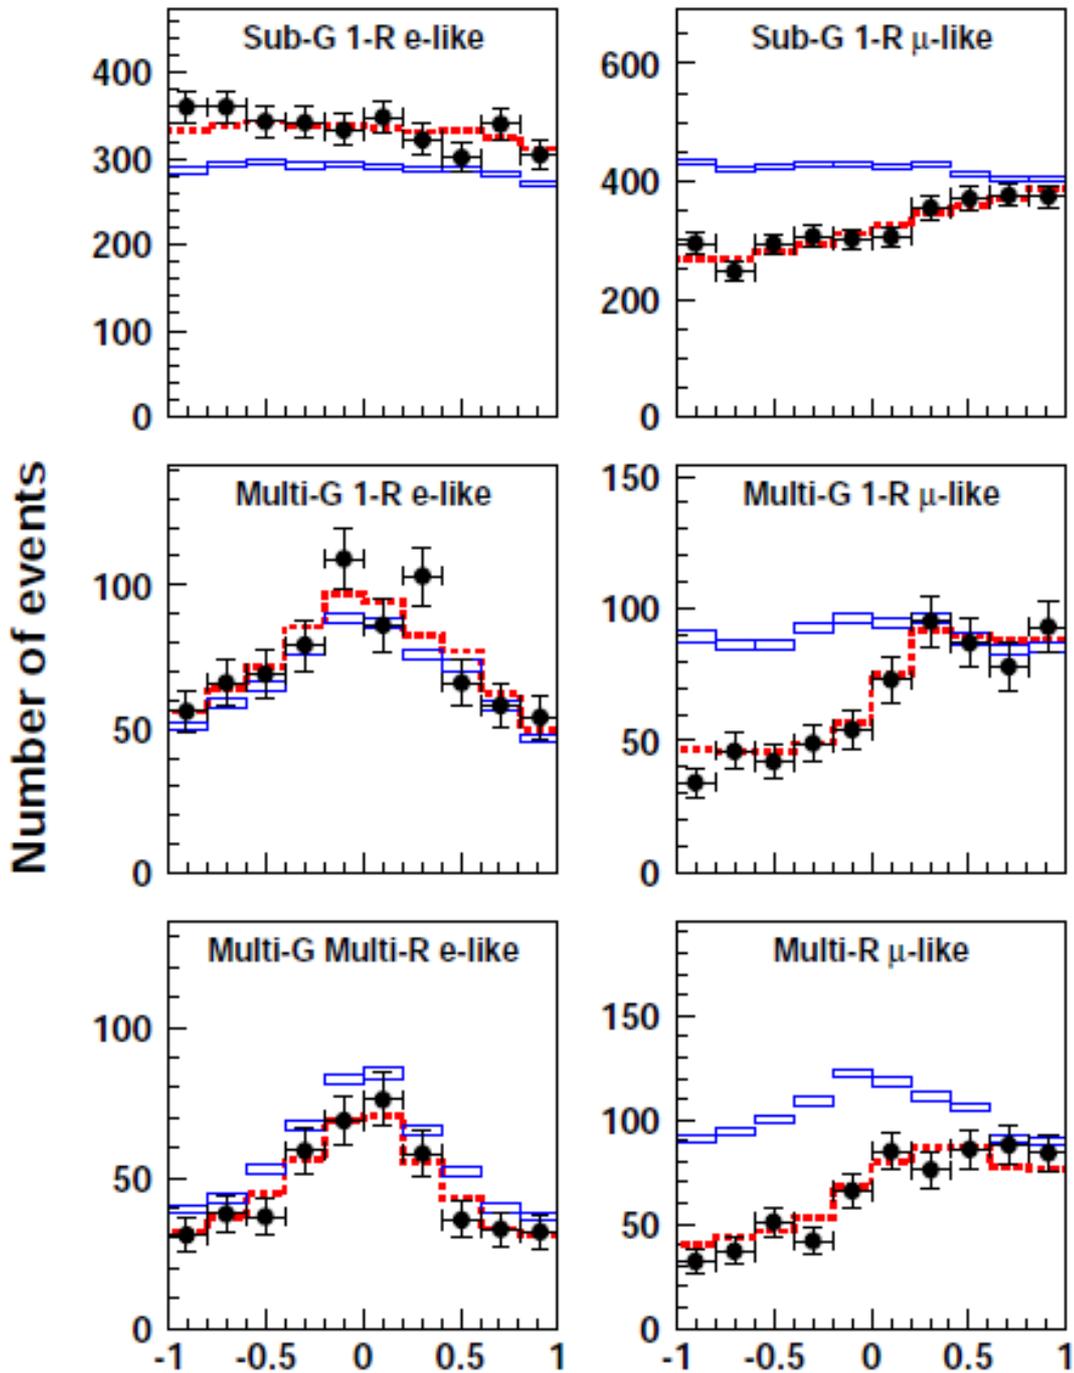


Figure 1.13: Zenith angle distributions for fully contained atmospheric neutrino candidates detected at Super-K. Data points are signified with black circles and Monte-Carlo simulations assuming no oscillations are provided in blue. The left graphs relate to electron-like events and the right graphs muon-like events. Figure taken from [44].

the only viable means of precision measurements of both the CP violation and the neutrino mass hierarchy are comparisons of ν and $\bar{\nu}$ oscillations between muon and electron neutrinos from an accelerator-based source [40].

To date, accelerators have made a valuable contribution to the history of neutrino experiments. The Super-K detector, in addition to its numerous applications described previously, also acts as a far detector to a neutrino beamline produced at the High Energy Accelerator Research Organisation (referred to as KEK) located in Japan in the city of Tsukuba, Ibaraki prefecture. KEK provided a ν_μ beam from a 12 GeV proton synchrotron, 250 km from the Super-K detector. The experiment was named K2K (KEK to Kamioka) and provided an opportunity to confirm the solar neutrino oscillation results of Super-K with a well-understood artificial beam under experimental control. The protons produced from the synchrotron were incident on an aluminium target. Interactions with the target material produced pions which underwent focussing in a magnetic horn. Pions subsequently decayed into muons and muons into neutrinos.

K2K took data over a period of 5 years and provided best-fit measurements of $\sin^2 2\theta_{23} = 1$ and $|\Delta m_{23}^2| = 2.75 \times 10^{-3} \text{ eV}^2$ and concluded that the probability of their results being consistent with expectations for no oscillations was 0.0015% (4.3σ) [45]. The K2K results are summarised in figure 1.14.

The Main Injector Neutrino Oscillation Search (MINOS) experiment provides a muon neutrino beam to steel-scintillator sampling calorimeter near and far detectors, 1.04 and 735 km from source, 0.98 and 5.4 kTon in mass. The neutrino beam produced from the NuMI beamline at Fermilab has a peak energy of 3 GeV. MINOS has measured $\sin^2\theta_{23} > 0.87$ and $|\Delta m_{23}^2| = 2.74_{-0.26}^{+0.44} \times 10^{-3} \text{ eV}^2$. An analysis of electron neutrino appearance has produced results, summarised in figure 1.15, which show a normal hierarchy best fit for $\sin^2 2\theta_{13}$ just beneath the CHOOZ limit, assuming $|\Delta m^2| = 2.43 \times 10^{-3} \text{ eV}^2$ and $\sin^2 2\theta_{23} = 1.0$ for all values of δ_{CP} . Here Δm^2 refers to a mixture of $|\Delta m_{13}^2|$ and $|m_{23}^2|$ which is unresolved [41].

Recently, MINOS has detected indications of a discrepancy in $\bar{\nu}_\mu$ flux compared with that of ν_μ . Configuring the NuMI beamline for antineutrinos, MINOS reconstructed values of $|\Delta \bar{m}^2| = 3.36_{-0.40}^{+0.46}(\text{stat.}) \pm 0.06(\text{syst.}) \times 10^{-3} \text{ eV}^2$ and $\sin^2(2\bar{\theta}) = 0.86_{-0.12}^{+0.11}(\text{stat.}) \pm 0.01(\text{syst.})$ in a two-flavour oscillation analysis where only one mixing angle and mass splitting is required. These are consistent with the ν_μ measurements at the 2% confidence level.

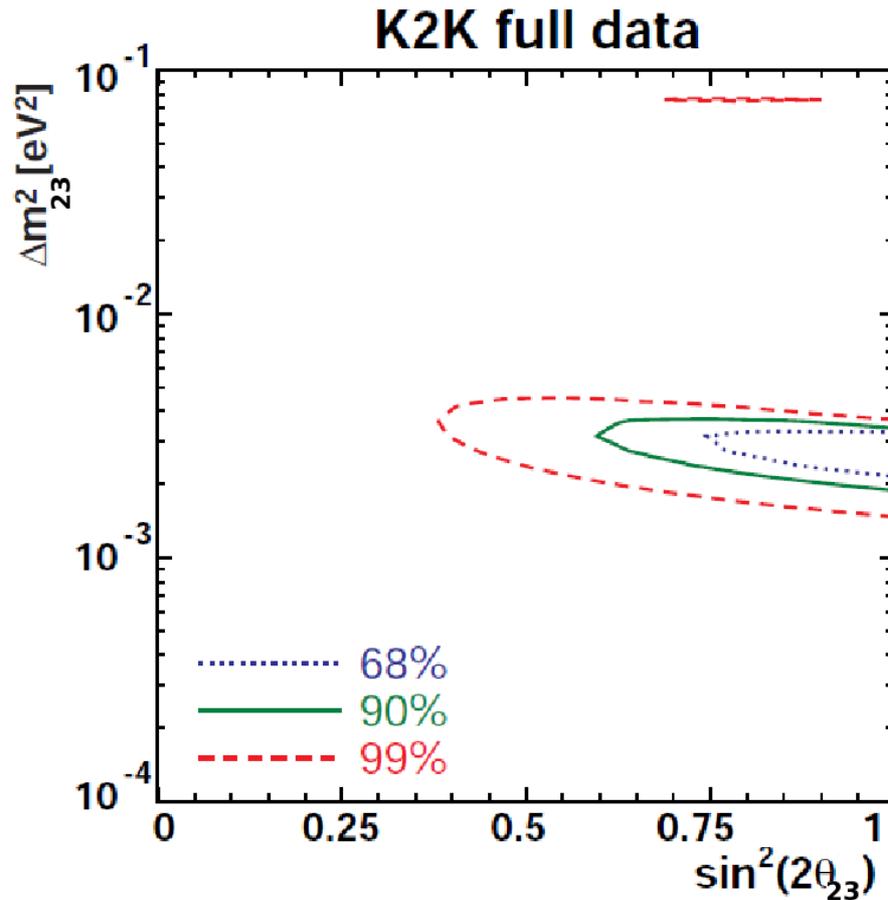


Figure 1.14: Parameter space showing oscillation parameter regions consistent with K2K results. Figure from [45].

Further statistics are required to confirm this result [41].

The preceding Liquid Scintillator Neutrino Detector (LSND), cited at Los Alamos, provided additional controversy. An 800 MeV proton beam was fired through a steel tank containing 167 tonnes of liquid scintillator surrounded by 1,220 uniformly spaced 8" Hamamatsu PMTs [47]. The liquid scintillator is CH₂ mineral oil plus 0.031 grams/litre of butyl-PBD affording both Cherenkov light and isotropic scintillation light detection. The experiment suffers from a significant cosmic ray background mitigated by a veto shield surrounded by 292 uniformly spaced 5" EMI PMTs with a 6 hit threshold [48]. LSND's results suggested oscillation of muon to electron neutrinos with a relatively large mass splitting (0.1 eV²) which is inconsistent with existing experimental data characterising neutrino oscillations, a problem referred to in the literature as the "LSND puzzle". Attempts to resolve

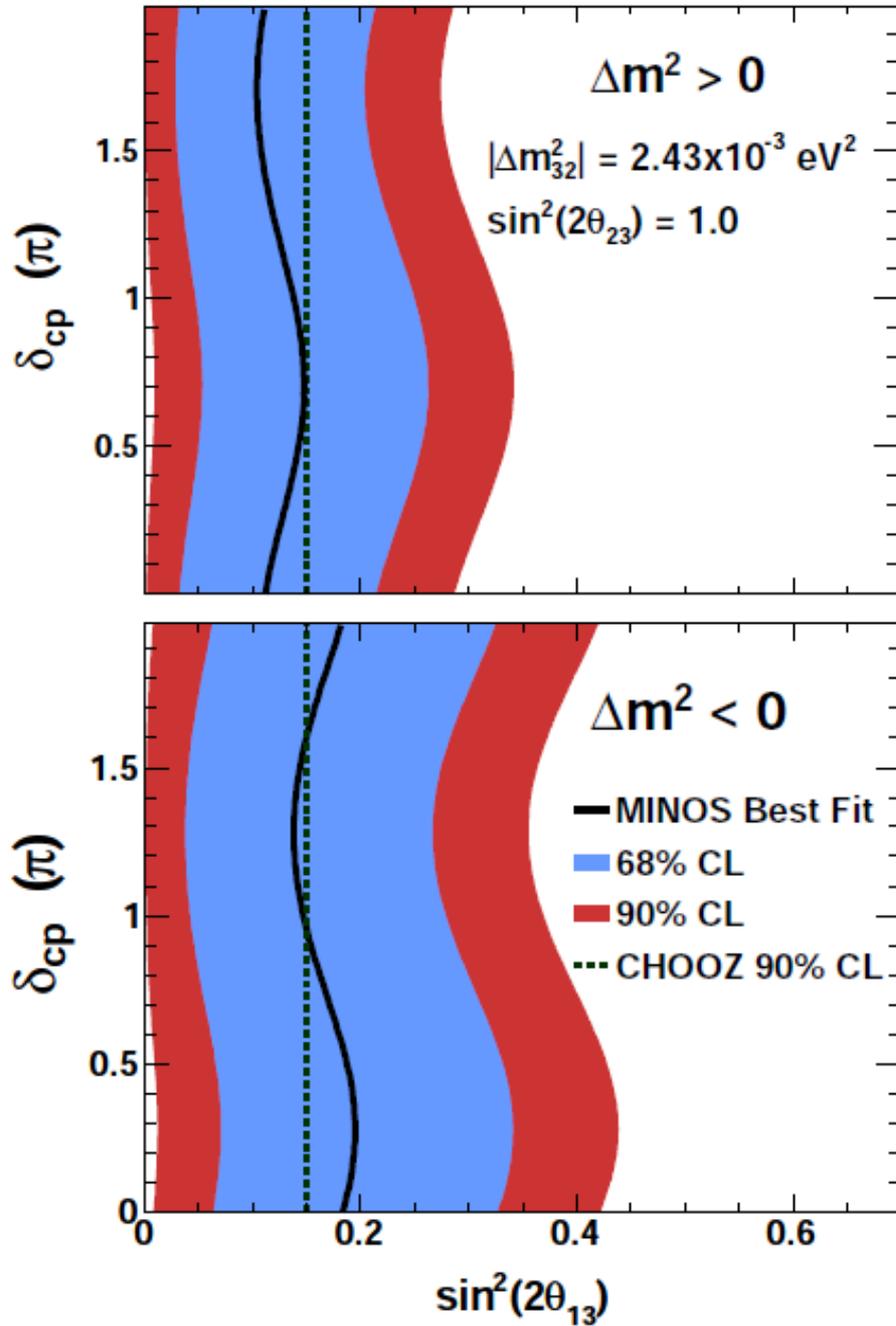


Figure 1.15: Parameter space of $\sin^2(2\theta_{13})$ and δ_{CP} consistent with event rates seen at MINOS assuming normal mass hierarchy (top) and inverted mass hierarchy (bottom). Figure from [41].

the problem of the anomalously high mass squared difference typically rely on explanations involving one or more heavy (1 eV scale) sterile neutrinos [49, 50].

MiniBooNE (Miniature Booster Neutrino Experiment) and its proposed upgrade BooNE have been designed to clarify the LSND results definitively. MiniBooNE comprises a 1 GeV

neutrino source positioned 500 m from an 800 ton mineral oil detector with 1,280 PMTs lining the inside of the tank. Initial data taking with a muon neutrino beam found no evidence of the LSND $\nu_\mu \rightarrow \nu_e$ effect [51], however, subsequent running with antineutrinos did find evidence of $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ consistent with observations at LSND [52]. Future work at the upgraded BooNE facility, ICARUS [53] at CERN or OscSNS [54] at Oak Ridge National Laboratory could provide conclusive evidence for sterile neutrinos at a 5σ discovery threshold.

Further large scale accelerator-based experiments are required to clarify the parameters which govern neutrino oscillations and search for exotic oscillatory behaviour. The CERN Neutrinos to Gran Sasso (CNGS) experiment operates a beamline between the CERN facility and OPERA [55] and ICARUS (a liquid argon time projection chamber) detectors in Italy. OPERA (Oscillation Project with Emulsion-tRacking Apparatus) is constructed from bricks comprising photo-emulsion film and lead plates. The presence of tau neutrinos is inferred from observations of kinks in tau decays corresponding to undetected neutrinos. This provides a more direct basis for studying $\nu_\mu \rightarrow \nu_\tau$ oscillations than simply observing a deficit of ν_μ candidate events. The first ν_τ candidate event is shown in figure 1.16 [56].

The T2K (Tokai to Kamioka) experiment will take a mono-energetic 0.5 – 0.9 GeV off-axis ν_μ beam from J-PARC to the Super-K detector in Kamioka, a distance of 295 km. Placing detectors off-axis to ν_μ neutrino beams is an established strategy to limit ν_e contamination. The experiment seeks to measure oscillations of $\nu_\mu \rightarrow \nu_e$ and measure θ_{13} [57]. Recently T2K have announced results based on an analysis of 1.43×10^{20} protons on target which suggest that $0.03(0.04) < \sin^2 2\theta_{13} < 0.28(0.34)$ at the 90% confidence level for normal (inverted) mass hierarchy [58] however further work at higher statistical precision is required to confirm this result. The present schedule of T2K must be reviewed following a complete assessment of damage to JPARC during the earthquake and tsunami of March 2011.

The NO ν A experiment is a complimentary near-future accelerator-based experiment exploiting the Fermilab NuMI beamline which produces neutrinos 14.6 mrad off-axis to a 15 kton far detector in Minnesota (a distance of 810 km), also examining electron neutrino

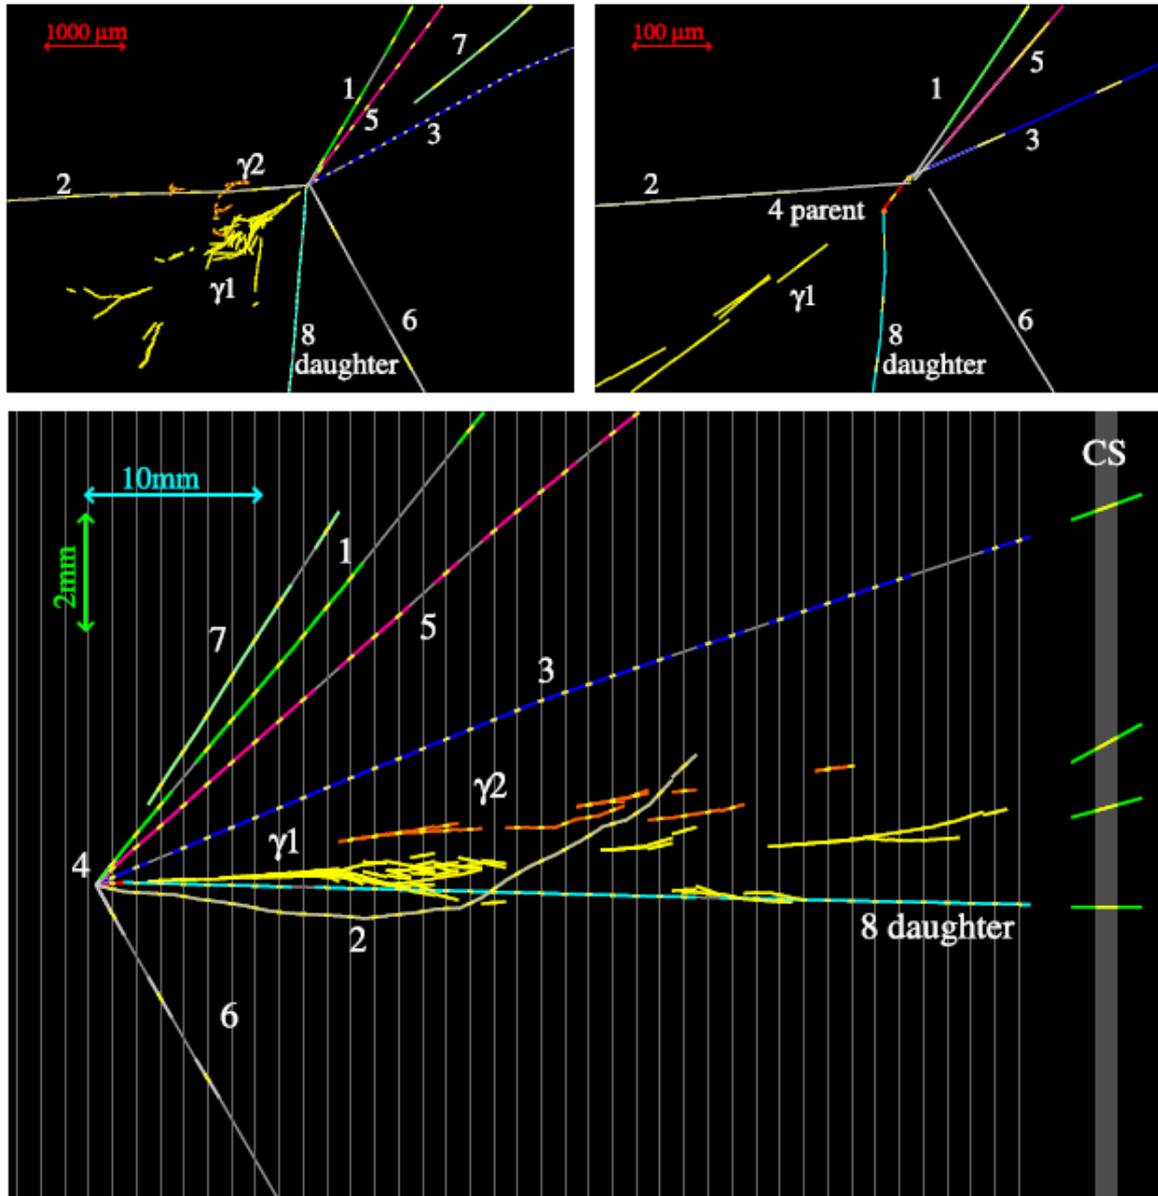


Figure 1.16: The first ν_τ candidate event in OPERA. Top right and top left: views transverse to neutrino direction. Bottom: view longitudinal to neutrino direction [56].

appearance. NO ν A is also supported by a 215 ton near detector 1 km from source at the Fermilab site. It is hoped that by comparing neutrino and antineutrino oscillation channels NO ν A may provide some insight into the mass hierarchy and the CP-violating phase δ in addition to the possibility of measuring the mixing angle θ_{13} [59].

Up to now, the only accelerator based neutrino beamlines have been from pion and kaon decay. Upgraded neutrino facilities from pion and kaon decay, with megawatt class proton beams, are termed Superbeams and are being studied in Europe [60, 61], the USA [62] and Japan [63, 64]. Future beams from muon decay (as in the proposed neutrino factory [65, 60] which would produce ν_μ and $\bar{\nu}_e$) or alternatively radioactive ions (beta beams [66, 60] producing ν_e or $\bar{\nu}_e$) are also currently under study. These are accelerated in storage rings with long decay straights which are aligned with one or more target detectors. Neutrinos produced from the decay process are uninhibited by the magnets constraining the muons or ions to the path of the ring and therefore continue along the path of their decay. It is common to arrange for both near detector and far detector facilities, where the latter may even be at a distance of order 10^3 - 10^4 km from the decay ring. A diversity of detector paradigms have been put to use in accelerator experiments including but not limited to water Cherenkov, liquid argon time projection, emulsion-lead and segmented scintillator systems. These are detailed in reference [40].

These near future and future experiments will be of great importance in mapping the future direction of neutrino physics research, with possibilities of confirming new particles and either measuring oscillation parameters or placing new limits on their possible values. In the long term the community is committed to a programme of research and development work to enable the option of constructing a large scale neutrino factory facility, described fully in chapter 2. Such a facility would allow unambiguous clarification of the neutrino mass hierarchy; provide record-breaking precision for mixing angle measurements and enable an investigation of possible CP-violation in the lepton sector. Alternatively, the community retains the option to build a beta-beam or superbeam facility, also described in chapter 2.

1.6 Conclusions

Decades of experiments have built an unassailable case for the existence of neutrino oscillations and by implication, neutrino mass. A summary of the state of the art is given in table 1.1.

Parameter	Best Fit Value
$\sin^2\theta_{12}$	$0.312^{+0.017}_{-0.015}$
$\sin^2\theta_{13}$	$0.010^{+0.009}_{-0.006}$ (norm) $0.013^{+0.009}_{-0.007}$ (inv)
$\sin^2\theta_{23}$	0.51 ± 0.06 (norm) 0.52 ± 0.06 (inv)
$\Delta m_{21}^2 [10^{-5} \text{ eV}^2]$	$7.59^{+0.20}_{-0.18}$
$\Delta m_{32}^2 [10^{-3} \text{ eV}^2]$	2.45 ± 0.09 (norm) - $2.34^{+0.10}_{-0.09}$ (inv)

Table 1.1: Best fit ($\pm 1\sigma$) values of the oscillation parameters. Results from both normal and inverted hierarchy assumptions are shown. Table from [67].

Nevertheless our understanding of the nature and behaviour of the neutrino remains incomplete and further precision is required. In particular the community looks with great interest towards potential discovery of a non-zero CP violating phase δ . These issues necessitate a new generation of neutrino experiments with unprecedented precision the results from which promise to augment the existing standard model of particle physics. We will now discuss such a facility, the neutrino factory, and the research and development efforts required to realise it in the following chapters.

Chapter 2

Neutrino Factory

2.1 Introduction

The discovery of neutrino oscillations together with their implications, described in section 1.3, raises a number of as yet unresolved issues which confound our understanding of the neutrino. In particular work remains to develop our understanding of neutrino oscillation parameters and clarify the mass hierarchy of the neutrino eigenstates. In addition, the neutrino provides a means to probe the phenomenon of leptonic CP violation.

The neutrino factory facility (Figure 2.1) is an accelerator complex, first proposed over thirty years ago [68, 69], that aims to create intense beams of neutrinos from the decay of a beam of accelerated muons contained in decay rings. The decay rings will be aligned with long baseline detectors and will produce a flux of 10^{21} neutrinos per year. Much research and development work [70, 71] has been carried out to define the parameters of the neutrino factory and a summary of these efforts is presented below.

2.2 The Case for a Neutrino Factory

As described previously, there now exists a rich body of experimental data which confirm the existence of neutrino oscillations. The neutrino factory aims to further develop our understanding of neutrino oscillations by: measuring the CP violating phase δ_{CP} , measuring

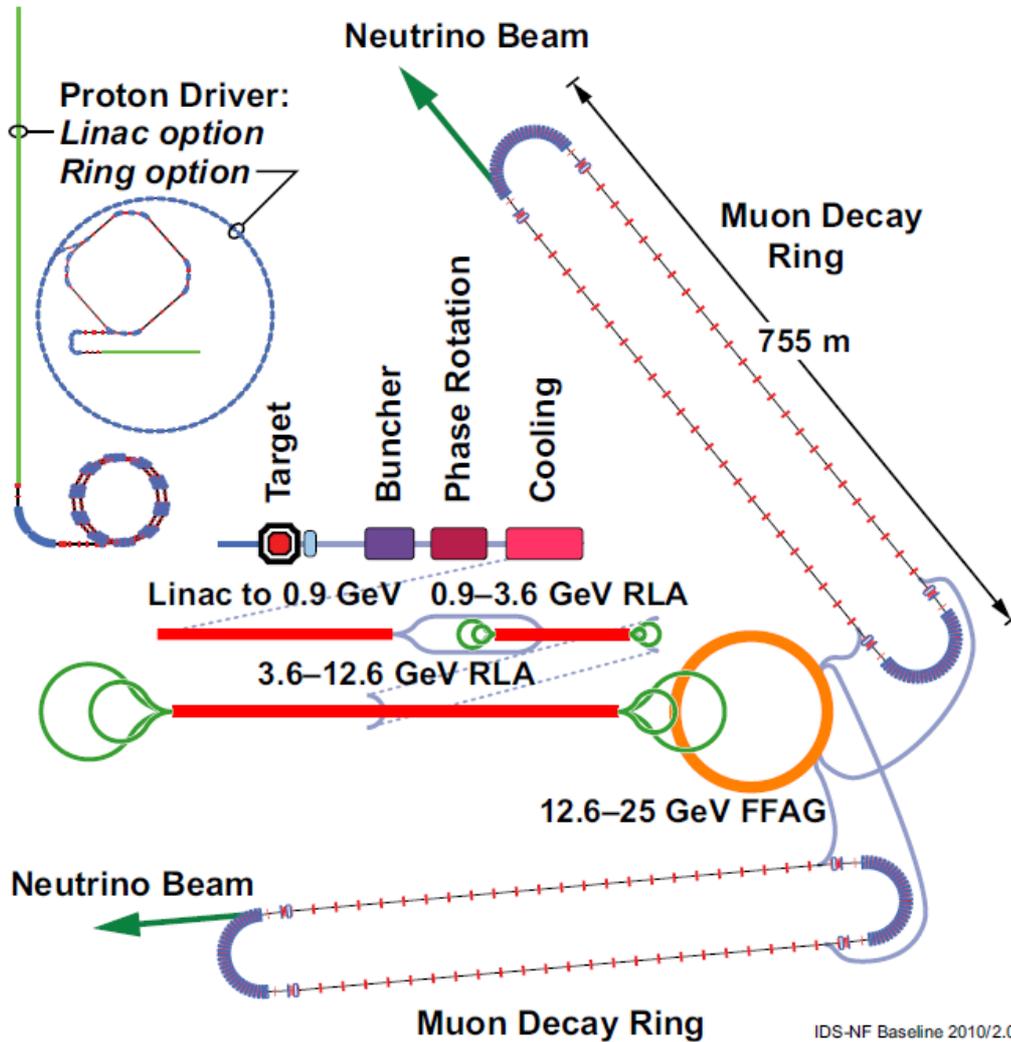


Figure 2.1: A schematic of a future neutrino factory facility [70].

$\sin^2 2\theta_{13}$ and determining the hierarchy of neutrino mass eigenstates. These terms have already been described in detail together with an explanation of their relevance to neutrino oscillation physics, in section 1.3.

The neutrino factory offers twelve oscillation channels for precision study. The experimental appearance and disappearance channels a neutrino factory is expected to be sensitive to are listed in table 2.1. Of these, the golden channel provides the primary means of investigation into these worthy areas of physics interest. “Wrong-signed” muons, detected at a far detector, possessing opposite charge to the muons in the corresponding storage ring at that time provide a clear signal for a magnetised detector with strong charge identification

capabilities.

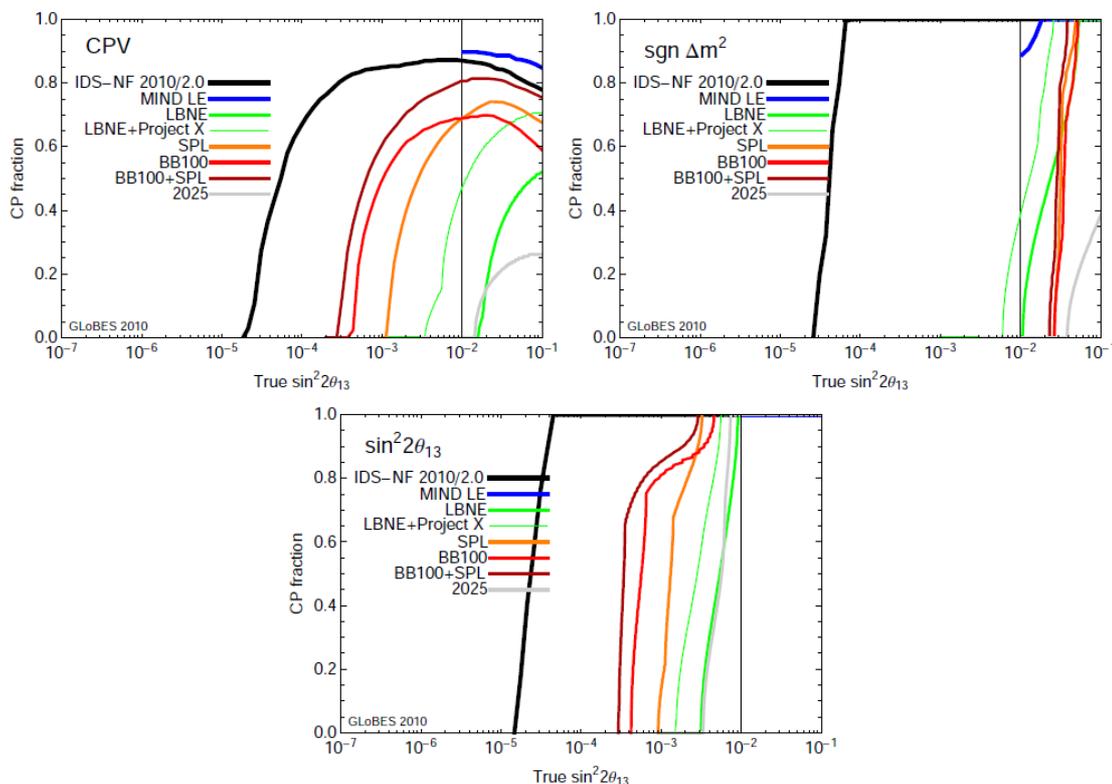


Figure 2.2: The physics reach of a range of possible future facilities. IDS-NF refers to a baseline neutrino factory described in this chapter, with a 100 kTon magnetised iron neutrino detector (to be described in section 2.3.7.2) at 4000 km from the storage rings and a similar 50 kTon detector positioned at a distance of 7500 km. BB100 refers to a Beta-Beam with $\gamma = 100$. MIND LE refers to a single baseline neutrino factory optimised for large $\sin^2 2\theta_{13} > 0.01$. The curve 2025 refers to the generation of experiments presently under construction. The remaining acronyms are consistent with those used in this thesis. Figure from [70].

Figure 2.2 illustrates the discovery reach of the neutrino factory, assuming a design baseline of a 100 kTon magnetised iron neutrino detector (to be described in section 2.3.7.2) at 4000 km from the storage rings and a similar 50 kTon detector positioned at a distance of 7500 km [70]. Each distance is referred to as a ‘baseline’ and the latter distance is known as the ‘magic baseline’ where solar and interference terms cancel in equation 1.40 leaving only a term dependent on θ_{13} . Similarly the ideal baseline for δ_{CP} discovery is around 4000 km. Figure 2.3 summarises the current predicted discovery sensitivities as a function of baseline.

The beta beam is a possible alternative facility which produces beams of neutrinos by accelerating beta unstable radioactive ions, rather than muons, and allowing them to decay

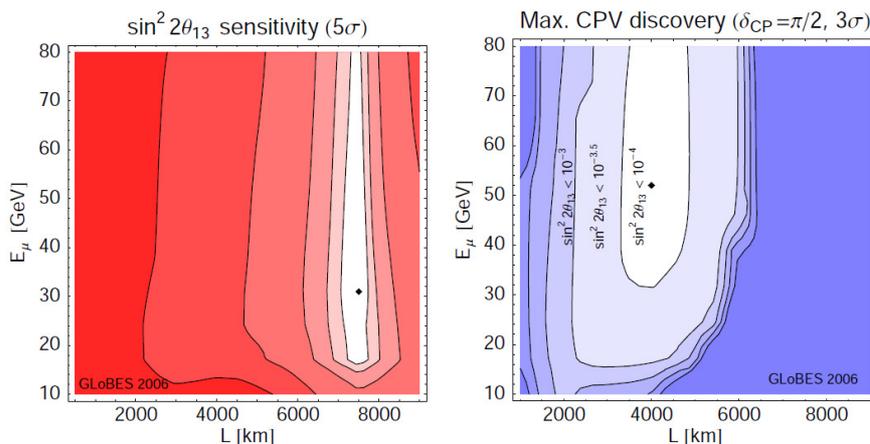


Figure 2.3: Predicted discovery sensitivities. The left plot illustrates the $\sin^2 2\theta_{13}$ sensitivity as a function of energy and baseline, with the diamond indicating optimal sensitivity of $\sin^2 2\theta_{13} \simeq 5.0 \times 10^{-4}$ with contours showing factors of 0.5, 1, 2, 5 and 10 of optimal sensitivity. The right plot illustrates the potential for discovery of maximal CP violation with respect to $\sin^2 2\theta_{13}$. Note from the plots that some flexibility exists of order a few hundred kilometres within the optimal region around the diamond point. [72]

$\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$	$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$	Description	Measurement
$\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$	$\nu_\mu \rightarrow \nu_\mu$	Disappearance	Atmospheric parameters
$\bar{\nu}_\mu \rightarrow \bar{\nu}_e$	$\nu_\mu \rightarrow \nu_e$	Appearance: “Platinum” channel	New oscillation phenomena
$\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$	$\nu_\mu \rightarrow \nu_\tau$	Appearance	Atmospheric parameters
$\nu_e \rightarrow \nu_e$	$\bar{\nu}_e \rightarrow \bar{\nu}_e$	Disappearance	Atmospheric parameters
$\nu_e \rightarrow \nu_\mu$	$\bar{\nu}_e \rightarrow \bar{\nu}_\mu$	Appearance: “Golden” channel	θ_{13} , mass hierarchy, CP Violation
$\nu_e \rightarrow \nu_\tau$	$\bar{\nu}_e \rightarrow \bar{\nu}_\tau$	Appearance: “Silver” channel	New oscillation phenomena, resolve degeneracies

Table 2.1: Experimental oscillation channels in a neutrino factory.

in a storage ring producing pure beams of ν_e (from β^+) or $\bar{\nu}_e$ (from β^-) with small divergence [66]. In the proposed standard configuration, ^{18}Ne and ^6He ions (for β^+ and β^- beams respectively), are produced from proton interactions with an ISOL (Isotope Separation On-Line) target. Since the neutrino possesses only a small share of the energy of the parent ion, it is necessary to provide substantial acceleration to the ion. The majority of beta beam design work to date has centred on exploitation of existing CERN infrastructure to produce a beta beam which accelerates ions to $\gamma \sim 100$, illustrated in figure 2.4. The design target

for such a facility is production of 2.9×10^{18} anti-neutrinos and 1.1×10^{18} neutrinos per year and research and development is currently underway within the EURISOL [73] and EUROnu [60] frameworks to achieve this performance. A beta beam also presents a need for careful consideration of strategies for dealing with degenerate solutions. The neutrino factory is able to resolve these through simultaneous treatment of both muon polarities and its dual baseline architecture.

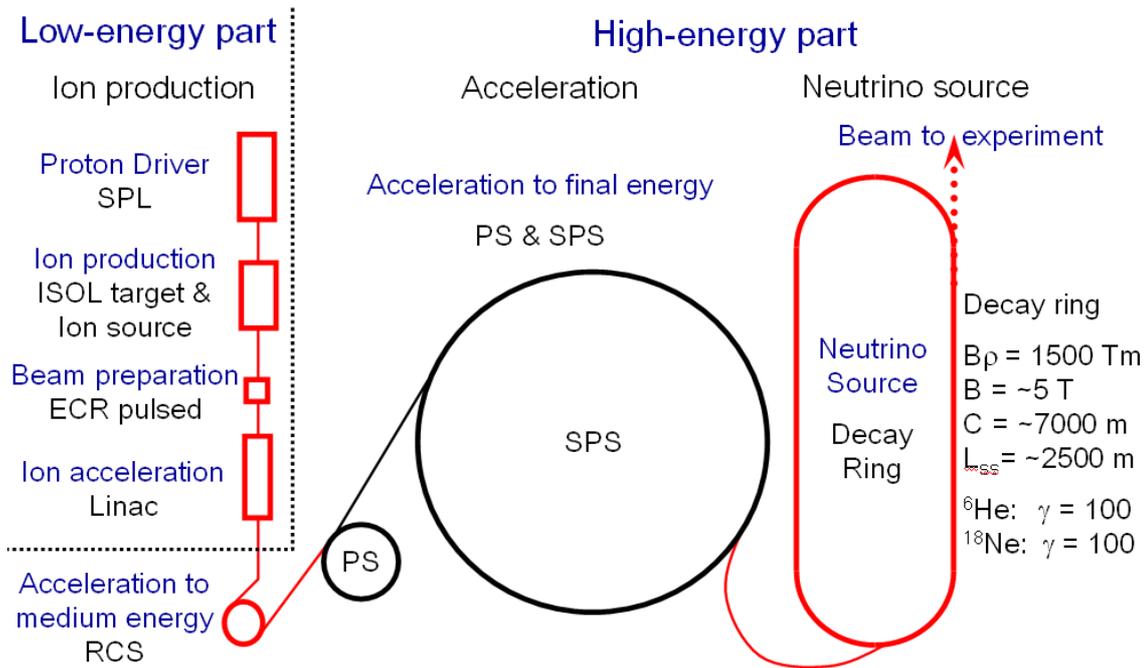


Figure 2.4: A possible design of a future beta beam facility, a rival proposal to the neutrino factory.[74]

Superbeams also offer an alternative to, or possible venue for interim study in advance of, a neutrino factory. Studies are presently being undertaken as part of EUROnu [60] and the Large Baseline Neutrino Experiment (LBNE) frameworks [75] into horn-focused pion beams which produce neutrinos through decay, to detectors at distances of 130 km [76] (EUROnu) or 1,300 km (LBNE) [62]. In common with a neutrino factory, the Superbeam requires a target capable of withstanding a multi MW proton beam. Neutrino beams produced at a superbeam source would be aimed at a water Cherenkov detector of 200 (LBNE) or 440 (EUROnu) kTon fiducial volume, however the LBNE study retains the option of a liquid argon detector with a smaller mass. The EUROnu study specifies a proton energy of 5 GeV

whilst the LBNE study specifies 120 GeV. We note that the beta beam offers a pure beam with well understood systematics, whereas as superbeam cannot exclude some contamination. Similarly a neutrino factory offers an opportunity to study a pure beam but with improved flux, reduced statistical errors and the advantage of two baselines to reduce degeneracies. It is clear from figure 2.2 that the neutrino factory outperforms all rival future facilities for each of our parameters of interest, over the greatest parameter space. Accordingly the neutrino factory is increasingly promoted as ‘the facility of choice’ as the community looks ahead to future venues for advancement of our understanding of neutrino oscillations.

2.3 Design Description

2.3.1 Proton Driver

The parameters of the proton driver are imposed by the goal of 10^{21} muons per year. In order to maximise pion and therefore muon yield the proton beam energy must be in the range 5-15 GeV depending on the choice of target, with an average power of 4 MW and a repetition rate of 50 Hz. This poses significant challenges to the design of an appropriate target, detailed in section 2.3.2. The period of the proton beam will consist of three short bunches with 1-3 ns RMS at intervals of $120 \mu\text{s}$. A well understood timing structure is necessary for efficiency in capture and acceleration of muons. These figures are summarised in table 2.2. Note that recent HARP results [77] suggest only 6-8 GeV is necessary for tantalum and lead (high Z atomic number) targets and 3 GeV for beryllium, copper and aluminium (low Z) targets.

Parameter	Value	Motivation
Proton Kinetic Energy	5-15 GeV	Pion production, space charge
Beam Power	4 MW	High neutrino yield
Repetition Rate	50 Hz	Radio-frequency cavities, kickers, space charge
Bunches per Train	3	Space charge
Total Time for Bunches	$240 \mu\text{s}$	Beam loading
Bunch Length at Target	1-3 ns RMS	Muon capture, space charge
Beam Radius	1.2 mm RMS	Target size
Geometric Emittance	$< 5 \mu\text{m}$	Focus at target

Table 2.2: Proton driver parameters, taken from [78].

Both linear and ring-based accelerators are under consideration. A linear option may exploit the proposed Project X facility, which would provide multi-MW proton beams at Fermilab [79]. Alternatively, a neutrino factory proton driver could make use of a synchrotron or series of synchrotrons already in existence such as the ISIS synchrotron at the Rutherford Appleton Laboratory or the CERN [70] however both would require significant upgrades.

2.3.2 Target

The target will be met by 3 pulses from a 4 MW proton beam at a repetition rate of 50 Hz, a significant challenge to the integrity of the system. The present baseline calls for a free mercury jet target, the effectiveness of which has been tested successfully in the MERIT experiment [80]. The design is illustrated in figure 2.5 and the expected performance of a mercury target when compared with other targets is shown in figure 2.6.

There exists considerable experience with solid targets. Solid targets consisting of bars exchanged between pulses have been shown [81] to survive the impact of the beam and promise simpler mechanisms for handling radiation exchange. However, such a system would require a greater operational volume than afforded by the pion capture and decay system and place restrictions upon the strength of magnetic field. Conventional liquid targets in the case of candidate materials such as mercury could, through production of pressure waves in the liquid caused by interaction with the beam, compromise pipes containing the liquid. Other targets which are being or have been considered include a metal powder jet [82], a rotating tantalum toroid [83], a static or fluidised packed bed target [84, 85] and a rotating system of tungsten bars exchanged between pulses [86].

The target system will be supported by a mercury pool proton beam dump, the vessel for which will be replaced periodically due to high radiation damage. Having produced pions from collisions of protons with the target, a solenoidal capture system provides a decay length for the pions to decay to muons.

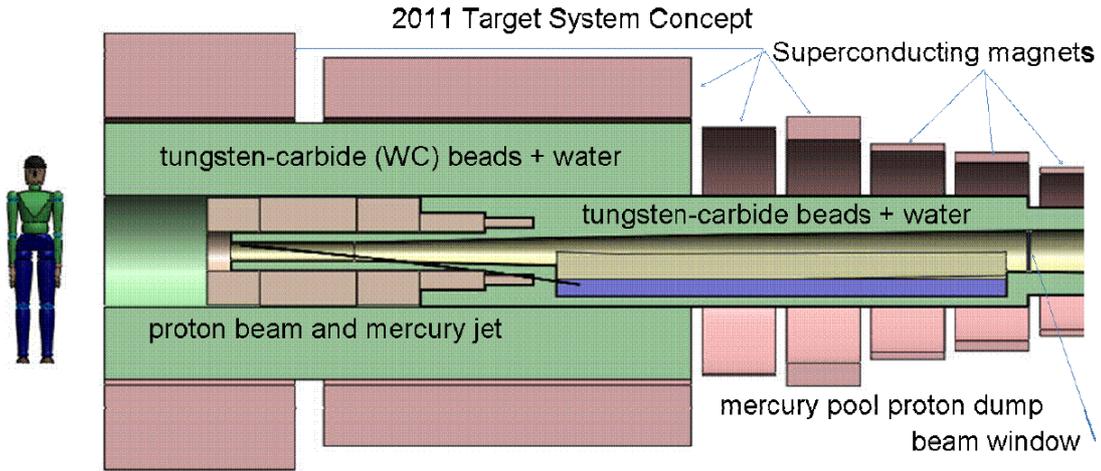


Figure 2.5: The proposed updated baseline of the neutrino factory target based on a free flowing mercury jet placed within a solenoidal magnet. The acronym SC here stands for superconducting magnet. Image taken from [87].

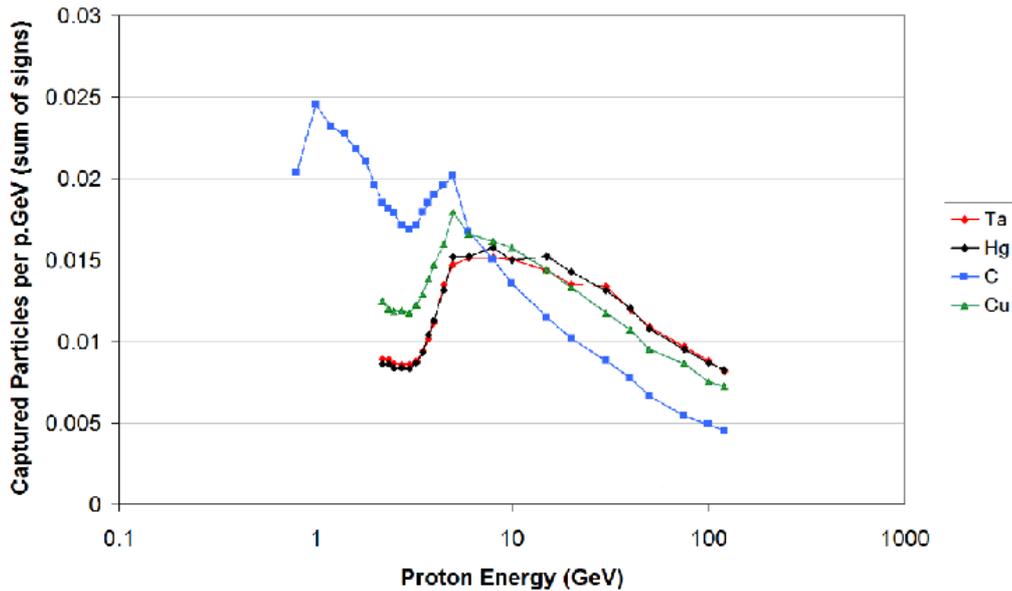


Figure 2.6: The muon yield of various target materials, taking into account the neutrino factory acceptance. Plot from [88] with discussion in [77].

2.3.3 Bunching and Phase Rotation

The target, buncher, phase rotation and cooling sections of the facility are referred to collectively as the muon front end. The front end is responsible for suitably matching the beam produced from the target and pion capture system to the acceleration system. At this point, the muons in the inbound beam cover a wide range of energies (100-300 MeV) and a short

range of time intervals (1-3 ns) which makes it infeasible to accelerate them efficiently in the subsequent RF cavities. Figure 2.7 illustrates the longitudinal change in the beam through the neutrino factory drift, buncher and phase rotation sections.

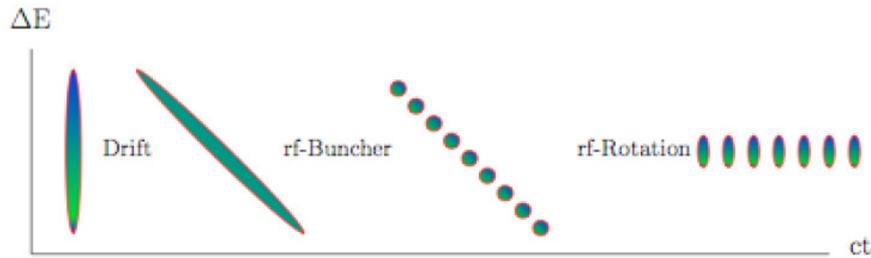


Figure 2.7: The buncher and phase rotation sections of the neutrino factory act to reduce the energy spread and increase the time spread of the muon beam [89, 90]. Figure from [70].

In the first stage, muons develop a position-energy correlation over a 57.7 m drift space. The buncher then has the responsibility of separating the bunch into ‘micro-bunches’ of equally spaced, discrete energies through a sequence of RF-cavities. The bunch train will be less than 80 m long.

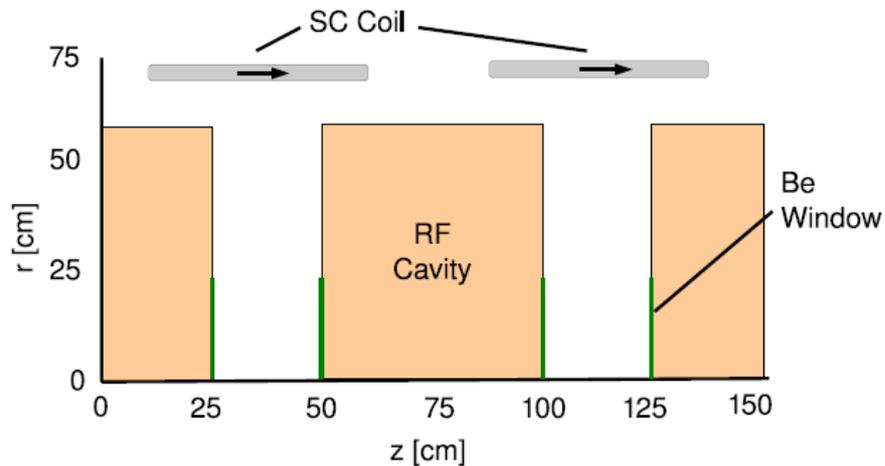


Figure 2.8: A diagram of a single cell of the phase rotator. Figure from [70].

The rotator consists of an additional chain of 56 RF cavities, 42 m in total length which deliver a variable RF phase to each of the micro-bunches to accelerate (in the case of early bunches) or decelerate them (in the case of late bunches) to an equal spacing in ct and a uniform central energy at $p=233$ MeV/c. Over the course of the phase rotator’s length the

cavity frequencies reduce from 230.2 MHz to 202.3 MHz. A single cell is illustrated in figure 2.8.

The phase rotator increases the number of muons meeting the acceptance of the acceleration system by a factor of 4 and can produce bunches of both muon polarities to similar intensities.

2.3.4 Muon Cooling

The muon beam is a tertiary beam produced from the decay of pions which themselves result from interaction of protons with the target. At the point where the muons are created, the divergence of the beam is significant. Work must be done to reduce the transverse size of the beam to fit within the acceptance of a feasible acceleration system.

The limited lifetime of the muon excludes traditional cooling techniques such as stochastic cooling. A new technique is required which is able to substantially reduce the phase space volume of the beam in a short time. Ionisation cooling, through alternating series of absorbers and accelerating cavities works by reducing the transverse and longitudinal size of the beam through energy loss in the absorbers and subsequently restoring the longitudinal component in the accelerating cavities. Ionisation cooling is explained in further detail in section 3.3.

The need for cooling is predicated upon the size of the muon beam leaving the muon front end relative to the acceptance of the acceleration system. Clearly, if the accelerator acceptance were not a constraint then cooling would be unnecessary. If the acceleration system cannot support a beam size of roughly 45π mrad then some length of cooling channel is required. The cost of the accelerator increases with acceptance, resulting in a cost optimisation between the length of cooling channel and accelerator acceptance.

Figure 2.9 is taken from an early and preliminary study [91] which shows minimum cost in the absence of cooling for high acceptance accelerators. However, Feasibility Study 2 [71] calls for a transverse acceptance of 30π mm-rad in the non-scaling FFAG accelerators. According to figure 2.9 this would require 80 m of cooling channel to produce the required

$0.17 \frac{\mu}{\text{proton}}$ which is necessary to meet the goal of $10^{21} \frac{\nu}{\text{year}}$. In the co-ordinate system of figure 2.9 the cooling channel would begin at 215 m and end at 295 m, reducing the required acceptance from 45 to 30 mm rad.

Beyond this acceptance one encounters longitudinal phase space distortions on account of the relationship between transverse amplitude and time of flight.

2.3.5 Muon Acceleration

Acceleration must occur within a fraction of the muon lifetime of $2.2 \mu\text{s}$. This places significant restrictions upon a design solution for the acceleration system, effectively excluding traditional synchrotrons. The baseline acceleration system for the neutrino factory is shown in figure 2.10. A linear accelerator, dual recirculating linear accelerators (RLAs) and Fixed Field Alternating Gradient (FFAG) accelerator combination is favoured. The linear accelerator has a greater phase space acceptance than the RLAs and increases the beam energy from 0.24 GeV to 0.9 GeV whilst decreasing the beam size. The RLAs have a dogbone geometry where magnetic arcs with varying magnetic field settings facilitate multiple passes of the radio-frequency (RF) cavity acceleration straights (shown in red in figure 2.10), reducing the necessary cost of each section. The arcs require the reduced beam size and momentum spread assured by the preceding linear accelerator. The first RLA increases the beam energy from 0.9 GeV to 3.6 GeV, and the second further increases this to 12.6 GeV. Each RLA is composed of superconducting RF cavities and quadrupoles, whereas the linear accelerator is a series of cryo-modules of superconducting RF cavities and iron-shielded solenoids. Transfer from the linear accelerator to the first RLA, and from the first RLA to the second, is via two double chicanes each of which combines a vertical dipole (for spreading), horizontal bending magnets, quadrupoles and an additional vertical dipole (for combining the beam).

The beam is injected into the non-scaling FFAG accelerator [93] at 12.6 GeV. Scaling FFAG accelerators produce a magnetic field which increases as a power law of radius from the beam axis centre whereas the fields associated with non-scaling FFAGs vary linearly with radius. Non-scaling FFAGs offer a reduced orbital excursion, which is desirable, and are

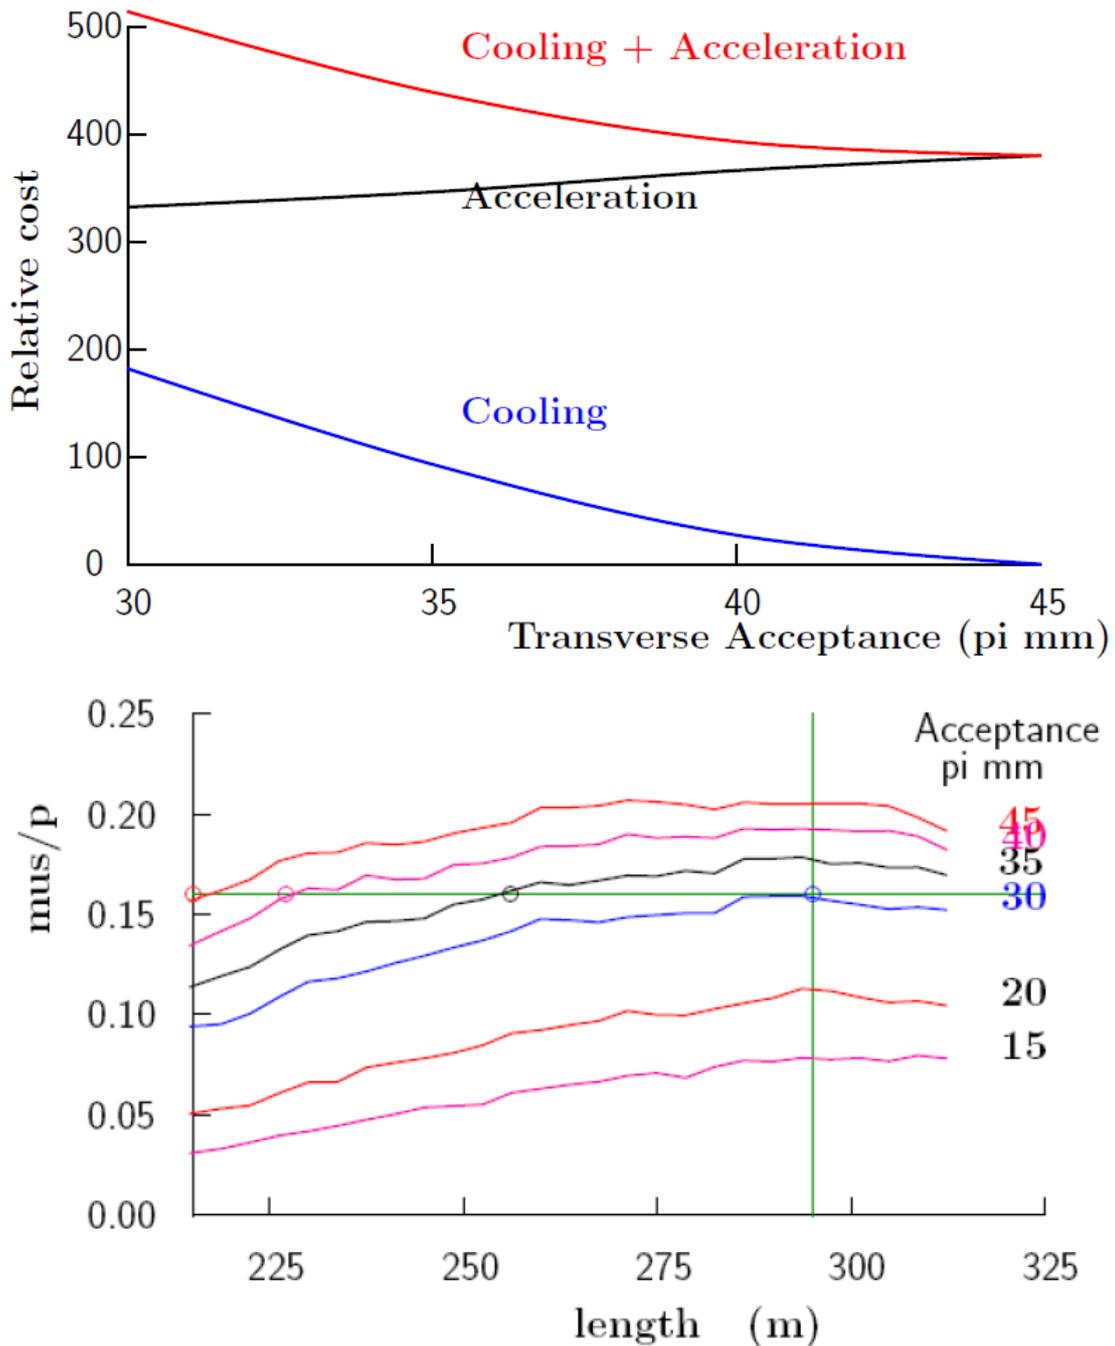


Figure 2.9: Above: Relative cost with and without cooling. Below: Muons per proton within the accelerator acceptance as a function of cooling channel length for a variety of acceptance profiles [92]. The cooling channel begins at 215 m, which is the distance from the target.

compatible with 200.15 MHz RF cavities they nevertheless presently feature a lower acceptance than their scaling equivalents. Whilst a smaller acceptance provides a cost advantage, injection becomes more complex in a confined space.

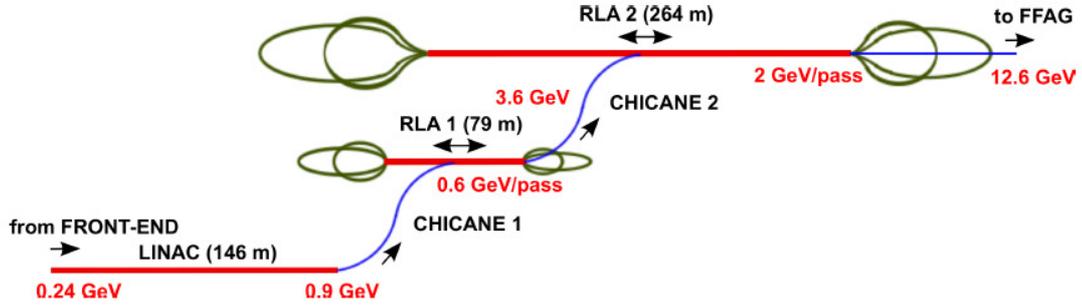


Figure 2.10: A view of the accelerating path within a neutrino factory, taken from [70]. Radio-frequency (cavity) acceleration straights are shown in red.

2.3.6 Muon Storage

The Neutrino Factory baseline calls for a dual racetrack shaped muon storage ring system (Figure 2.11) with a capacity of 10^{21} muons, in order to store both polarities of muon. Muon bunches are injected into either ring according to their polarity. Each ring has a 600.2 m decay straight which points into the ground toward a distant detection facility and may be positioned up to 444 m underground in the case of a 7500 km baseline. The long straight sections allow the muon to decay creating a beam of neutrinos described by the following decays:

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \quad (2.1)$$

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu \quad (2.2)$$

Those muons which do not decay are matched in the arcs and traverse the return straight where they undergo tune control and collimation. The racetrack geometry features a 37% efficiency in producing neutrinos in the desired trajectory.

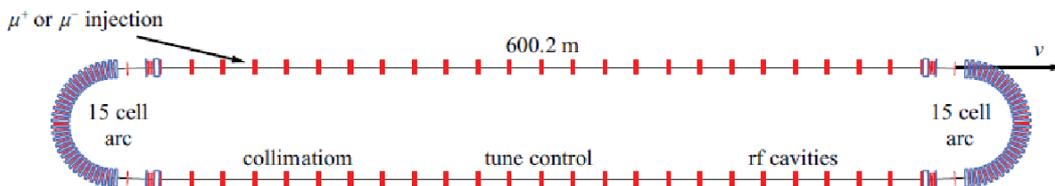


Figure 2.11: The racetrack design of the neutrino factory muon storage rings. The top decay straight produces a neutrino beam from the decay of muons. The return straight performs acceleration, tuning and collimation. Each decay ring may store either polarity of muon. Figure from [70].

2.3.7 Detectors

2.3.7.1 Near Detectors

One to four near detectors are specified in the neutrino factory baseline, corresponding to each of the straight sections of the muon storage rings. Near detectors allow for characterisation of the neutrino beam in the absence of oscillation effects, in addition to measurements of neutrino flux, neutrino-nucleon scattering cross-sections, charm production background, electroweak parameters such as the weak mixing angle $\sin^2\theta_W$, parton distribution functions and new physics through non-standard interactions. Measurement of the charm background is needed since it is one of the most pronounced background signals camouflaging the oscillation measurements at a far detector and can be measured effectively with a vertex detector.

At present there are two candidates for the high resolution section of the near detector: a scintillating fibre tracker and a transition-radiation straw-tube tracker. Work is currently underway to predict the performance of each [70].

2.3.7.2 Far Detectors

The Magnetised Iron Neutrino Detector (MIND) research and development effort has made significant progress describing the performance and parameters of the neutrino factory far detectors, illustrated in figure 2.12 and parametrised in table 2.3. Two far detectors based on the MIND design are required to measure δ_{CP} and θ_{13} with optimal precision, placed in line with the decay straights of the muon storage rings. The first will be situated roughly 4,000 km from the neutrino factory and the second at roughly 7,500 km. The choice of

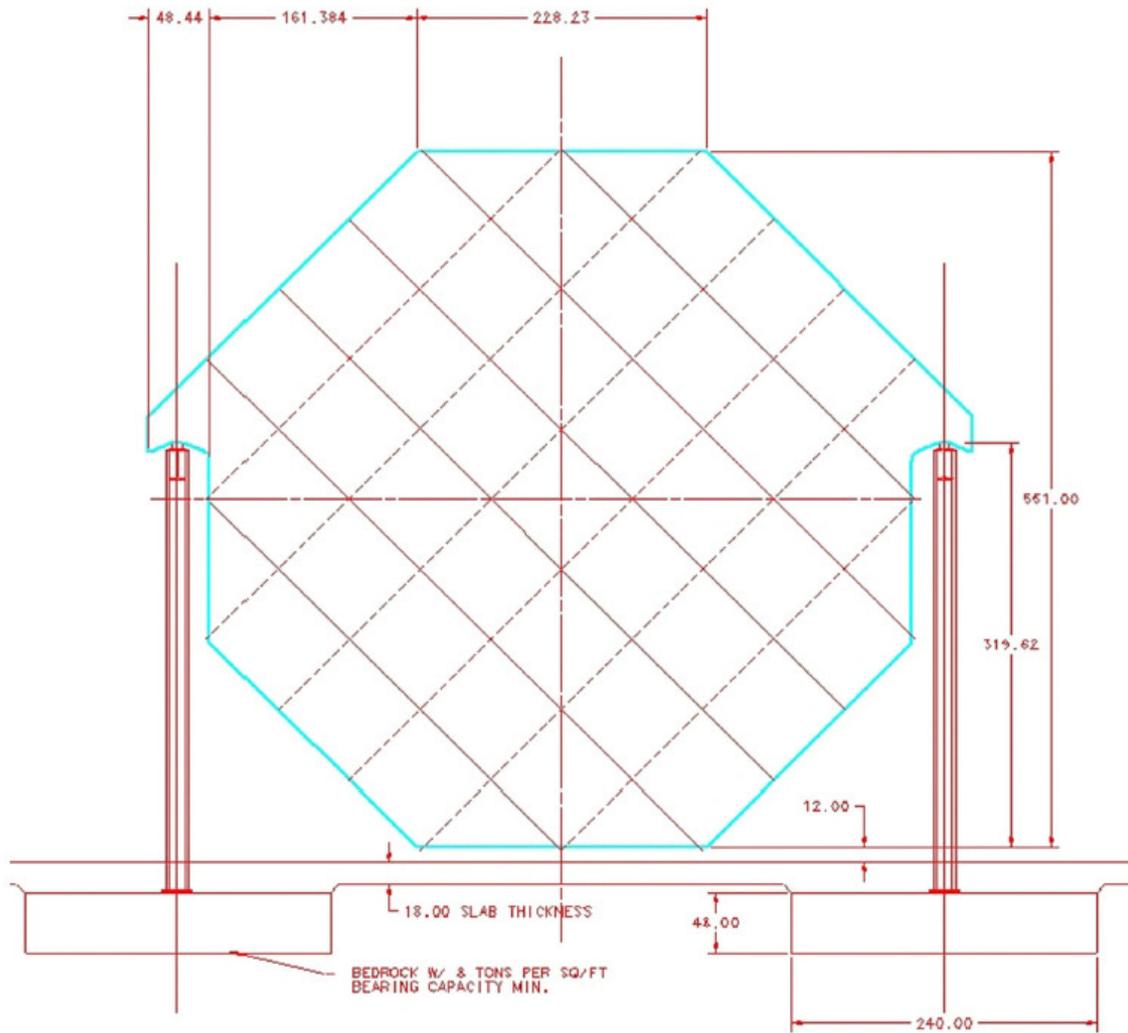


Figure 2.12: A single plate in the Magnetised Iron Neutrino Detector (MIND).

two far detectors is driven by the need to eliminate degeneracies. The 7,500 km baseline is referred to as the ‘magic’ baseline since it is at this distance that CP-violation and matter effects cancel, providing an opportunity to measure θ_{13} in a clean way.

The MIND detectors observe neutrino oscillation through detection of muons with opposite polarity to those present in the corresponding storage ring at that time. In so doing MIND must contend with background signals from meson decays in hadronic showers, misidentification of neutral current events and misidentification of muon charge. The expected performance of MIND as a function of energy is illustrated in figure 2.13. MIND can readily distinguish between a muon and the hadronic particles inherent in the background as muons are Minimum Ionising Particles (MIPs) which will travel a greater distance in iron. The ex-

Parameter	MIND 1	MIND 2
Distance (km)	3000 – 5000	7000 – 8000
Fiducial Mass (kTon)	100	50
Size Iron Plates (cm ³)	1500 × 1500 × 3	1500 × 1500 × 3
Length of Detector (m)	125	62.5
Number of Iron Plates	2500	1250
Dimensions of Scintillator Bars (cm ³)	1500 × 3.5 × 1	1500 × 3.5 × 1
Number of Scintillator Bars per Plane	429	429
Total Number of Scintillator Bars	2.14 × 10 ⁶	1.07 × 10 ⁶
Total Number of Readout Channels	4.28 × 10 ⁶	2.14 × 10 ⁶
Photon Detector	SiPMT	SiPMT
Magnetic Field (T)	> 1	> 1

Table 2.3: The current MIND baseline [70]

pected discovery reach enabled by MIND for $\sin^2 2\theta_{13}$ and δ_{CP} is between 10^{-5} and 10^{-4} . For further detail on MIND reconstruction and analysis the reader is referred to [70, 94].

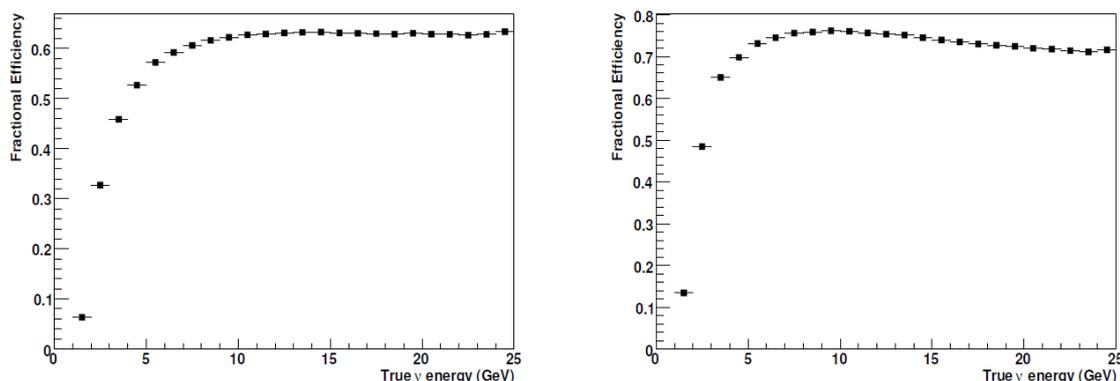


Figure 2.13: Figure by Laing [94] showing efficiency of signal identification of ν_μ (left) and $\bar{\nu}_\mu$ (right) charged current interactions as a function of energy.

2.4 Future Work

There exists a collaboration, the International Design Study for a Neutrino Factory (IDS-NF)[95], whose mandate is to deliver a Reference Design Report for a Neutrino Factory by 2012/13. This document must include a detailed specification of the physics sensitivity at a neutrino factory as well as of the detector, accelerator and diagnostic subsystems required to realise the facility. The document, whilst not recommending a specific site, will provide cost estimates and a detailed schedule which will open the gate for a decision on funding. An In-

terim Design Report (IDR) has already been published and is freely available [70]. The IDR details the significant research and development milestones that have been reached to date, including but not limited to substantial research and development work towards a suitable target converging on the current mercury jet baseline and an analysis of the performance of MIND.

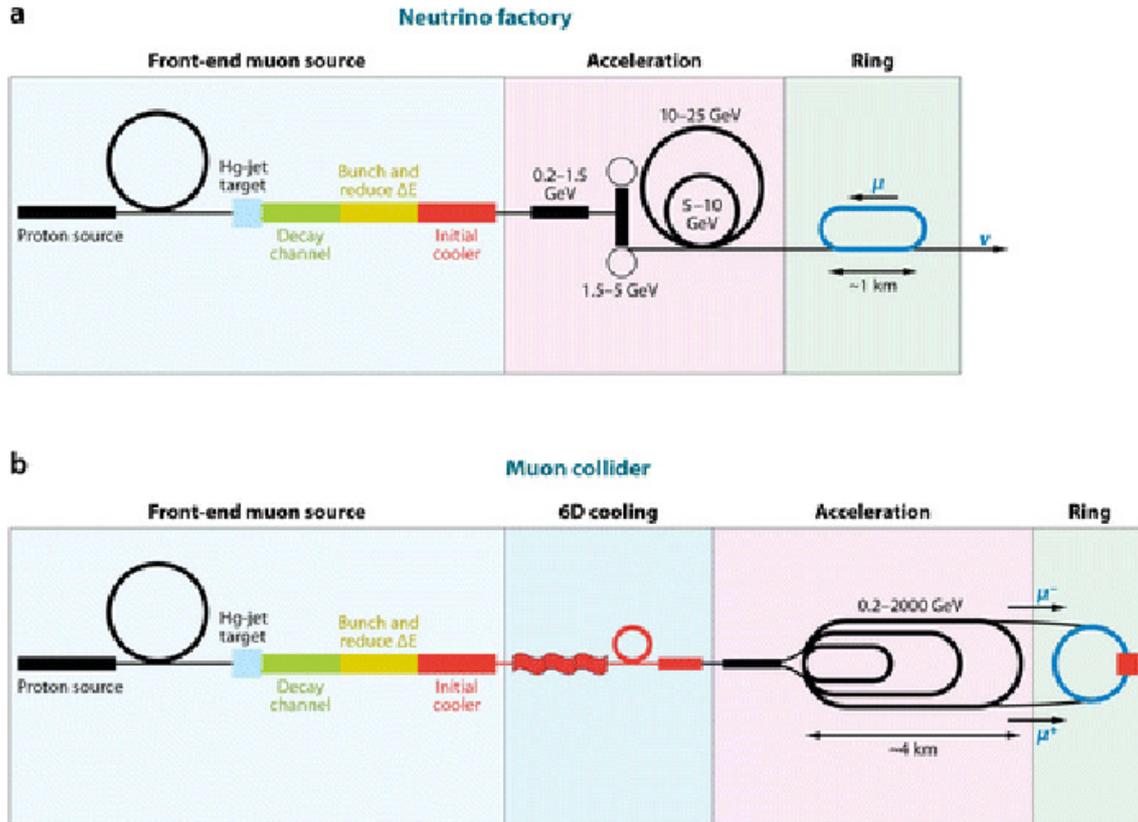


Figure 2.14: Diagrams of (a) the neutrino factory and (b) muon Collider which illustrate their similarity upstream of the muon storage rings. The present baseline of the neutrino factory front-end is identical to that required for a muon Collider. Image sourced from [96].

There is considerable interest in a muon collider, a future facility which could be an upgrade to a neutrino factory to deliver an accelerator at the energy frontier. The schematic of a muon collider, shown in figure 2.14 and originally described in [97] includes significant overlap with a neutrino factory, with a similar front end. The current research and development efforts are therefore strongly correlated. Cooling for example is a requirement for both facilities, however 6D cooling is required for a muon collider. A muon collider is a multi-TeV centre of mass energy scale lepton anti-lepton collider and could be realised as

part of a future neutrino factory upgrade program.

The point-like nature of the muon enables low backgrounds at a muon collider when compared with hadron colliders operating on a similar energy scale. In addition, the higher mass of the muon when compared with the electron ($\frac{m_\mu}{m_e} \approx 211$) provides an advantage to the design of the muon collider over rival lepton-anti-lepton facilities offering energies comparable to hadronic facilities in circular rings smaller than electron colliders. The muon also produces far less synchrotron radiation than electrons, making a circular accelerator possible. A linear collider is the only acceptable alternative for an electron-positron accelerator operating at the TeV level.

Further research and development is required to realise a muon collider and define a stable baseline optimised for cost efficiency and performance. Among these challenges is the need for a higher performance cooling channel than is required for a neutrino factory.

2.5 Conclusions

Effort is presently underway to cost and schedule a future neutrino factory facility which will provide unparalleled reach into the δ , θ_{13} , Δm^2 parameter space. A summary of the current baseline for the proton driver, target, muon cooling, muon acceleration, muon storage and neutrino detection subsystems has been presented and discussed. The wider context of an effort to realise a post-LHC high energy lepton-anti-lepton collider has also been introduced, with reference to its synergy with neutrino factory research and development. At present the next major milestone for the project is the delivery of a Reference Design Report by 2012/13 which will provide a basis for delivery of a future neutrino factory facility, providing a detailed study of the cost, engineering and sensitivities required. An essential part of this work is the demonstration of ionisation cooling and the construction of a single cooling cell. This is the goal of the Muon Ionisation Cooling Experiment (MICE) which will be discussed in the following chapter.

Chapter 3

The Muon Ionization Cooling Experiment (MICE)

3.1 Introduction

The Muon Ionisation Cooling Experiment (MICE) is critical to the design of a future Neutrino Factory facility. In order to produce sufficiently intense beams of muons within the acceptance of the acceleration system, one must demonstrate a means of reducing the transverse size of the beam in a single pass. Conventional cooling techniques, useful for long lived particles such as protons, will not prove effective here. Ionisation cooling is a never before demonstrated technique that is effective over the comparatively short muon lifetime of $2.2\mu s$. Although the physics of ionisation cooling, described in section 3.3, is not in doubt, it has never been demonstrated in practice and MICE will provide a venue for the first ever systematic study.

In addition to assessing the suitability of ionisation cooling for the Neutrino Factory, MICE will design, engineer and build a fully working prototype of an ionisation cooling cell, a significant milestone towards a neutrino factory front end and a crucial contribution towards the cost of the full facility.

MICE is supported by an international collaboration of institutions, with the Rutherford Appleton Laboratory in Oxfordshire acting as the host laboratory. The MICE beamline, described in section 3.5, operates parasitically to the ISIS¹ synchrotron.

3.2 Emittance

The Neutrino Factory design must present a tertiary muon beam within the acceptance of the acceleration system. This beam is the product of pion decay and the pions are in turn produced from the collisions of protons with the neutrino factory target. The phase space volume of the beam is referred to as its emittance and is a useful parameter when considering the passage of a beam through an accelerating system.

We refer to statistical mechanics and invoke Liouville's theorem which holds that such a volume remains constant in the absence of non-conservative forces. Therefore, emittance is a conserved quantity in the absence of acceleration. The area which is occupied by the beam is $\pi\epsilon$, where ϵ is the emittance.

Following [98], we contain the particles of a beam in phase space with a two-dimensional ellipse, shown in figure 3.1 and parametrised in the Courant-Snyder invariant [99] by the familiar Twiss parameters α , β and γ :

$$\gamma x^2 + 2\alpha x x' + \beta x'^2 = \epsilon, \quad (3.1)$$

where the Twiss parameters are defined in the usual way:

$$\beta = \frac{\langle x^2 \rangle}{\epsilon}, \quad (3.2)$$

$$\gamma = \frac{\langle x'^2 \rangle}{\epsilon}, \quad (3.3)$$

$$-\alpha = \frac{\langle x x' \rangle}{\epsilon}. \quad (3.4)$$

¹ISIS is not an acronym but rather is named after the Thames River's passage through Oxfordshire.

Here $\langle x^2 \rangle$, $\langle x'^2 \rangle$, $\langle xx' \rangle$ denote variances in the position variable x and $x' = \frac{p_x}{p_z}$. Equation 3.1 is known as the Courant-Snyder invariant. The area of this ellipse is the emittance.

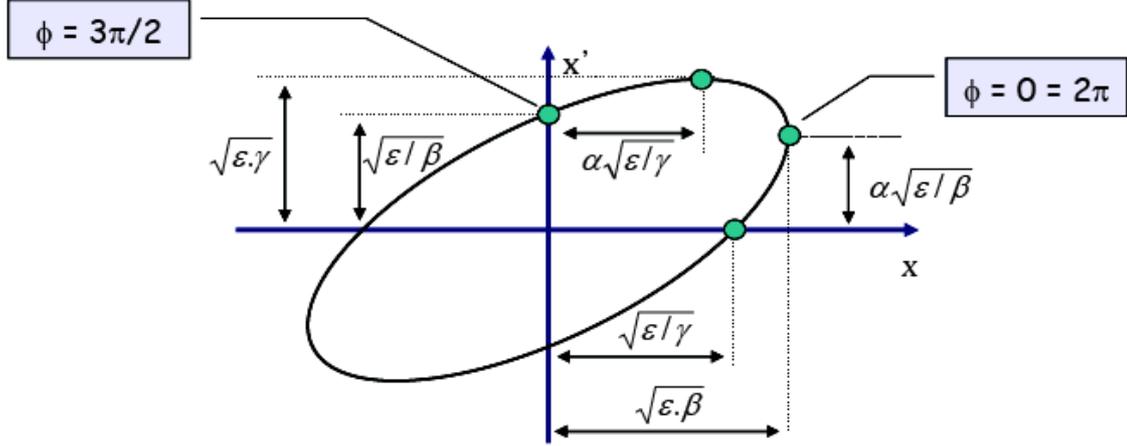


Figure 3.1: A geometrical representation of emittance in phase space [100].

The equation for an ellipse of any number of dimensions may be written in matrix form as:

$$\mathbf{u}^T \sigma^{-1} \mathbf{u}, \quad (3.5)$$

taking σ as a symmetric matrix yet to be defined. In two dimensions, $\mathbf{u}_{2D} = (x, x')$, and the equation for the ellipse is:

$$\sigma_{22}x^2 + 2\sigma_{12}xx' + \sigma_{11}x'^2 = \epsilon^2. \quad (3.6)$$

We define σ after comparison of equations 3.1 and 3.6:

$$\sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{pmatrix} = \epsilon \begin{pmatrix} \beta & -\alpha \\ -\alpha & \gamma \end{pmatrix} \quad (3.7)$$

The phase space area of the beam is given by equation 3.8.

$$Volume_{2D} = \pi \sqrt{|\sigma|} = \pi \sqrt{\sigma_{11}\sigma_{22} - \sigma_{12}^2} = \pi \epsilon_{2D} \quad (3.8)$$

The conservation of emittance, when defined in this manner, is not valid in the case of

acceleration, described by a process called adiabatic dampening where emittance varies as:

$$\epsilon = \frac{1}{1 + \frac{\Delta E}{E_0} z} \epsilon_0 = \frac{p_0}{p} \epsilon_0, \quad (3.9)$$

where $\frac{\Delta E}{E_0}$ provides the fractional increase in energy over an accelerator length z . In such circumstances we rely upon a truly invariant quantity known as normalised emittance which is given by equation 3.10.

$$\epsilon_n = \beta_L \gamma_L \epsilon = \frac{\epsilon p}{m_0 c} \quad (3.10)$$

β_L and γ_L in equation 3.10 signify the Lorentz factors associated with special relativity ($\beta_L = \frac{v}{c}, \gamma_L = (1 - \beta_L^2)^{-1/2}$) and should not be confused with the Twiss parameters. From this point onwards whenever emittance is discussed, normalised emittance is implied.

It is useful for us to expand our discussion of emittance to four dimensions and six dimensions where we have $u_{4D} = (x, p_x, y, p_y)$ and $u_{6D} = (x, p_x, y, p_y, E, t)$. This description in phase space substitutes the previous one in trace space by replacing the inclination angles with x and y momenta. Assuming a paraxial beam, the angles are defined:

$$x' = \frac{p_x}{p_z} \quad (3.11)$$

$$y' = \frac{p_y}{p_z} \quad (3.12)$$

We may now define transverse emittance as follows:

$$\epsilon_{2D} = \frac{1}{m_\mu c} \sqrt{|V_{2D}|} \quad (3.13)$$

$$\epsilon_{4D} = \frac{1}{m_\mu c} \sqrt[4]{|V_{4D}|} \quad (3.14)$$

$$\epsilon_{6D} = \frac{1}{m_\mu c} \sqrt[6]{|V_{6D}|} \quad (3.15)$$

where $|V_{ND}|$ is the determinant of the N dimensional covariant matrix in phase space related to the one in trace space by $|V| = p^2|\sigma|$. Normalised longitudinal emittance is a two dimensional quantity in the MICE experiment and follows from the definition of a two dimensional (E,t) emittance:

$$\epsilon_{long} = \frac{1}{m_{\mu}c} \sqrt{|V_{2D}|} = \frac{1}{m_{\mu}c} \sqrt{\langle E^2 \rangle \langle t^2 \rangle - \langle Et \rangle^2} \quad (3.16)$$

Depending upon the context, for convenience emittance is often quoted in various units [101] such as mm, mm rad and π mm rad. One may define the area (or some fraction of the area) of the ellipse in figure 3.1 as the emittance or choose the definition in equation 3.8. Where the emittance is calculated from the area of the ellipse (given by $\pi \times A \times B$ where A and B are half the length of the major and minor axes respectively) the factor π is included in the units purely as a flag of the method used rather than a numerical factor. To avoid confusion, throughout this thesis emittance will be quoted in units of mm rad.

3.3 Ionisation Cooling

It is necessary to reduce the transverse emittance of, or “cool”, the muon beam before it is injected into the accelerating system of a Neutrino Factory. A possible successor to the Neutrino Factory, the Muon Collider, relies even more heavily upon beam cooling. The short lifetime of the muon ($2.2 \mu s$) precludes the possibility of stochastic [102], laser or electron cooling [103] and its large mass prohibits radiation dampening as a viable alternative. A new technique, ionisation cooling, is favoured in the case of the muon because of its long interaction length.

Ionisation cooling relies upon particles crossing an absorbing material, where they experience a loss of total energy through ionisation and restoration of the longitudinal component of the momentum through radio-frequency (RF) accelerating cavities immersed in focusing magnets. The net effect, for a properly chosen absorber, is to reduce the transverse phase space of the muon beam, as in figure 3.2. In figure 3.2, graph 1 shows a loss of total mo-

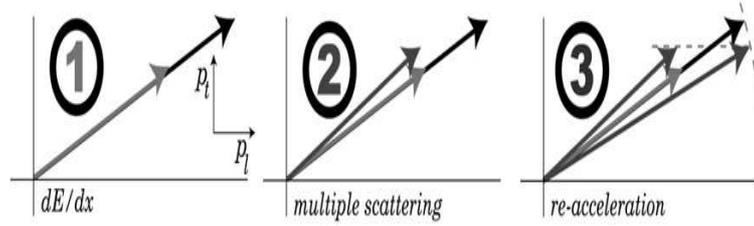


Figure 3.2: These graphs illustrate the principle of ionisation cooling, which MICE aims to demonstrate experimentally.

momentum through energy loss as a particle traverses the absorbing material. Graph 2 shows an increase in transverse momentum caused by multiple scattering. The goal of the cooling channel is to reduce transverse momentum, so an absorbing material with low density is chosen to minimise the effect of multiple scattering. Finally, the particle is re-accelerated in graph 3, where only longitudinal momentum is restored. This results in a reduction of transverse momentum and consequently a reduction of the transverse phase space size of the beam. This technique has yet to be demonstrated experimentally, hence the current interest in MICE. The stated goal of MICE is to measure fractional change of emittance of 10% to an error of 1%, which is extremely challenging.

Any appreciable change in emittance arises from either energy loss or multiple scattering. The former is modelled through the Bethe-Bloch function. The mean squared increase in divergence is described approximately through a Gaussian fit to the Molière scattering formula, the variance of which shall be taken as:

$$\frac{d \langle \Theta_x^2 \rangle}{ds} = \frac{(13.6 \text{ MeV})^2}{(p\beta_L)^2 X_0} \quad (3.17)$$

Here Θ_x is the angular spread in x , p the momentum of the beam or particle considered and X_0 the radiation length. Recalling our definition of two dimensional transverse normalised emittance in equation 3.14, the change in emittance resulting from traversing a given volume of material can be written as:

$$\frac{d\epsilon_n}{ds} = \frac{1}{2m^2\epsilon_n} \frac{d|V|}{ds}, \quad (3.18)$$

where the s is the longitudinal direction in the frame of the beam and the square root of the determinant is given as follows:

$$|V| = (\langle x^2 \rangle \langle p_x^2 \rangle - \langle xp_x \rangle^2) \quad (3.19)$$

Approximating $\Theta_i \approx \frac{p_i}{p_s}$ we can rewrite 3.19 as:

$$|V| = p_s^2 (\langle x^2 \rangle \langle \Theta_x^2 \rangle - \langle x\Theta_x \rangle^2) \quad (3.20)$$

Therefore, when we consider those terms contributing to the change in emittance:

$$\frac{d\epsilon_n}{ds} = \frac{-1}{E} \left| \frac{dE}{ds} \right| \epsilon_n + \frac{1}{2m^2\epsilon_n} \langle x^2 \rangle \frac{(13.6 \text{ MeV})^2}{\beta_L^2 X_0} \quad (3.21)$$

Defining the transverse beta function to be

$$\beta_{\perp} = \frac{\langle x^2 \rangle p}{m\epsilon_n}, \quad (3.22)$$

we may write:

$$\frac{d\epsilon_n}{ds} = \frac{-1}{E} \left| \frac{dE}{ds} \right| \epsilon_n + \frac{(13.6 \text{ MeV})^2 \beta_{\perp}}{2mX_0\beta_L^3 E} \quad (3.23)$$

This describes the effect of transverse ionisation cooling. The equation is the sum of two terms: a cooling term due to energy loss (therefore always negative), which is proportional to emittance, and a heating term due to multiple scattering. β_{\perp} is related to the dispersion angle, with a large β_{\perp} exacerbating multiple scattering. We note that the radiation length X_0 is a function of the Z number of the absorber material:

$$X_0 = \frac{716.4A}{Z(Z+1)\ln\frac{287}{\sqrt{Z}}} \text{ gcm}^{-2} \quad (3.24)$$

The choice of the absorber material is crucial to ensure cooling, rather than heating, of the beam. Heating may result from multiple scattering of particles within the absorbing material, therefore a low Z material with long radiation length should be chosen in order to minimise

the effect of multiple scattering so that the balance is in favour of cooling. Given powerful focusing (provided through superconducting magnets) to reduce β_{\perp} and high gradient RF cavities we have a technique for cooling a muon beam effective over the short muon lifetime and the high muon mass.

3.4 Experimental Setup

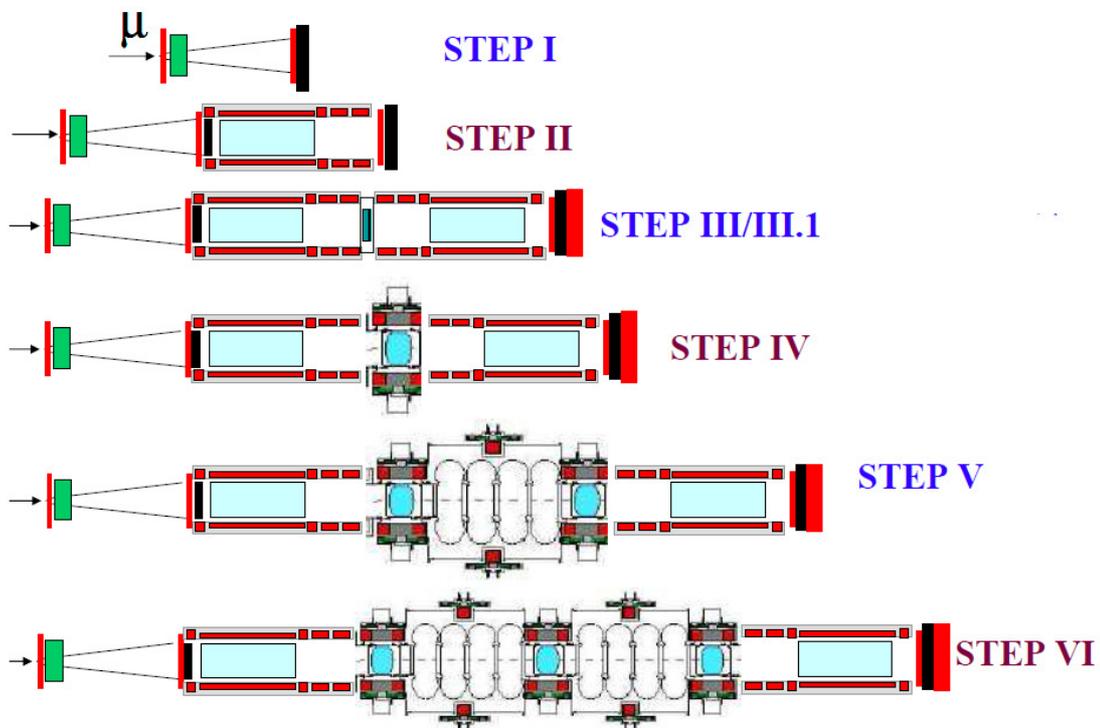


Figure 3.3: The staged progression of MICE is illustrated above. At time of writing, MICE is taking data in a step I configuration [104].

The Muon Ionization Cooling Experiment is a staged experiment, evolving gradually with time to provide a systematic study of ionisation cooling, as illustrated in figure 3.3. The MICE beamline, not pictured, feeds from the ISIS proton source at the Rutherford Appleton Laboratory. Pions are produced from proton collisions with a cylindrical titanium target. These pions are allowed to decay in the beamline and their momenta are selected using two dipole bending magnets.

MICE (figure 3.4) consists of a cooling channel with scintillating fibre trackers upstream

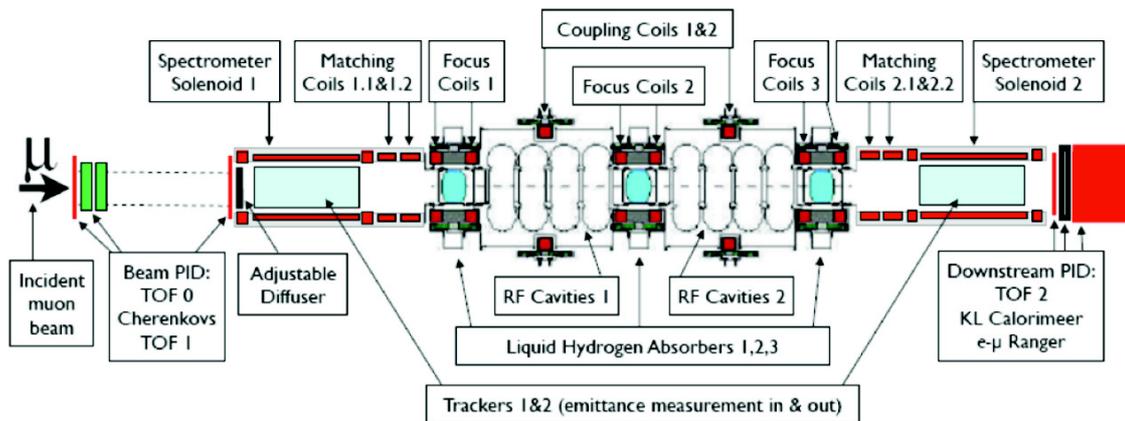


Figure 3.4: A detailed schematic of the final stage of MICE with the full cooling channel consisting of three absorber modules and two RF-cavities, braced by particle identification and tracking detectors. Flags indicate areas of responsibility [104].

and downstream of the channel for emittance measurement. Additional detectors, such as time of flight, Cherenkov detectors and calorimeters (the KL and electron muon ranger detectors) allow for particle identification and selection of a beam. MICE is a single particle experiment, and strong particle identification is crucial in constructing a beam during the offline analyses. The cooling channel itself consists of low density absorber materials such as liquid hydrogen or lithium hydride, surrounded by focusing coils, and radio frequency accelerating cavities surrounded by coupling coils.

Presently MICE is in its Step I configuration. Due to equipment failures, most recently with the spectrometer solenoids, the experimental schedule has been subject to delay and is presently under review. The spectrometer solenoids are the present limiting factor for the transition to step II. Spectrometer 1 was disassembled due to blockages in the re-condensing circuit, prior to failure of a high temperature superconducting (HTS) lead [105]. It is anticipated that the solenoids will be installed in 2011 and, possibly, that Steps II and III may be abbreviated to accelerate the remaining schedule towards Step IV.

3.5 Beamline

The MICE beamline operates parasitically to the ISIS 800 MeV proton synchrotron and is fully operational. Shown in figure 3.5, the design is based on that of a pion-muon decay

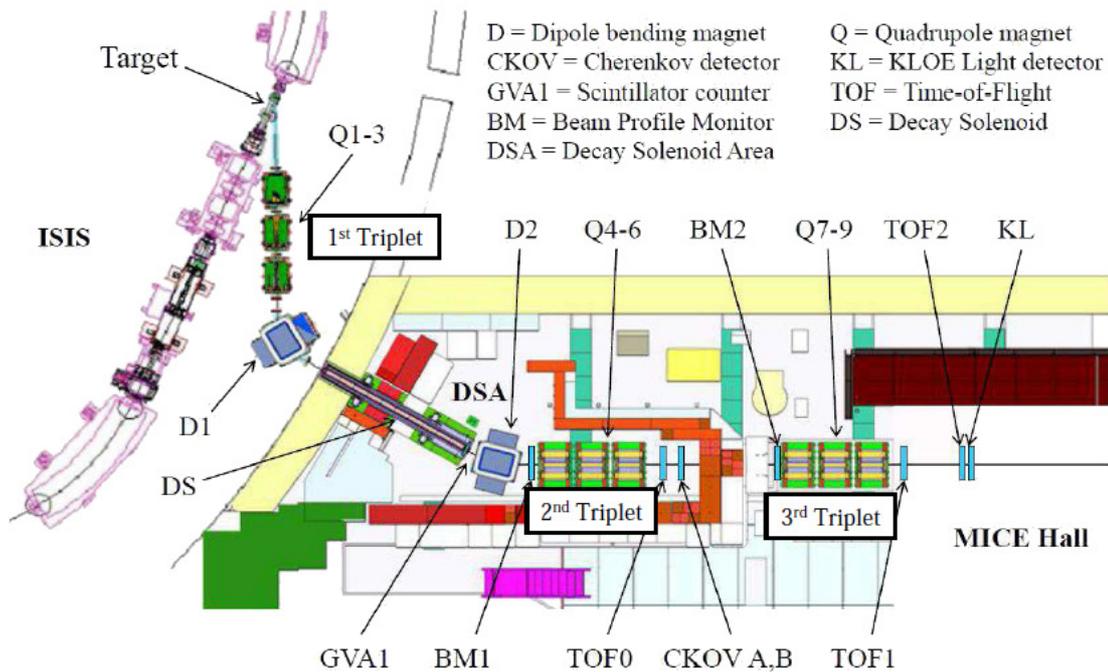


Figure 3.5: Sketch of the MICE beamline running parasitically to the ISIS synchrotron. Beamline counters and particle identification detectors (including the luminosity monitor not shown) were used to commission the MICE muon beam.

channel. The MICE Target is dipped into the halo of the ISIS beam at a rate of a few Hz, producing mainly pions but also deflected protons and muons which are captured in some fraction by the first quadrupole triplet (Type Q4 [106]) and transported to the first of two Type I NIMROD dipole bending magnets which provides momentum-selection, bending desired particles towards the remaining legs of the beamline. A 5T superconducting decay solenoid then captures pions and any muons which may already be present, and allows the pions to decay before a second dipole bending magnet which provides a second momentum cut, tuned to the distinctive momentum associated with muons which decay in the backwards pion hemisphere in the pion centre of mass frame (figure 3.6). This allows muons of a set momentum to be selected, whilst minimising the contamination due to protons.

At this point, the beam is mostly composed of muons. A final pair of quadrupole triplets (Type QC) transport our muon beam to the lead diffuser, which allows for correction between an ideally matched beam and that produced by the beamline. The diffuser has the effect of increasing the emittance of the beam through multiple scattering. Through varying the currents of the quadrupoles and dipoles, experimenters are able to describe characteristics of

the desired beam including beam particle species and momenta.

The beamline is designed to provide a range of normalised emittance settings between 1 and 12 mm rad in a momentum range of 140 to 240 MeV [107].

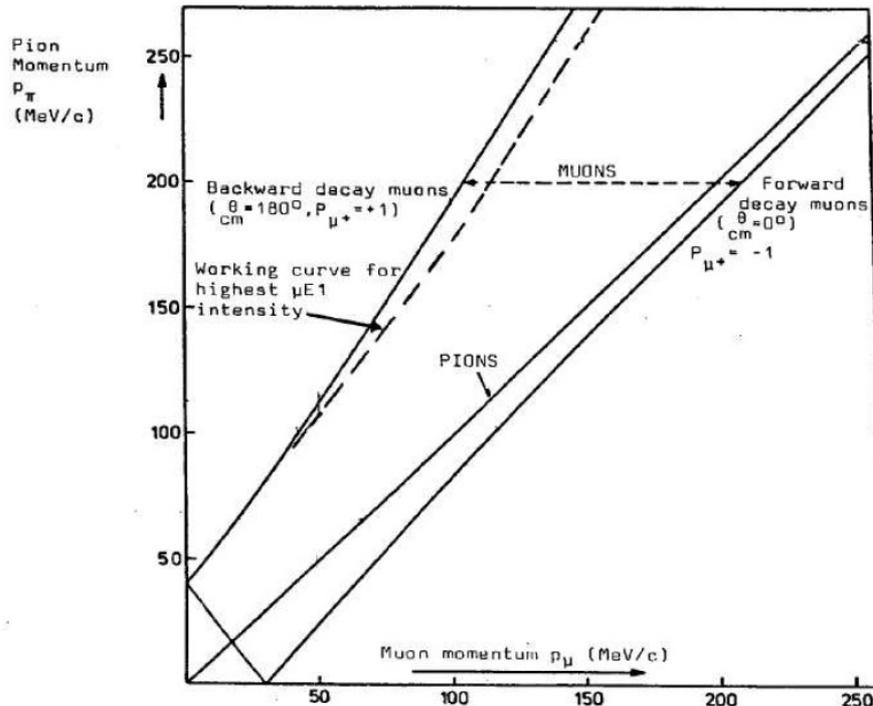


Figure 3.6: The first dipole is tuned to a given pion momentum. The second dipole is tuned to the distinctive momentum of muons which decay in the backwards hemisphere of the pion centre of mass frame. Note that both forward and backward muons travel in the forward direction in the laboratory reference frame. Figure taken from [108]

3.6 Target

The current MICE Target is a hollow cylinder of titanium, with 3 mm outer and 2.3 mm inner radius, which is dipped into the ISIS proton beam. A schematic of the target body is shown in figure 3.7. The target drive motor is of a brushless permanent magnet linear design. Magnets in the target shaft interact with the coils contained within the stator, causing the shaft to actuate. The target shaft has been given a diamond-like carbon (DLC) coating and the bearings are made from VESPEL® material manufactured by DuPont [110]. Previous prototype targets had employed a DLC coating for both, producing significant dust.

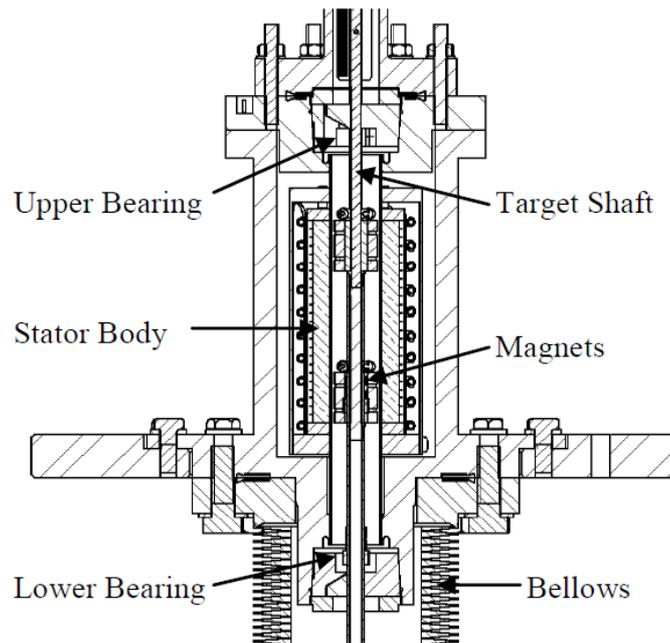


Figure 3.7: A schematic of the MICE target body. Magnets in the shaft interact with the stator coils, causing actuation [109].

A replica exists in a laboratory environment in order to provide assurance that actuations of the currently deployed target do not result in a measurable release of dust into ISIS due to excessive wear on the VESPEL® bearings. The replica target has performed 2.15×10^6 million pulses or actuations without difficulty, corresponding to continuous running for one month. The target positioned in the ISIS vault has so far run for approximately 2×10^5 pulses, through intermittent use. The replica has therefore exceeded the anticipated number of actuations for the year ahead.

Previously MICE had employed a fin shaped target of 10mm length and 1 mm thickness, which although featuring a similar average material, offers the beam a greater average thickness due to the geometry.

3.6.1 Operation of the MICE Target

ISIS accelerates protons from kinetic energies of 70 MeV at injection to 800 MeV at extraction. Each injection-extraction cycle lasts for 10 ms, with the next cycle starting 10 ms later. The beam envelope shrinks to a minimum prior to extraction. It is the responsibility of the

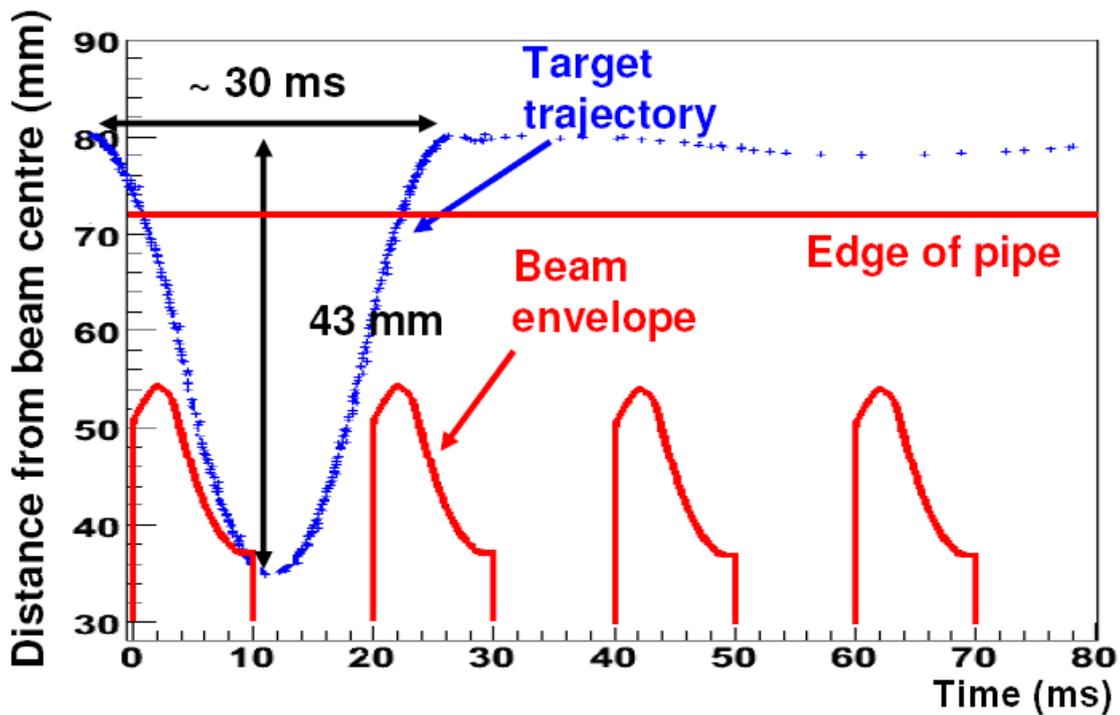


Figure 3.8: The target intersects the beam during the last 2 ms of the ISIS cycle.

target mechanism to overtake the shrinking beam envelope to intersect the beam during the last 2 ms prior to extraction and then raise the target before the next cycle begins (figure 3.8). This places considerable demands on the system, which must provide acceleration of order 90 g.

In a live experiment the target depth and delay are set according to experimental requirements by a designated target expert. These parameters are stored in a configuration database (see chapter 4) with other parameters which characterise the experimental run.

3.7 Absorbers

The effect of the absorbing material in the MICE Cooling Channel is to reduce the total momenta of incident particles. Low density absorbers are required in order to minimise multiple scattering, which inflates transverse momenta. Liquid hydrogen has been chosen as the material in the final three Step VI absorbers, with the baseline vessel 23 litres in volume and 350 mm long. Beryllium safety windows are situated around the absorbers and are

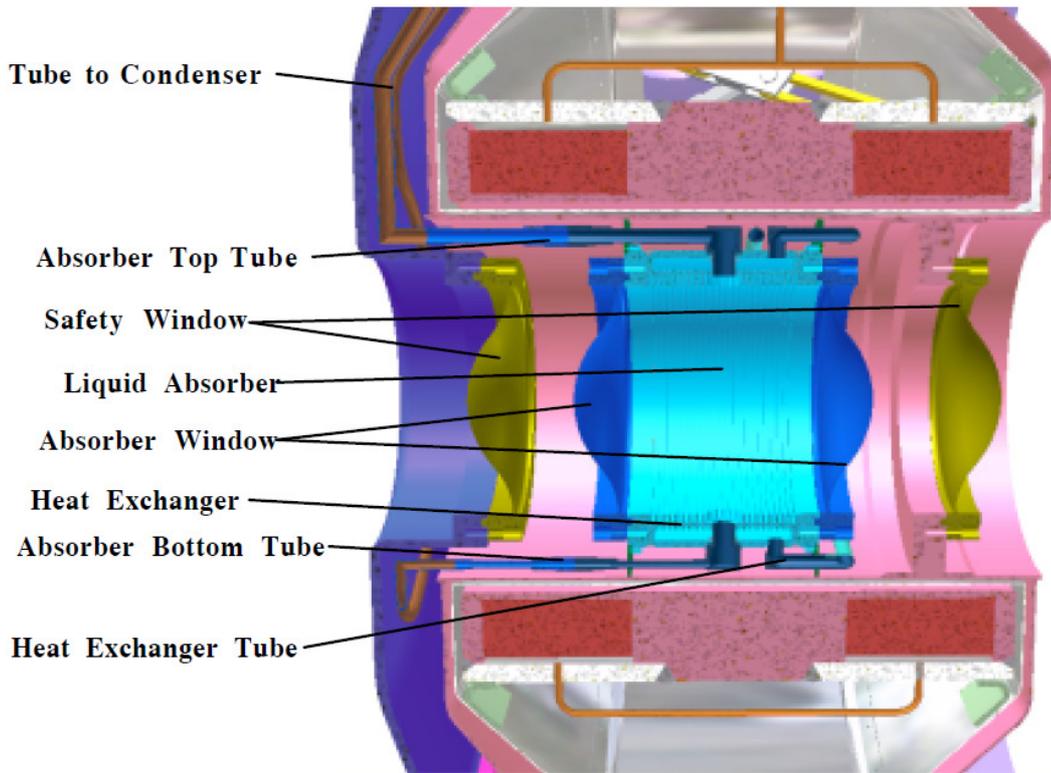


Figure 3.9: MICE liquid absorber cross-section. The Beryllium safety windows are curved to withstand pressure. The absorber windows are made from 6061-T6 aluminium [111].

curved in order to withstand increased pressure. The liquid absorber is shown in figure 3.9.

Each absorber is contained within an Absorber Focus Coil (AFC) module (figure 3.10) manufactured by a private contractor, Tesla. The AFC module is designed in such a way that the absorbers themselves may be readily removed and replaced, indeed there is an option to use a solid absorber of Lithium Hydride in Step III.

Work is currently underway within the collaboration to design and fabricate a wedge absorber which would extend the physics reach of MICE [112]. Whereas conventional MICE absorbers are only able to contribute to transverse cooling of the beam, a wedge shaped absorber would allow for emittance exchange and longitudinal cooling for a dispersive beam (achieved in MICE through a statistical re-weighting of a “real” beam at the offline analysis stage). Emittance exchange rests on higher energy particles passing through more material, resulting in a reduction in longitudinal emittance in addition to a transverse emittance reduction.

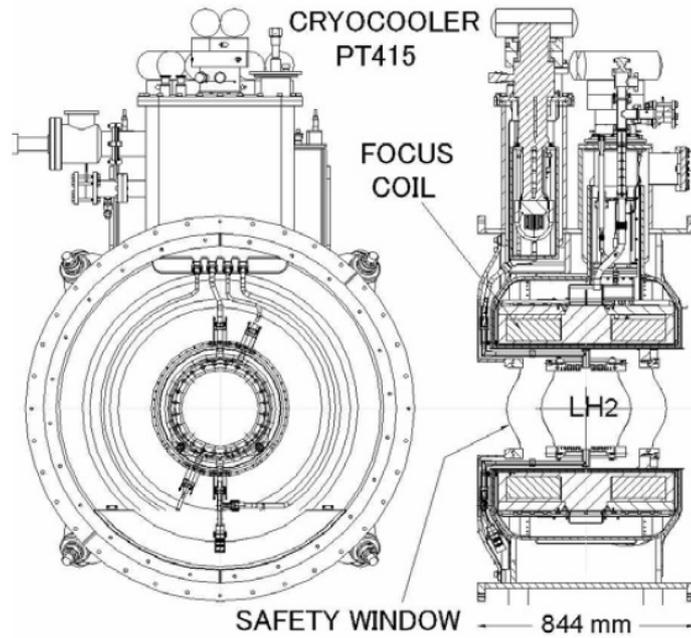


Figure 3.10: The MICE Absorber Focus Coil (AFC) module. [113].

3.8 RF Cavities

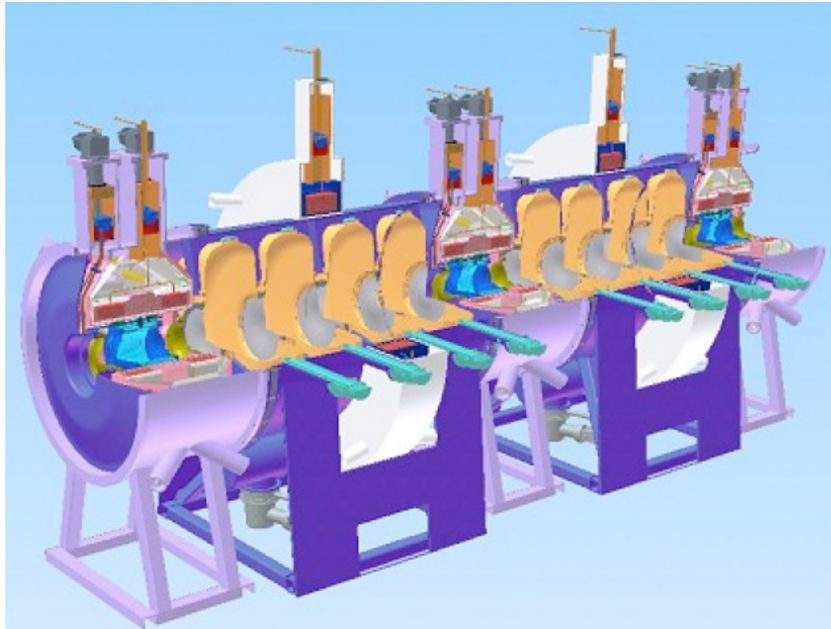


Figure 3.11: Cross section showing two RF cavity modules, with four cavities per module, braced by three AFC modules [114].

Radio Frequency (RF) cavities (Figure 3.11) are a typical means of providing acceleration to particles through microwave signals. An individual prototype cavity is shown in

figure 3.12. The eight normal conducting cavities in MICE will act to restore longitudinal momentum lost in the absorbers. The final Step VI configuration of MICE supports two RF cavity and coupling coil modules, each with four 201 MHz RF cavities. These must provide high accelerating gradients of up to 8 MV/m. It should be noted that this is less than the 16 MV/m required for a Study-II cooling channel [71], due to financial restraints in supplying RF power. The MICE RF system must operate in fields of up to 4 T with minimal production of bremsstrahlung from low energy electrons which can interfere with the detector measurements. The spectrometers and the cavities are only separated by 35 cm of absorber. Nevertheless, the trackers must measure emittance to one part in 10^{-4} . There are significant engineering challenges associated with the use of the RF Cavities in the context of ionisation cooling. The LH_2 , though cryogenically cooled, can suffer a net heating effect from the aforementioned particles produced by dark currents. Careful engineering is required to overcome such challenges for the first time. Power is supplied through a superconducting coaxial loop coupling coil, with one coupler per cavity. Each coupler is coated in TiN to prevent damage from gas discharge or multipactoring.

Ten cavities have been fabricated in total. Five cavities have been completed at Lawrence Berkeley National Laboratory (LBNL). An additional five have been fabricated by Applied Fusion Ltd. [115] under the supervision of LBNL.



Figure 3.12: A single prototype cavity with attached water cooling pipes. [114]

3.9 Detectors

3.9.1 Luminosity Monitor

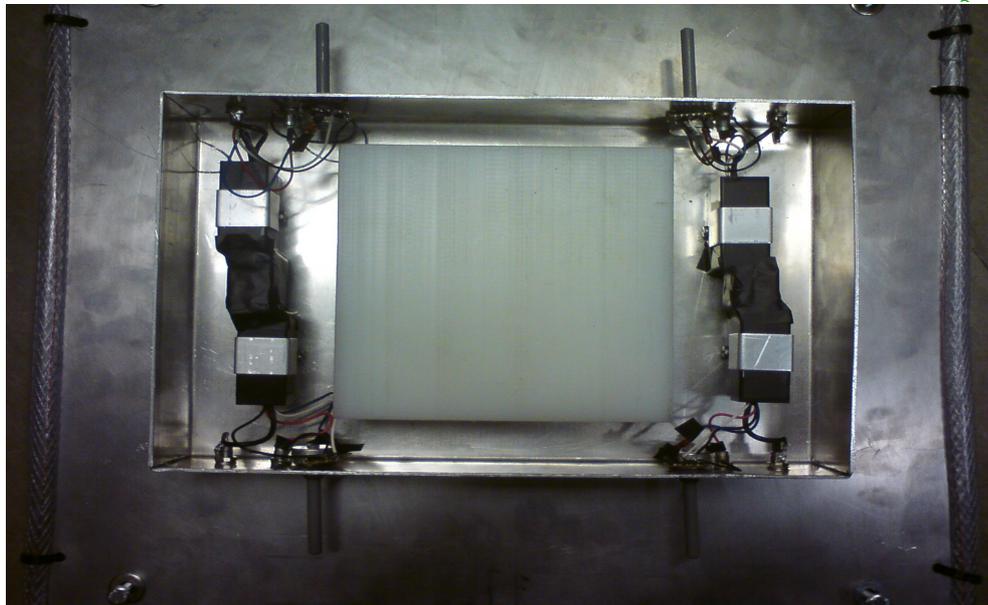
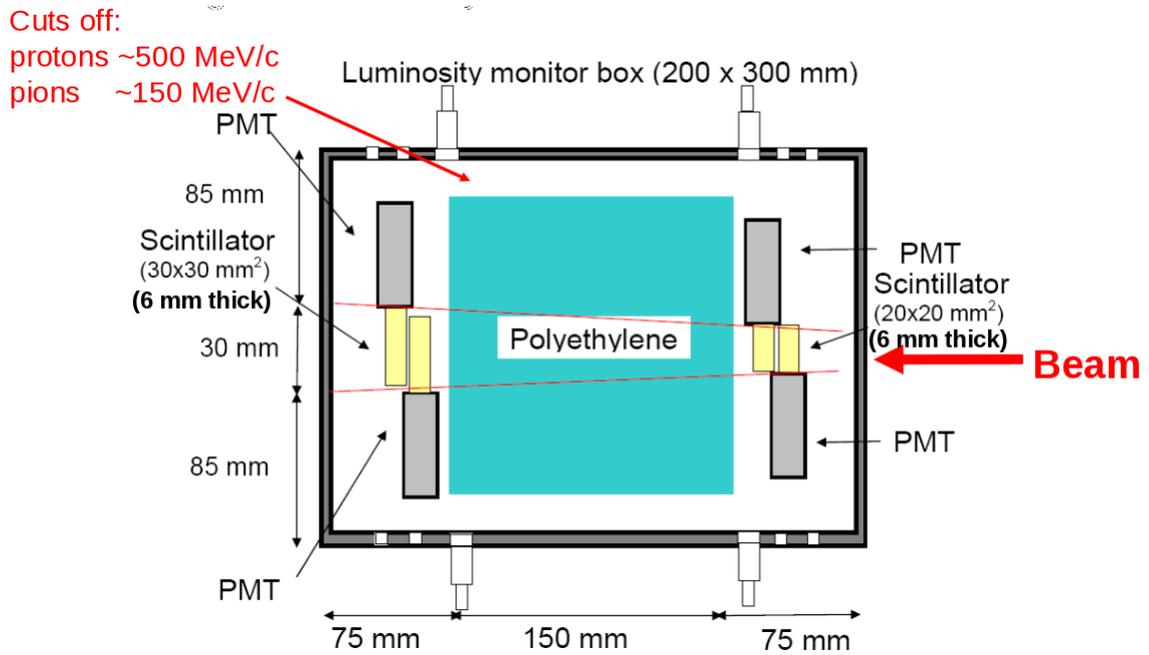


Figure 3.13: The design of the luminosity monitor features four low noise photo multiplier tubes in sets of two. We measure counts from coincidences of each pair of PMTs, and all four together. A block of polyethylene provides a filter for protons below 500 MeV/c and pions below 150 MeV/c.

The Luminosity Monitor consists of two pairs of photomultiplier tubes coupled to scintillators at either end of a slab of polyethylene, suppressing protons below 500 MeV, as detailed

in Figure 3.13. Placed ten metres from the target (distance determined by rate), at the same angle as the cooling channel but within the area enclosed by the synchrotron, it provides a means of relating the ratio of particles in detectors along the MICE Beamline and MICE channel to the number of protons interacting with the target, independent of simulation models. A signal is produced upon the coincidence of either pair of PMTs, or both together on three separate scaler channels.

Further detail on the luminosity monitor is provided in chapter 6.

3.9.2 Particle Counters

There are three detectors used primarily for counting purposes. The Geneva1 (GVA1) detector consists of a single slab of scintillator with $18 \times 18 \text{ cm}^2$ active area and 1 cm thickness read out from a photomultiplier tube. GVA1 is positioned directly after the decay solenoid. Beam Position Monitors 1 and 2 (BPM1 and BPM2) each comprise two planes of Kuraray scintillating fibres. Each plane is read out by a multi-anode PMT. BPM1 has an area of $20 \times 20 \text{ cm}^2$ whereas BPM2 has an area of $45 \times 45 \text{ cm}^2$ as shown in figure 3.14.

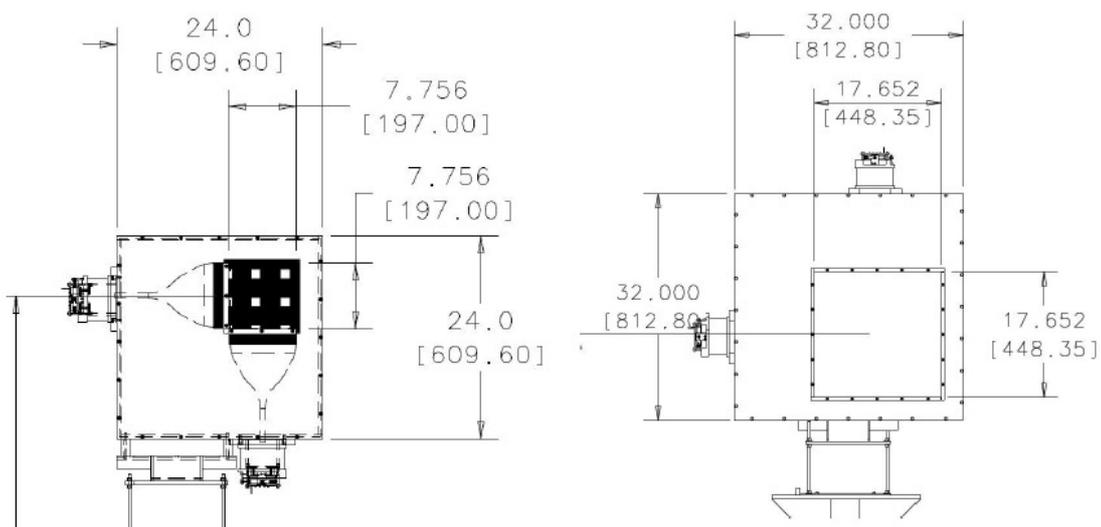


Figure 3.14: Engineering diagram of the BPMs, taken from [116]. Measurements enclosed by square brackets are in units of millimetres, other measurements are in units of inches.

3.9.3 Particle Identification

In order to perform effectively as a single particle experiment MICE requires strong particle identification to differentiate between pions, electrons, protons and muons. This is provided by means of three time of flight counters and an upstream and downstream calorimetry system.

3.9.3.1 Time of Flight Detectors

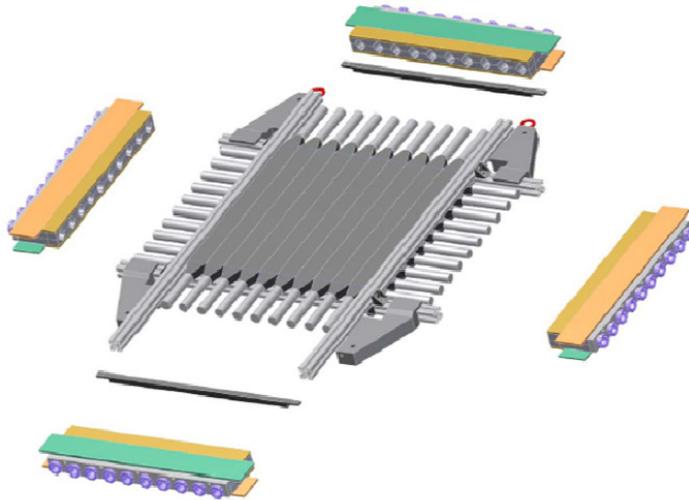


Figure 3.15: The TOF2 detector features a horizontal and vertical plane of 10 scintillator slabs. Each slab has a dual PMT readout. [117]

Three time of flight (TOF) detectors have been installed in the MICE Hall and calibrated using electron and pion beams [118]. TOF0/1/2 each consist of two planes of 10/7/10 scintillator slabs, one arranged horizontally and the other vertically, with each slab attached to two photomultiplier tubes as shown in figure 3.15. TOF0 and TOF2 have a timing resolution of 50 ps, however due to an inferior set of photomultiplier tubes the resolution of TOF1 is presently limited to 60 ps ahead of repairs in the near future. TOF0, TOF1 and TOF2 have an active area of $40 \times 40 \text{ cm}^2$, $42 \times 42 \text{ cm}^2$ and $60 \times 60 \text{ cm}^2$ respectively [119, 120].

To calculate the time of flight between two stations one must quantify a series of corrections: the timing resolution for each bar (σ_t), the transmission time for an individual

PMT (σ_{tt}) measured during calibration and a 'time walk' correction to account for the use of leading edge discriminators which introduce dependence upon the discrimination crossing time [121]. Let t_1 be the time for a particle, travelling at the speed of light, c , to travel between TOF0 and TOF1. t_2 is the time taken for a particle travelling at speed v to travel the same distance. t_1 and t_2 are related to the momentum as follows:

$$p = \frac{m}{\sqrt{\left(\frac{t_2}{t_1}\right)^2 - 1}} \quad (3.25)$$

An illustrative example is given in figure 3.16.

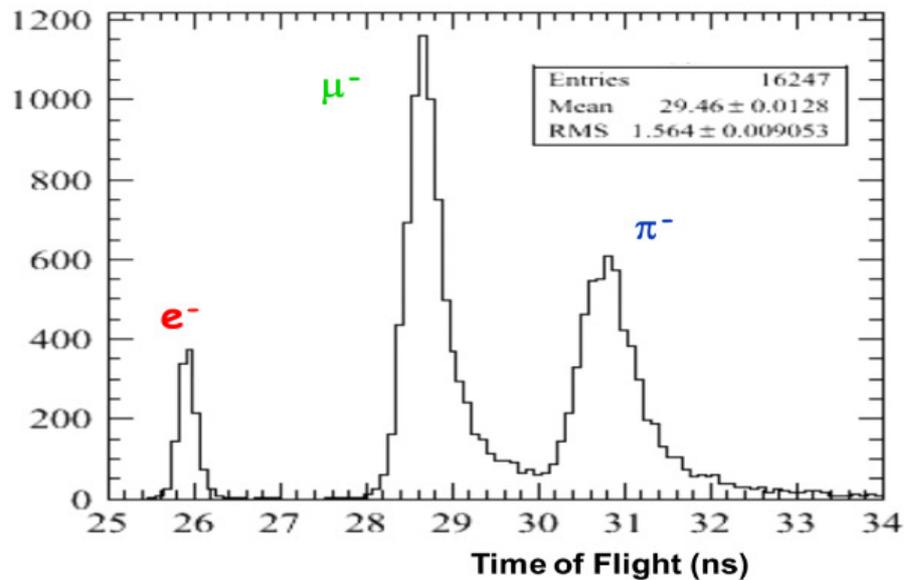


Figure 3.16: The MICE time of flight system is able to differentiate between muons, pions and electrons. Histogram by Rayner [122].

A coincidence of TOF0 and TOF1 is sufficient to make a time of flight measurement, with TOF2 providing coincidence to identify those particles reaching the end of the cooling channel. The more upstream TOF0 is the most frequently used trigger for the data acquisition system during normal data taking in the current Step I configuration.



Figure 3.17: Left: Single CKOV detector, picture taken in direction of beam trajectory. Right: Two CKOV detectors in position between the second quadrupole triplet and TOF0.[123]

3.9.3.2 Cherenkov Detectors (CKOV)

TOF0 and TOF1 alone are not sufficient to provide upstream separation of pion and muon signals for the full range of momenta in the MICE experiment. For this reason, MICE features two aerogel Cherenkov (CKOV) detectors, installed since August 2007 (figure 3.17), which allow greater discrimination between pions and electrons above 240 MeV/c. Cherenkov light is emitted when the velocity of a particle passing through a dielectric exceeds the velocity of light in that medium. One may construct a Cherenkov detector with a dielectric material, or radiator, with a carefully chosen refractive index which satisfies the Cherenkov condition (equation 3.26) preferentially for particles of different mass.

$$n\beta_L \cos\theta = 1 \quad (3.26)$$

No single radiator could be found which would produce Cherenkov light but completely exclude pions over all momenta which led to the design of a dual detector configuration. Two CKOV detectors are positioned sequentially, immediately downstream of TOF0, differing only in the refractive index of the aerogel used as a radiator. Each CKOV is constructed from 15 mm steel plates and features four PMTs at each lateral face. A glass window is placed prior to the Matsushita aerogel block. Conical mirrors guide the Cherenkov light to the PMTs. A diagram is shown in figure 3.18.

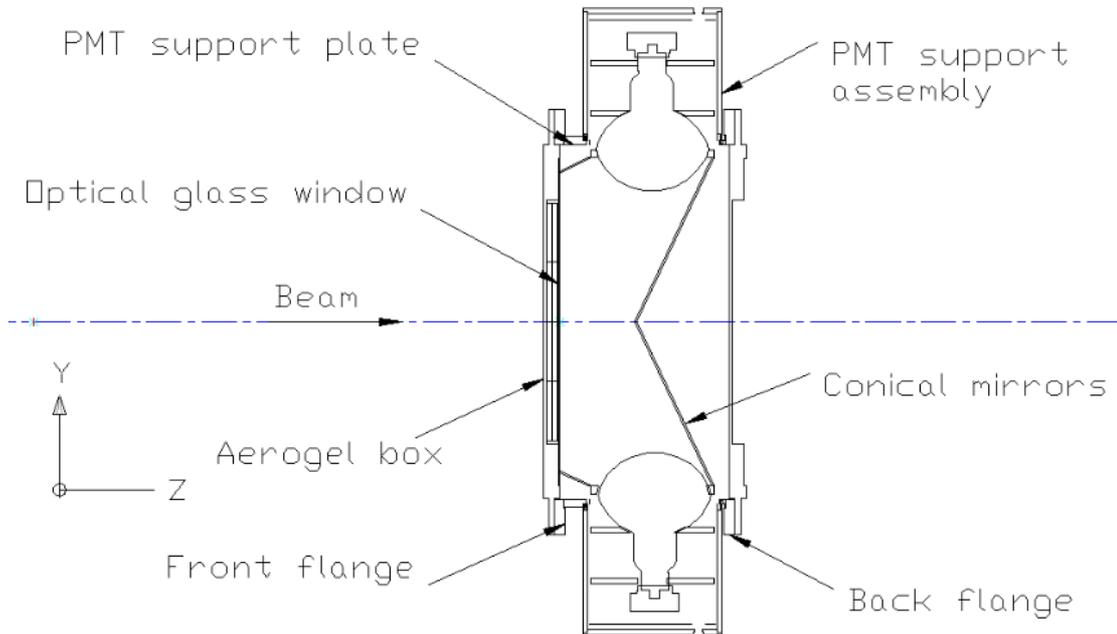


Figure 3.18: Cross-section of a single CKOV unit. An additional two PMTs are positioned horizontally to the beam axis.

CKOV1 has a refractive index of 1.07 and CKOV2 has a refractive index of 1.12. Given the refractive indices chosen, figure 3.19 demonstrates that only CKOV1 will signal muons in momentum region I, whilst in momentum region II a muon will cause a coincidence of both CKOVs. In momentum region III pions will also produce a signal, however this region is of little interest to MICE. For momenta below region I the TOF system is sufficient to separate pions and muons.

Calibration work has been carried out and has quantified a photon yield of 26.4 photo-electrons per electron for CKOV1 and 30.07 for CKOV2 [123]. The two CKOV detectors together with TOF0 and TOF1 provide the upstream particle identification system of MICE.

3.9.3.3 KLOE-Light Calorimeter (KL)

The KL, shown in figures 3.20 and 3.21 closely follows the design of the calorimeter for the KLOE experiment at Frascati National Laboratories (LNF) in Italy [126]. The KL operates downstream of the MICE cooling channel next to the Electron Muon Ranger (EMR) which itself partners with the KL to form the downstream calorimetry system of MICE. Although

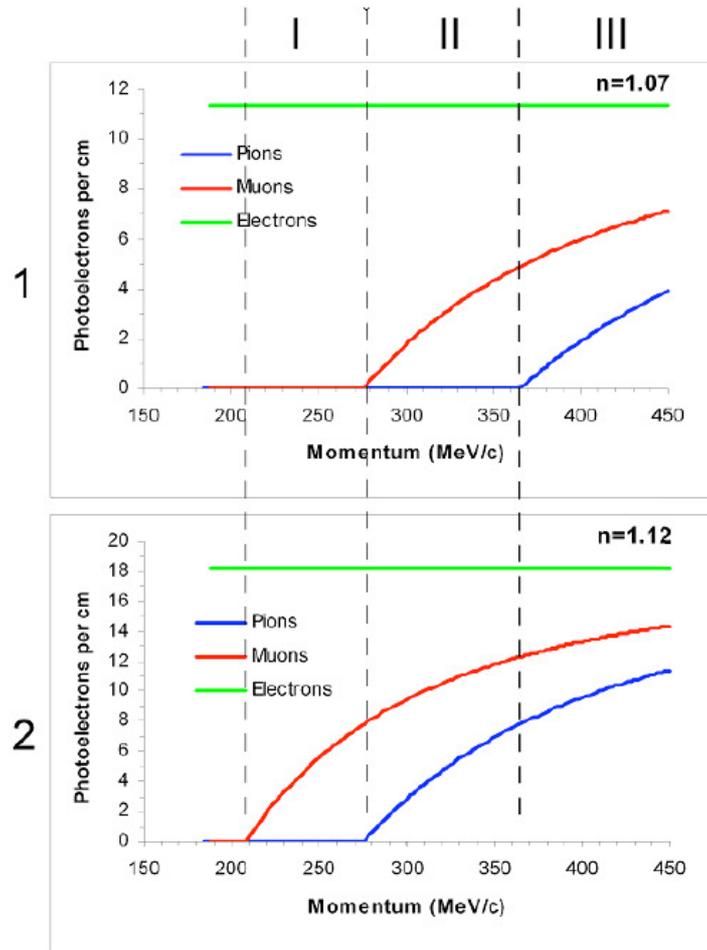


Figure 3.19: The response of the two aerogels to pion, muon and electron events. A muon event in momentum region I will cause a signal in CKOV2 only. A muon event in region II will cause a signal in both CKOV detectors. Region III is not of practical interest in MICE.[124]

few pions should be transmitted downstream of the cooling channel, downstream calorimetry is nevertheless necessary due to the risk of μ decay. Around 1% of muons will decay in the channel, producing electrons which can introduce a significant error to precision muon emittance measurements.

The KL comprises 1 mm diameter BF12 scintillating fibres glued between 0.3mm thick grooved lead plates and is $120 \times 120 \text{ cm}^2$ transverse to the beam axis with 4 cm thickness. The detector has an energy resolution of $\frac{\Delta E}{E} = \frac{7\%}{\sqrt{E(\text{MeV})}} = 7\%$ and timing resolution of $\frac{70 \text{ ps}}{\sqrt{E}}$. It has been operational in the MICE Hall since commissioning in Autumn of 2008.

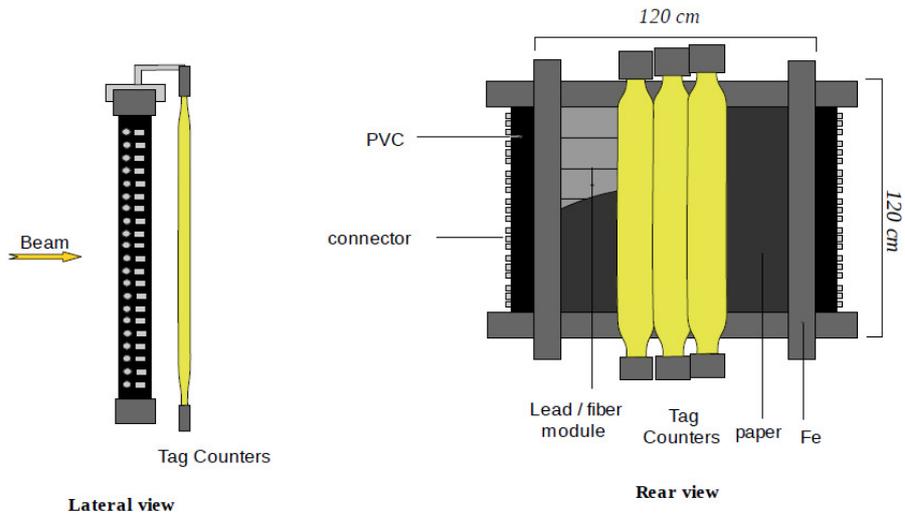


Figure 3.20: Lateral and rear view of the KL detector [125]



Figure 3.21: The KL under test in the MICE Hall at Rutherford Appleton Laboratory.

3.9.3.4 Electron Muon Ranger (EMR)

For cost reasons, it is necessary to discriminate between electrons, which bias our emittance measurement, and muons, functionality which the EMR affords with wide acceptance. The design features 50 planes, each of dimensions $1\text{ m} \times 1\text{ m} \times 1.7\text{ cm}$ with triangular cross-section scintillator bars which are read out through Hamamatsu R7600-00-M64 EG multi-anode photomultiplier tubes as shown in figure 3.22. The EMR will be installed and commissioned in 2011 and positioned immediately downstream of the KL. A study has shown that in addition to providing a means of veto for electrons, the EMR also allows experimenters to infer the muon momentum from the range measurement, to an accuracy of 2.5% [117].

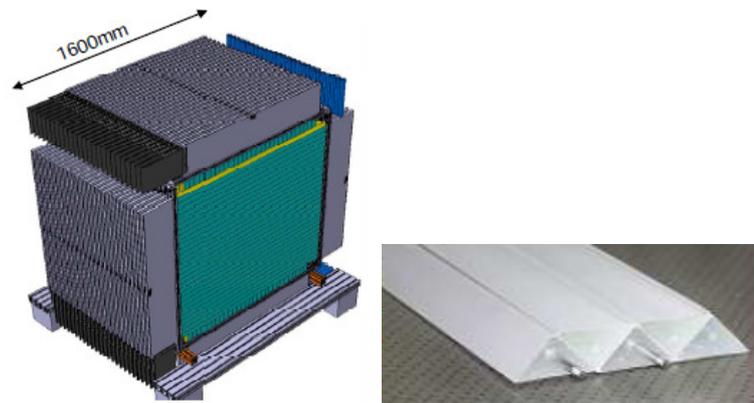


Figure 3.22: Left: Schematic of the EMR calorimeter, featuring an acceptance of 1 m^2 . Right: The triangular segmented scintillator bars are shown. These bars are arranged in 50 planes inside the EMR.

3.9.4 Particle Tracking

The goal of MICE is to measure emittance reduction between 10% and 15%. For our results to be meaningful for a fully realised Neutrino Factory cooling channel, we must be confident in our absolute emittance measurement to an error of 0.1%, which has never been achieved before. A high resolution particle tracking system is therefore of paramount importance in meeting the goals of MICE.

The final step VI configuration of MICE will feature two scintillating fibre trackers, described below, bracing either end of the cooling channel. These trackers provide for space

point and momentum reconstruction, and therefore emittance measurement, before and after ionisation cooling has taken effect, allowing for a calculation of reduction in emittance.

3.9.4.1 Scintillating Fibre Trackers

The MICE trackers [127] are used for position and momentum reconstruction. From these parameters we form a covariance matrix which is the basis of our emittance measurement. Each tracker has an identical design, featuring five stations each with a 40 cm diameter carbon fibre frame supporting three pairs of scintillating fibre planes (figure 3.23). Those fibres are $350\mu\text{m}$ in order to minimise multiple scattering. However this produces a low light yield necessitating a high quantum efficiency from the Visible Light Photon Counters (VLPCs) which are connected to the stations through internal and external light guides. The VLPCs are examples of an impurity band silicon-avalanche photo-detector which convert light into a detectable electrical signal read out through a 1024 channel VLPC cassette system, divided into 8 modules with 128 channels each. They maintain a constant temperature of 9K by means of a cryocooler and were previously deployed for use in the D0 experiment [128].

In 2008 and 2009 studies were performed on the performance of the MICE Trackers using cosmic rays at the Rutherford Appleton Laboratory [129]. Reconstruction was performed using the same software to be used by the experiment for real emittance measurement data. The space point efficiency of each of the five stations was assessed by taking data from the other four to construct a track which crosses the station under examination, residuals were therefore constructed using the nearest valid space point on the tested station. The space point efficiency was $99.7 \pm 0.1\%$ and the track residual rms $661 \pm 2\mu\text{m}$ (see figure 3.24).

Each tracker module is immersed in a 4T uniform solenoidal field [130] in order to produce curved tracks suitable for momentum measurement. Although both trackers have been constructed and tested successfully, there are outstanding issues with the solenoid magnets, precipitating substantial delay in the transition to Step II where one tracker will be used to reconstruct emittance.

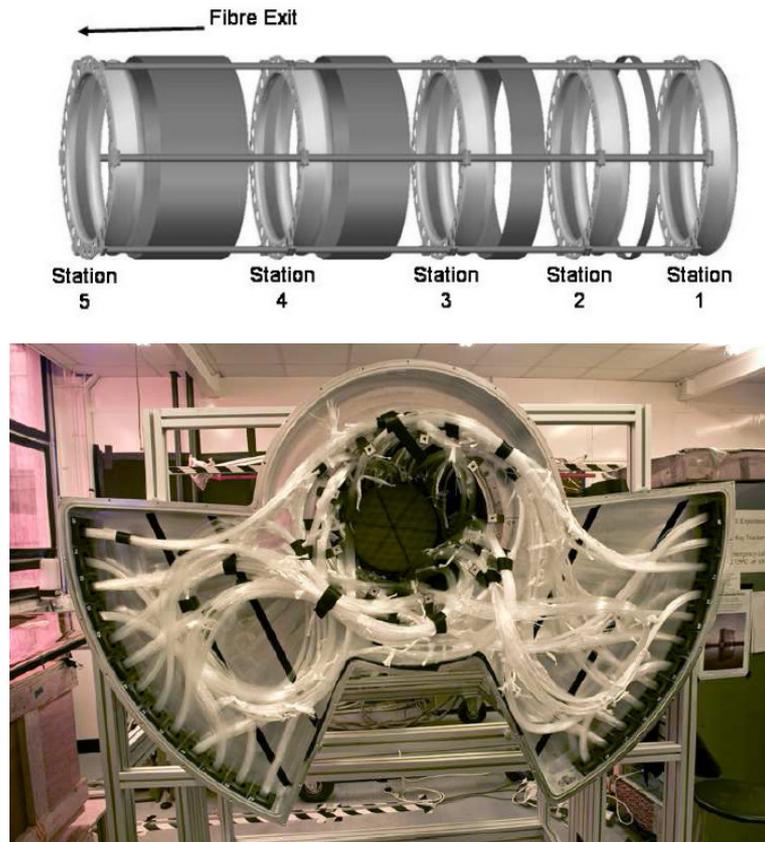


Figure 3.23: Above: Each tracker is comprised of five stations which each support three pairs of scintillating fibre planes. Below: Fibre optic readout is sent through a patch panel to a VLPC system.

3.10 Proton Absorber

It has been noted that a significant proton contamination exists at TOF0 raising the risk of detector saturation and impeding particle identification [131]. A proton absorber has been proposed exploiting the fact that for low momenta, protons lose energy in material faster than pions.

The proton absorber consists of three sheets of polyethylene of 15, 29 and 54 mm thicknesses which may be raised or lowered using a cable pull. Zero or more plastic sheets may be used in concert providing a total thickness appropriate to the estimated proton momenta for a given experimental run. The proton absorber is placed between the decay solenoid and second dipole bending magnet.

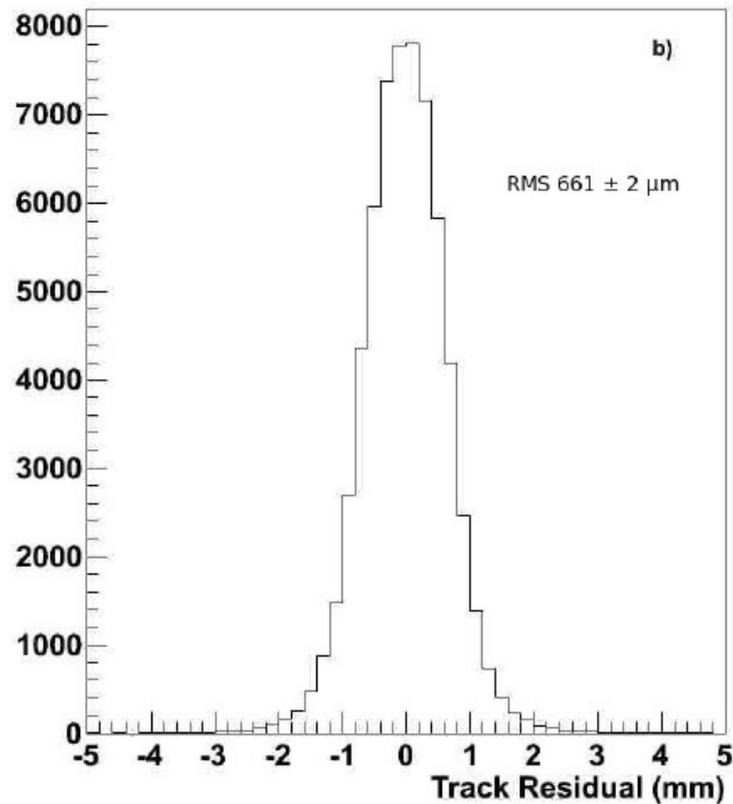


Figure 3.24: Results of tracker performance study using cosmic rays showing the track residual distribution [127].

3.11 Software

3.11.1 Simulation & Analysis

The simulation, design and analysis of the experiment are supported by a variety of software tools developed in-house. Most prominent are G4MICE [132] and G4Beamline [133], described here, which both make use of the GEANT4 toolkit containing models of physical processes useful for particle tracking. As their development has matured, the use of these software tools has overtaken that of applications such as ICOOL and Turtle which were used to a greater extent in earlier phases of the experiment. G4MICE is developed in-house, as is, through the formal inclusion of Muons Inc. [134] within the collaboration, G4Beamline.

3.11.1.1 G4MICE

G4MICE provides an extendable framework for offline simulation, reconstruction and analysis as well as online monitoring and is the standard software in the experiment for all analysis work. G4MICE contains configuration files which describe the state of the experiment with respect to geometry (including detectors, magnets, absorbers and cavities), detector cabling and detector calibrations. Acting as a client to the Configuration database, G4MICE may reconstruct any such aspect of the experimental configuration which was in force during data taking.

In addition to the GEANT4 toolkit for high energy physics simulation, from which the software partially derives its name, G4MICE also exploits the Fermilab BeamTools package for descriptions of magnets, absorbers and cavities. CLHEP provides classes useful for random number generation and other applications including but not limited to mathematical functions (complemented in this regard by GSL) and units. The ROOT analysis framework is used for developing graphical user interfaces and performing data analysis using a wide variety of libraries including classes for drawing histograms, file persistency and graph fitting.

3.11.1.2 G4Beamline

Developed by Muons Inc. [134], G4Beamline provides high statistics simulation of particle behaviour in a variety of fields and materials for some geometry and set of beamline optics. G4Beamline is implemented in C++ although its associated user interfaces are implemented in Java. The user interacts with the application exclusively through sets of data cards describing the parameters and geometry of a run, and is not enabled to write their own C++ code exploiting built in libraries. G4Beamline also includes visualisation features which allow users to view and explore the geometries constructed from their data cards, as well as study the particle species and trajectories of a sample of particles at any point in space. An example visualisation is given in figure 3.25

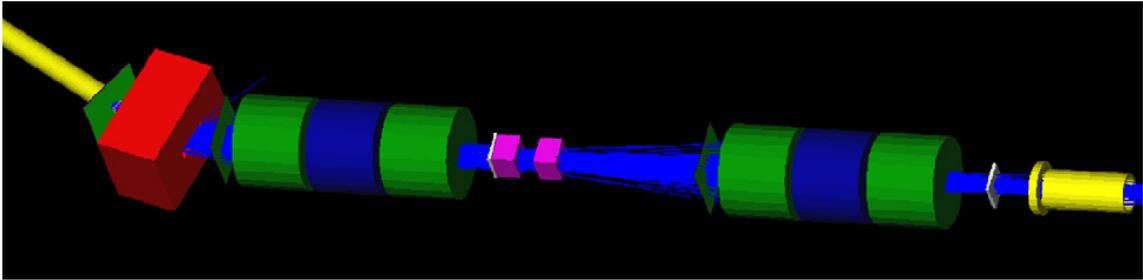


Figure 3.25: A visual representation of part of the MICE beamline geometry in G4Beamline [117]. Here solenoids are shown in yellow, quadrupoles in green and blue and a dipole in red. Smaller boxes and planes denote particle counters and detectors useful for tuning the simulation.

3.12 Outlook

MICE is currently in Step I, and has completed several user runs taking real data useful for analysis of the beamline and a rough measurement of emittance using momentum and position information from the TOFs. Progress to Step II rests primarily upon the successful installation of a working solenoid for the scintillating fibre tracker, with the existing solenoid currently undergoing repair.

MICE will have a legacy of infrastructure in the present MICE Hall which houses the beamline from ISIS which will be appropriate for future cooling studies. In particular, MANX (the Muon collider And Neutrino factory eXperiment) has been designed to study 6D cooling through the use of a helical cooling channel [135]. Discussions have begun to officially classify the MICE hall as a UK Science facility, administered by STFC for further cooling research and a proposal exists to run MANX as a joint experiment between Fermilab in the United States and the Rutherford Appleton Laboratory, with the latter providing the MICE beamline and Fermilab providing magnet and detector upgrades [136].

3.13 Conclusions

The motivations and primary features of MICE have been described, as has its role as a research and development experiment within the context of a future Neutrino Factory facility. MICE is presently in Step I and data taking is currently underway, with a view to moving to Step II in Autumn of 2011.

Chapter 4

The Configuration Database

4.1 Introduction

MICE has been described as a systematic study of ionisation cooling, evolving in a step-wise fashion over time. The positions of detectors and other modules, associated electronic channel mappings and calibrations are therefore subject to incremental change. In addition, beamline settings and alarm handler limits will vary over time depending on the beam conditions selected or as part of an iterative tuning process. In order to be able to enquire about the state of the experimental configuration at any given time, a configuration database (CDB) system has been proposed, designed, implemented and tested.

These configurations exclude both detector data from the Data Acquisition System (DAQ) and fast changing monitoring parameters such as temperatures. The database exists conceptually as an interface between online and offline software; providing information crucial for offline analysis of data and also saving and loading configurations to and from the Experimental Physics and Industrial Control System (EPICS) software [137]. A full diagram of MICE data flow is included in figure 4.1. The configuration database is complementary to a range of other systems and subsystems.

Considerable time was spent analysing use cases of the MICE experimental domain and the results of this are summarised in section 4.3. The design is justified in terms of these use cases in section 4.4. The database exists on a webserver which provides remote access to

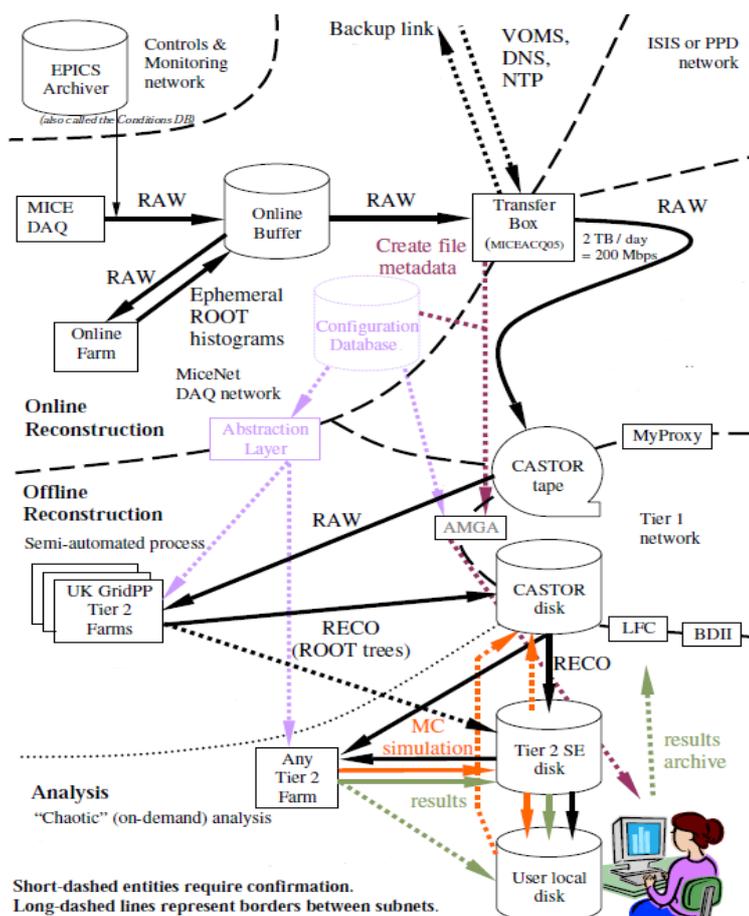


Figure 4.1: An analysis of the data flow of the MICE Experiment[139]. The area we are concerned with here relates to the configuration database and interface only. These are highlighted in purple. This diagram is included to illustrate that the database is integral to data flow in MICE, and that other systems exist to handle non-configuration data.

client applications through the SOAP protocol specification[138] and to users through a web interface. Remote access through SOAP is described in detail in section 4.5.

4.2 Relational Databases

The dominant species of database in public research as well as the corporate marketplace is the relational database. Relational database theory [140] considers database records, or tuples, as lists of (attribute,value) pairs of inhomogeneous type. Sets of these records are contained in structures known as relations. A practical database may contain one or more

relations, each of which have different sets of attributes. Typically tables collect semantically linked data. To give a simple example, a table named 'Student' might contain a matriculation number, name and grade point average, and a table named 'Course' may contain a course number, course title, a number of credits and lecturer. The granularity of tables within the database is typically decided on balance of semantics and efficiency [141].

Each relation should define a primary key: the criterion under which each of its records are considered unique. Best practice compels the database designer to define the minimum set of attributes required to identify a unique entry. In the case of our Student table, the primary key is the matriculation number.

One may relate one or more tuples in one relation to one or more tuples in another relation as mappings between primary keys. For example, our database may also contain a Lecturer relation which has as attributes staff number (the unique identifier and therefore the primary key), name, research group and rank. The course table contains an attribute called lecturer, which would simply be the primary key (or staff number) of the relevant tuple in the lecturer table, in this case the staff number. An attribute which is the primary key of another relation is referred to as a foreign key. The inclusion of the staff number as a foreign key allows one to perform a join operation over the tuples in both our Course and Lecturer relations, which returns rows of course-lecturer pairs where the join condition (here, `Course.Lecturer=Lecturer.StaffNumber`) is satisfied. Including the lecturer staff number within the Course relation is an example of a 'one to many' relationship. We hold that one lecturer may have many courses, but, so far as our example is concerned, a course may have only one lecturer. One to many, or 1-N, relationships are common in databases however effective examples of 1-1 relationships are rare as this would generally suggest that they should be integrated into the same relation, unless one is an optional component of the other.

Our database is modelled in figure 4.2. Here relations are represented as boxes containing their attributes and relationships are indicated by lines between relations. These lines are annotated with a single line when one tuple of the relation participates in the relationship, and branch slightly from a circle when many tuples participate, as can be seen in the one to

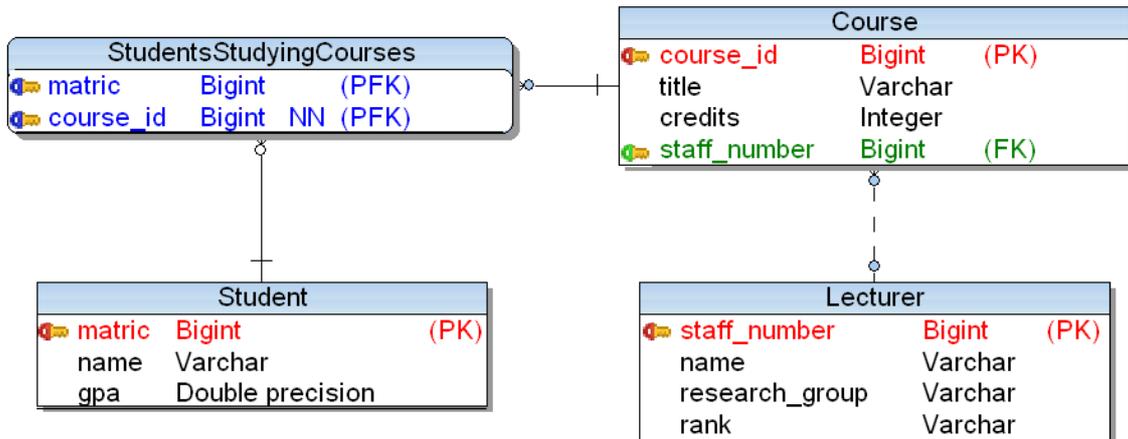


Figure 4.2: An example of a relational database containing four relations: Student, Lecturer, Course and StudentsStudyingCourse. A one to many relationship exists between Course and Lecturer and a many to many relationship exists between Student and Course. Red keys indicate the primary key for that relation. A green key indicates that the attribute is a foreign key from another relation. A blue key, not featured here, would indicate that a foreign key is also part of the primary key for that relation.

many relationship between Course and Lecturer. An example of a 'many to many' or M..N relationship is given between Student and Course, using an intermediate table composed only of primary key pairs from both tables. Expressed simply, the relationship implies that one student may have many courses and one course may have many students.

Modern database management systems [142, 143, 144] are based, with varying degrees of strictness, upon Structured Query Language (SQL)¹ standard. SQL enables definition of database theory concepts and structures and allows for querying operations over their contents. Within the vernacular of SQL, relations are referred to as tables, tuples as rows and attributes as columns.

SQL allows the construction of queries that operate over one or more relations. The user can provide a list of attributes to be returned, the tables from which they should be found, the condition(s) which must be met and joins which should be performed. An experienced database user can interact with a database solely through the use of SQL for any routine operation including but not limited to table definition, querying, inserting data and updating

¹It should be noted that strictly SQL is not an acronym and therefore should not be expanded, nevertheless it has become common to offer the expansion of 'Structured Query Language' when introducing the basic concepts.

existing data.

4.3 Analysis and Use Cases

The CDB has been developed from scratch around a careful study of the collaboration's requirements over several disparate domains. Care has been taken to ensure that the information in the database is necessary and complete. That is, a run configuration is uniquely defined by its records in the CDB, and that information not relevant to the configuration is not stored, avoiding duplication of information stored on other systems.

At all times, the concept of a configuration database must be kept separate from that of a conditions database. In the CDB for example, the magnet currents requested by the users running the experiment are stored. These magnet currents are the same as have been entered in the control system. Together they represent part of the set configuration of the beamline which results in some data output. However, the actual measured currents may differ from those requested. The CDB does not store the actual measured current but simply the value requested by the shifter. The measured values of the magnet voltage and current are stored on the EPICS data archive. It should be noted however, that the CDB also stores the alarm handler limits which constrain the measured current to some range centred on the set current. EPICS uses these same limits to warn the shifter if the measured values drift from the set values by an agreed amount.

As a guideline, all information in the CDB should be useful for offline analysis. Much of the analysis of the use cases was focused on establishing what information met this criterion and designing the database schema accordingly. For this approach to succeed it was necessary to gather an understanding of the various problem domains in the experiment, cataloguing and collecting use cases from the analysis group, the beamline and target groups as well as the software, online and detector groups. Stakeholder involvement was critical from the outset.

The various areas generating configuration data of the experiment were identified, and are referred to in this chapter as the domains of the CDB. These are: Geometry, Calibration,

Cabling, Set Values, Tagged Set Values and Alarm Handler Limits. These are discussed in more detail in their respective sub-sections. However, in examining these domains more closely, certain common features of configuration data emerge which ease system design and further justify the decision to restrict the remit of the CDB to configuration data only.

In particular, it is clear that information will be written to the configuration database at no higher frequency than the end of each run, which in the maximal case of a succession of failed runs could reach once per minute. Such information will always originate from the MICE local control room and no requirement exists for remote write access. Read accesses through remote and web based clients should not exceed a few per second but could originate from a significant number of existing or future systems, in any geographical location. In several domains such as geometry, calibration and cabling the ability to update information but retain previous values is desirable. These requirements, combined with those discussed in the remainder of this section characterise a secure ‘bi-temporal’ database with web based remote access through HTTP, the features of which will be explained in the following sections.

4.3.1 Beamline Settings

Without the database one must manually read, enter and record beamline parameters such as magnet settings in the MICE control room. A key use case requires the database to allow users to retrospectively determine what settings were in force at a given time or in a particular run of interest and reinstate those settings. Users must also be able to aggregate settings and associate them with a given tag, which may or may not be associated with a real run. The settings contained within this tagged run may also be loaded in the control room prior to a run, and then recorded in the database as the real settings which were in force.

There is therefore a need for two types of beamline setting information. Clearly we must be able to query the database for the settings information for a particular run or timestamp. However, some configurations may be reused for several runs. It is possible that one may wish to reuse the configuration previously used on a particular run, for a particular momentum etc. We refer to our two types of set value information as run and reference, or tagged,

settings. Run settings are timestamped and associated with a run number and can be readily saved and retrieved. Reference settings are indexed by a human-readable tag and are not necessarily associated with a timestamp but may be used in 0 or more runs. It is envisioned that certain beamline settings will be marked for re-use which may or may not have previously been instantiated as run settings. Before the start of a run, the EPICs user can load reference settings for use in a real run.

4.3.2 Geometry

Geometry is defined here as position, dimension and material information for a module. A module is placed into a context, potentially exposed to a magnetic field and/or contained by a parent volume. Detectors, magnets, absorbers and RF-cavities are a few examples of complete modules in the MICE domain.

The geometry of the experiment is expected to evolve in a stepwise fashion. New modules will be added and subtracted from the MICE hall as part of regular preparations for normal running. It is necessary to be able to establish a full understanding of what the geometry of the experiment was at a given time.

In addition, it is expected that post-run surveys of the hall will reveal misalignments of modules. These may occur, for example, due to the presence of magnetic fields, improvements in the accuracy of survey methods or human error. The database must store the corrected geometry information in addition to the previous incomplete information. Therefore, in addition to storing a history of state, the database must also support a history of what was known about the state.

4.3.3 Calibrations

MICE contains many elements requiring calibrations including a diverse community of detectors and detector elements (such as photon detectors) contained within those detectors. In order to generalise the concept of a calibration and minimise unnecessary design and implementation overhead, we require that a calibration may have an arbitrary number of pa-

rameters of any type. It is understood that calibrations will pertain to definite time periods, and that more than one calibration may exist for a particular module at a particular time.

It is the responsibility of the detector groups to ensure that the calibrations in the database are sufficient and correct and that the default calibration returned represents an approved and agreed calibration.

4.3.4 Electronic Channel Mapping

Each detector necessarily combines at least one distinct cabling configuration associated to a variety of electronic modules (such as preamplifier pulse discriminators, scalers, time-to-digital (TDC) and analogue-to-digital converters (ADC) and other components) subject to variation with time. Changes in cabling for some detectors necessarily mandate a wholly new calibration, for example in the case of the TOFs due to their high timing resolution. Nevertheless, it is desirable to accommodate all cabling domains in a generic way in order to ease maintainability of the code and also to accommodate future components.

Analysis users require to reconstruct from the database the cabling settings in force for a given run and detector or time and detector combination. Additional use cases may exist for future cataloguing of components stored and deployed, providing a complete description of cabling, however these are not presently fully understood in terms of selection criterion or nature of client software.

4.3.5 Alarm Handler Limits

During running there exists sets of major and minor alarms for values of a list of critical parameters, such as magnet currents which are labelled process variables in the EPICS nomenclature. The alarm handler software, running in the control room, sounds the appropriate audible alarm if minor or major limits are exceeded.

There is an understood distinction between values set by a control room operator for parameters such as magnet currents and values measured from the actual magnets themselves. Some fluctuation may occur from the value originally set. Assuming the difference between

set and measured values is small, no action is taken. Should some parameter exceed a major limit, the run itself will become invalid. In all cases an audible alarm is heard in the control room.

The alarm handler software itself is independent of the CDB. However, the alarm limit information as well as being of immediate use in the control room will also be preserved for offline users and as a possible reference for future runs.

4.4 Design and Implementation

4.4.1 Bi Temporal Paradigm

Key use cases require that facts in the database be associated with a period of validity. For example, a typical requirement describes a user asking the API (Application Programming Interface) for the state of the experiment at some point in time. Therefore when we place a component such as a detector in the MICE Hall and later remove it, we can expect our period of validity to match the duration over which the component was present. Calibrations also are expected to have well defined periods of validity.

However, it is not enough to be able to reconstruct the state of the experiment at a given time. We must also be able to reconstruct what was known about this same state at another time. For example, in the case of geometry, one may later discover a component was misaligned. They must record in the database the new geometry information for this component. However, we wish to keep the old information for analysis, with its overlapping period of validity. In the case of calibrations, one may have a calibration valid for that time, but later discover another perhaps better calibration also valid for the same time.

Clearly we must retain both calibrations, and either calibration may be required to answer question about the history of what was known about the state of the experiment at a given time. It follows that as well as a valid time, or period of validity, entries in the database must also have a transaction time, which is the time they were added to the database. This allows for a history of the known state of the experiment to be reconstructed, and assures us that we

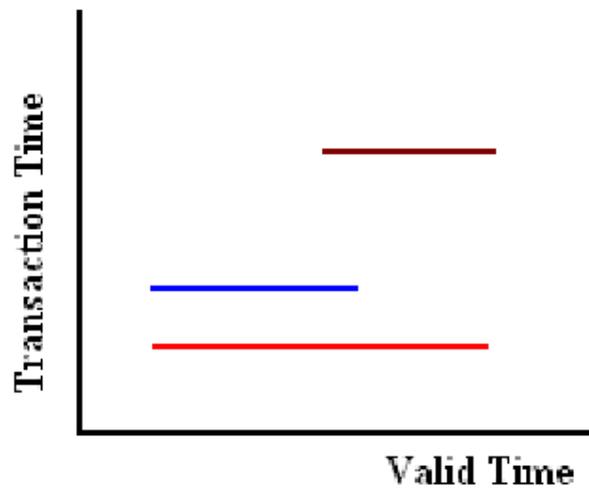


Figure 4.3: Here separate database records are represented by different coloured lines which have the same valid time period. Each record is added to the database at a different time, resulting in a different transaction time. These two time dimensions are orthogonal to one another. In the case of geometry information, valid times for different records representing the same module perfectly overlap in valid time, whereas calibrations may partially intersect one another in valid time.

are not obliged to delete any information. These two dimensions of time must be included as separate timestamps in the geometry, calibration and cabling areas of the database.

The database therefore can be described as bi-temporal in paradigm, storing information with respect to both valid time and transaction time, as shown in figure 4.3. Bi-temporal database theory originated with the work of Richard Snodgrass [145] who proposed modifications to SQL-92 [146], a standard for Structured Query Language (SQL) which is used for managing data in database systems, to support dimensions of time. Traditional relational databases rely on a unique attribute or set of attributes, called the primary key, which each record uses to distinguish itself from all others. SQL queries become more complicated when one accepts that it is necessary to address a record by one or more dimensions of time. Joins, common database operations which combine two or more sets of records (called tables) by means of key comparisons, become more resource intensive when made between relations with one or more temporal attributes in their primary key, as the nature of the temporal domain requires that the join be made upon inequality predicates rather than the equality predicates which suffice for non-temporal joins [147].

4.4.2 Beamline Settings

In the design and implementation we make a distinction between values set by shifters for an actual run and sets of values saved as a tagged run for future use. Figure 4.4 illustrates a schema for supporting real runs associated with a start and end time. The run table contains meta data describing the run itself, the nominal particle momentum and primary beam particle species are included here. The magnet table completely describes the optics of the experiment, storing current information for any magnet with a *magnetName* value which is always one of *ds*, *q1..q9*, *d1*, *d2* for the decay solenoid, quadrupoles and two dipoles respectively. At present, magnet settings are read into the database in a completely automated fashion as should continue to be the case when the additional cooling channel magnets (tracker solenoids and absorber focus coils) become available. However, some elements of the run table must be entered into an electronic form by an operator. It is envisaged that the ISIS, DAQ and Target tables will also be largely automated in the near future.

Client applications may interrogate the CDB system for set values in force for a given run or at a given time. In the case of tagged runs (Figure 4.5) we store a reduced set of values with respect to the case of a real run. From figure 4.4 to figure 4.5, only the Run and Magnet tables survive in any form as the others contain information (for example, relating to synchrotron conditions in case of the ISIS table) which cannot be set from the MICE control room, and must instead be recorded on a run-to-run basis. Due to safety reasons, the Target clearly must remain in human expert control and not relinquished to automated control. Furthermore DAQ settings do not describe any aspect of configuration which we should necessarily impose on future runs as DAQ versions will evolve with time, and gate width and trigger should be at the discretion of the experimental shift leader.

Client applications may request tagged settings by referencing their unique tag. In the case of both tagged and real settings, write access is granted exclusively to control room processes. Web interfaces, detailed in section 4.5.3 exist for both tagged and real set values.

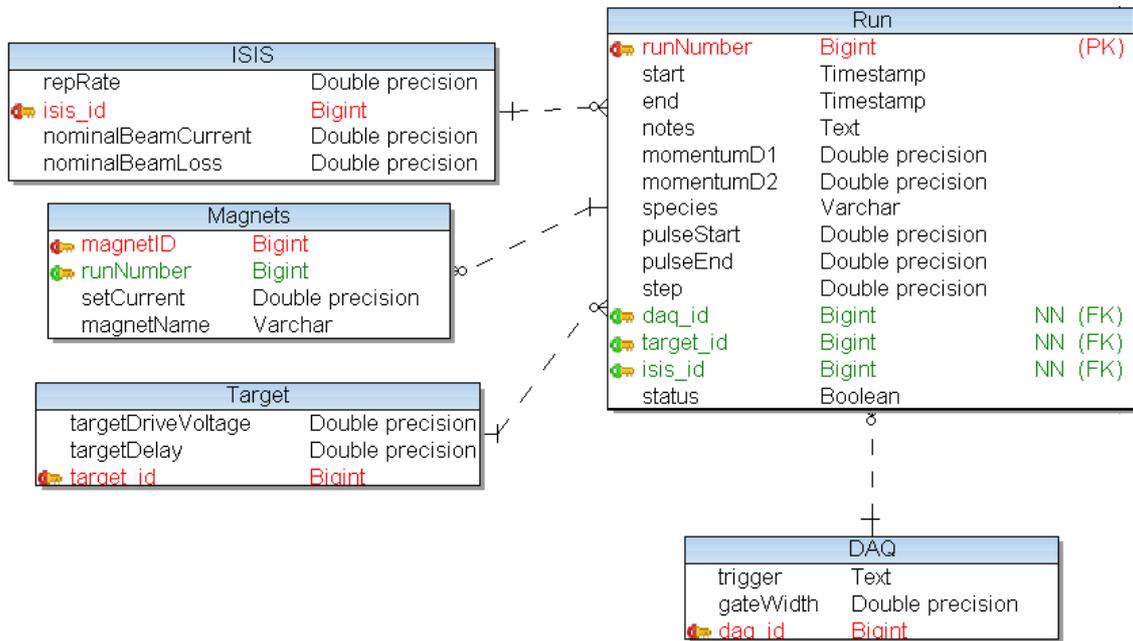


Figure 4.4: Run meta data is tied to set values for target, magnet, ISIS and DAQ settings. Each run has one row of the Target, ISIS and DAQ tables associated with it and N rows of the magnet table where N is equal to the number of magnet modules deployed in MICE, nominally 12 (9 quadrupoles, 2 dipoles, 1 decay solenoid from the beamline) in step 1, increasing when magnets are added to the cooling channel.

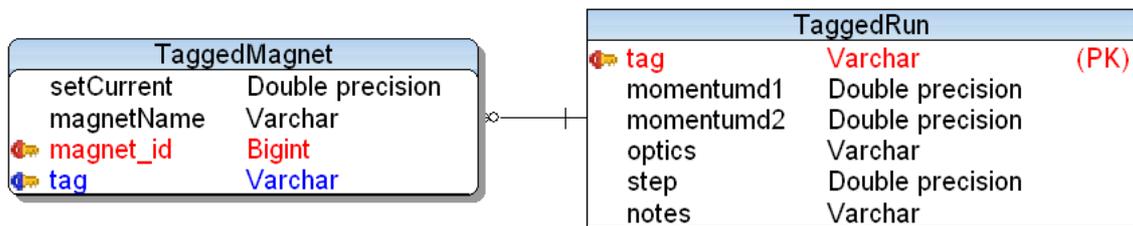


Figure 4.5: The settings associated with a tagged run are a reduced set of those available for a real run. We do not wish to allow the user to automatically apply all the settings which are normally recorded with the set values associated with a real run. For example, the target is a critical piece of MICE equipment on which experimental running depends and target settings are the key driver of radioactive activity in the hall. Therefore target settings are an example of values which must remain under human expert rather than automated control.

4.4.3 Geometry

A simple one table design was employed for modelling the geometry of the experiment as detailed in figure 4.6. This was chosen prior to the formal adoption of the API detailed in section 4.5.1 and the abstraction this afforded, in order to ensure that the database structure mirrored that of the client code G4MICE. In particular it was essential that the database not impose any constraints on G4MICE.

The database stores the positional, rotational and dimensional information of all modules that G4MICE uses for analysis, as well as other module specific information like field maps, material composition etc. The information for each module is given a valid time period, for example, the time that a module is in place. Should that module be found to be mis-aligned, the database is given a corrected record for the same valid time and the old record is retained. Both records can be easily discriminated and the correct one chosen however, through another field, the transaction time, which is the time the information was added to the database.

Geometries can be written to the database from G4MICE. This code exists but was not included within G4MICE during database testing. The geometry files in G4MICE are easily edited and are in a format that can be easily understood by a human being without programming expertise. In order to write the information to the database, G4MICE loads the geometry and translates it into XML [148], to be described in section 4.5.2. This XML file is sent to the database API which parses it and writes it to the database.

G4MICE will be used to reconstruct geometries from the database. In order to do that, it parses an XML response from the API. It is an obvious unit test to use G4MICE to write a geometry to the database, recall it, and test for differences. This has been tested and shown to work.

4.4.4 Calibrations

In keeping with our temporal model described in section 4.4.1, we store both the period of validity and the time that the calibration was added to or edited within the database. The

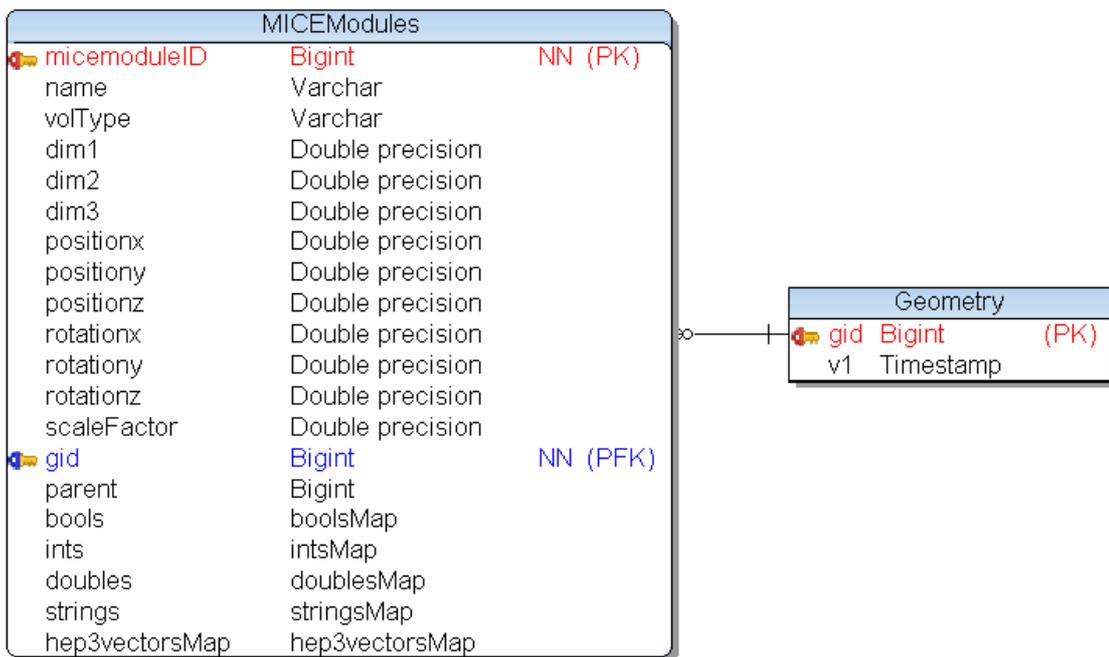


Figure 4.6: The geometry of the experiment is modelled as a hierarchy of modules with parameters described above. Each module has knowledge of its parent module through the link to its unique identifier. A parent id of 0 defines the mother volume. The module structure was developed to mimic the geometry structure in G4MICE rather than develop a new database specific model.

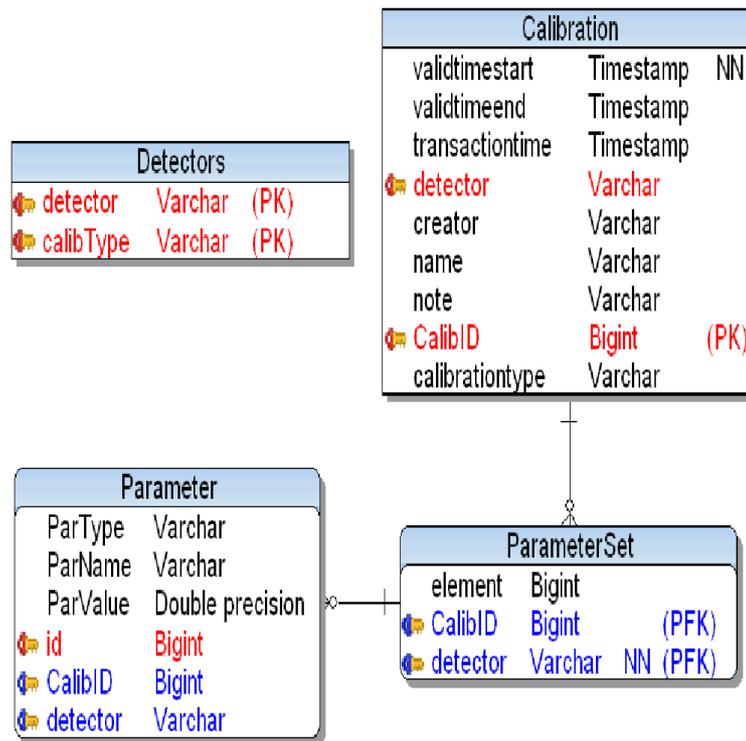


Figure 4.7: Support is given for calibrations comprising arbitrary length and potentially inhomogeneous lists of parameters. The relationship of parameters with parameter sets follows a "many to one" correspondence, and the description of the type of parameter is provided by the paramType attribute.

default calibration is presently defined as the most recently added calibration for a particular timestamp. However, specific calibrations may be requested through data cards in G4MICE.

The table structure (Figure 4.7) allows for flexible definitions of calibrations. One calibration of a given type for a given detector, may have one or more parameter sets. These parameter sets are inhomogeneous vectors of arbitrary length. For ease of querying, values of the parameters are stored as double precision numbers, although the parType attribute allows the user to specify any numerically represented type, eg double, boolean, binary or integer. Allowance is made for a detector to have more than one type of calibration (eg a calibration for each PMT). Calibrations may only be of a (type, detector) combination found in the Detectors table.

4.4.5 Electronic Channel Mapping

It is critical that the cabling schema adequately describe a generic detector rather than presuppose that time and expertise will exist to design new schemas during detector commissioning. With limited manpower resources for maintenance it is unlikely that each detector group would be able to nominate a database expert to maintain the representation of that group's unique detector schema.

Electronic channel mapping is a necessarily complicated domain to map directly to database tables. However, the requirements include only a very limited set of selection criterion which is of an identical format over all detectors. This criterion includes valid time information or run number, cabling type (under the assumption that one detector may have one or more discrete sets of cabling information) and detector name. The remaining information, that is the actual channel mapping themselves, may be stored in flat files. Flat files contain no enforced structure or relationships. Therefore, their structure is then free to change as a function of detector. An unimplemented design where each detector was treated individually is shown in figure 4.8. This example goes beyond the electronic channel mapping use cases which might involve cabling files being automatically read by G4MICE and includes information that would be useful for human reference only through an unknown client. Those latter use cases have been excluded in the current design, enabling the flat file approach in figure 4.9 which fits best when G4MICE is the only probable client, as the structure of those flat files may be identical to those used currently in G4MICE, reducing necessary design and implementation overhead at the client side. Had this approach not been taken, it would have been necessary for a control room client to convert a G4MICE file into a message the database can understand, the CDS to translate this into its tables, the CDS then to convert it back into message form and send it to a different version of G4MICE under the control of an analysis user, and for that user's version of G4MICE to translate the message into its own cabling file format. Taking the flat file approach, the control room version of G4MICE simply sends the file itself, with a few tags describing meta data such as valid time and cabling type, and the CDS returns this identical file to the G4MICE client. Reducing the number of repre-

sentations of channel mapping within the experiment's software framework assists channel mapping experts in ensuring that the information is complete at all stages.

4.4.6 Alarm Handler Limits

The schema supporting alarm handler requirements is illustrated in figure 4.10. Minor and major alarms are represented by four variables - lolo, lo, hi, hihi - which constrain measured values fed back to EPICS within operator set limits. Each set of variables is tied to a single process variable (PV) name and through the AlarmHandler table to a timestamp also. For each timestamp the PV name must be unique.

Should a major alarm begin during a run, it can be assumed by default that conditions during the run are not stable and the run may be flagged as unsuitable for analysis. The user is able to query the database for the alarm handler limits in force at a given time.

4.4.7 Implementation

The database management software is Postgres 8.3 [142], which is a sophisticated open source solution widely used in high performance applications in the commercial and scientific sector and contains temporal extensions to its implementation of SQL, however these are not presently exploited in the CDB system. The API is written in Java, although non-Java applications can readily interface with it, for example the control room EPICS software. It communicates to the database using the JDBC libraries which are independent of database platform. The input and output to the API is given in an XML file format in all cases. XML is composed of lists of elements and attributes in a manner analogous to other markup languages such as HTML. However, we may define and redefine the elements and attributes which each XML message can make use of. As such, XML is extendable; new elements may be added to its structure reflecting changes in the database structure, which is desirable for future proofing of client applications in the control room. Further information about XML is given in reference [148].

The database interface is situated on a webserver for remote access, described in section

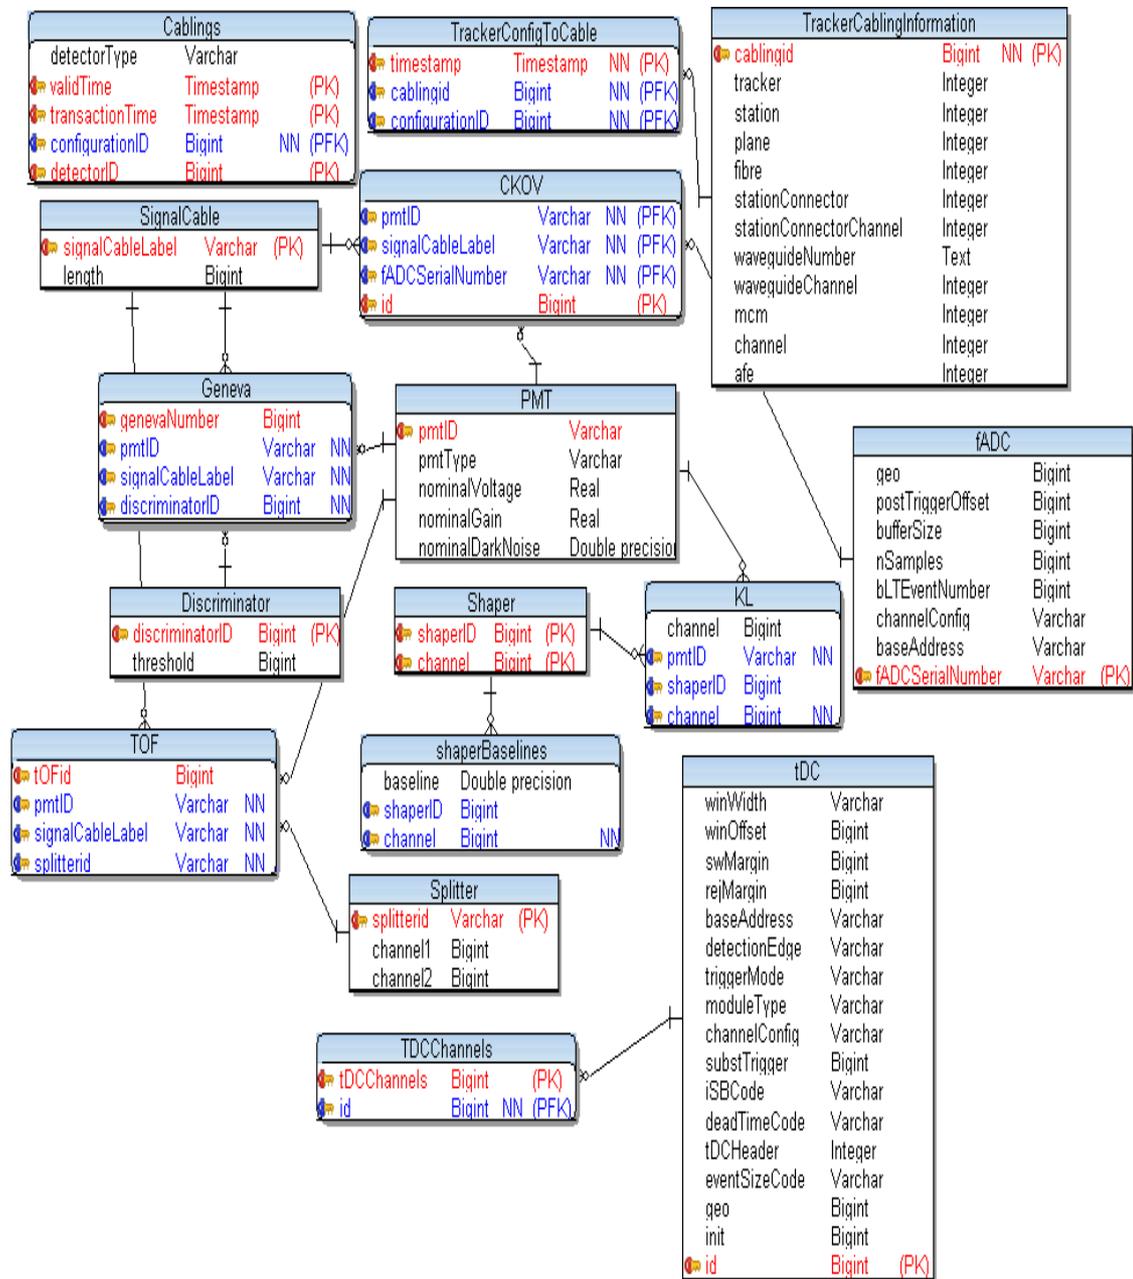


Figure 4.8: An approach taken without the benefit of the flat file channel mapping strategy. Detectors are considered roots of cabling. Here the KL, TOF, Geneva and Tracker cabling are illustrated. The diagram quickly becomes more and more complicated and difficult to query and maintain with more resolution and greater number of detectors.

Cabling		
 id	Bigint	(PK)
startTime	Timestamp	
name	Varchar	
type	Varchar	
detector	Varchar	
file	Bytea	

Figure 4.9: The flat file strategy. Here we have united the representation of the electronic channel mapping for all detectors, increasing maintainability and reducing complexity. The flat file is stored as a byte array and may be particular to some detector without introducing constraints on the rest of the table which remains sufficiently general as to be useful for all detectors.

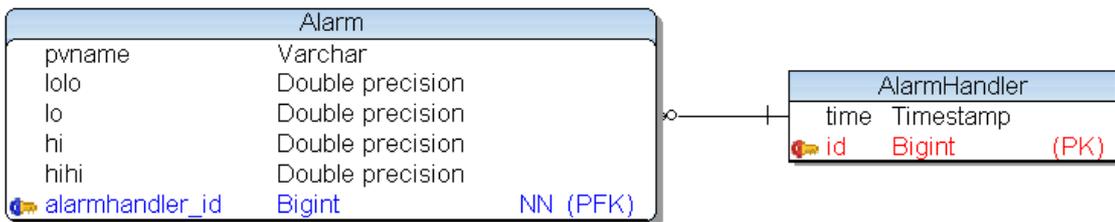


Figure 4.10: Every alarm is tied to a process variable with a unique name, 'pvname' in the Alarm table shown. An alarm structure includes low and high limits and very low and very high limits indicating the severity of the alarm. The AlarmHandler table allows us to recall the set of alarms which were in force at a given time.

4.5 which runs Apache Tomcat 6 [149] and uses SOAP 2.2 [138] (described in section 4.5.2).

4.5 Remote Access

4.5.1 API

No user should encounter a need to directly intervene with the currently running instance of the database. An Application Programming Interface (API) is presented here as a full and sufficient means for interacting with the database indirectly. For any use case which exists, the developer can provide a function within the API that matches the appropriate input and output whilst hiding the technical specifics of the database implementation from the user. All the user should require to know is the name (and appropriate input and output parameters) of the function which they should call. This should make calls to the database straightforward and reproducible.

In addition, the implementation of the database and of client applications may change. The user, whether a home user or a Grid user, will be using an unknown version of their application code. Control room applications may also be subject to changes. Since all interactions with the database are mediated through the API, the actual code should only be changed in one place (the API) at maximum, rather than many places at minimum.

A well documented API will also be robust to people turnover. The author is therefore focussing his resources not on writing many different pieces of code for a number of client applications, but on a single API containing all the detailed code which those disparate client applications can easily use.

4.5.2 SOAP Connectivity for Remote Access

SOAP [138], Simple Object Access Protocol, describes the transmission of XML [148] (eX-tensible Markup Language, a platform independent structured and general purpose encoding for textual information) via an 'application layer' protocol such as, in our case, HTTP [150]. The database API is configured as a web service, which is a software resource responsive

over the internet to, in this case, remote SOAP calls. Procedure calls and return values can be sent and received via SOAP and this forms the basis of all read and write operations involving the database API. Support is also provided for error feedback within the envelope of a SOAP message.

By describing communication between the database and its clients in XML, one secures platform independence, as illustrated in the simplest example of the Java [151] API communicating with the EPICS system, which is implemented in C++ [152]. In addition, one can rigorously define the interface between the database and the outside world, allowing separation of client and server logic one from another. Just as development of client applications may proceed independently of database expertise, but guided by the XML description of the interface with the database, so too are the needs of unknown future applications supported in the extensibility of our XML schema. Roll out of database connectivity to client applications is aided by the provision of SOAP message templates which developers of client applications can use to interface with the database.

4.5.3 Web Interface

Key use cases require mapping of run meta data to real configurations. Users must be able to project from the database those run configurations that correspond to given cuts or selections in meta data parameters such as timing, beam species, etc. In addition, given a run number, a user must be able to directly view the known state of the configuration of that run.

This information must be available to non-technical users who are not proficient in client applications such as G4MICE, EPICS etc or in database query languages such as SQL. Therefore, one page web interfaces have been developed which allow users to directly input their parameters and receive on-screen feedback.

A web interface has been implemented, comprising screens dealing with querying runs based on parameters given in the run table, tagged presets retrieval, a summary of all runs stored and an expert panel for direct read-only SQL requests as shown in figure 4.11. Users may request lists of runs matching some criteria, here attributes of the Run table in figure 4.4

Set Values Tagged Set Values Summary Expert

Configuration Metadata Viewer

"Give me the runs which meet these conditions..."

Give me the list of run numbers which match the following known conditions:

Run Number	>	2000
Species		pipius
Momentum at D1	=	300
Step	=	1

Add Condition Do It

Figure 4.11: Web interface for the Configuration Database, allowing lookup of previous run conditions, review of tagged presets, a summary of all runs previously undertaken, and specialist queries from developers and expert users.

are given in drop down menus. Users may query runs for values less than, greater than, or equal to each of the attributes given in the drop down menu. A list of runs matching this criteria is generated, and the user may click individual run numbers for a full summary of the beamline settings for that run. This interface is available at [153] and is currently running. It has been implemented using PHP and Javascript and in all cases their inputs are sanitised to avoid the insertion of malicious code into user entry boxes which would otherwise then be inserted into legitimate queries in an 'SQL injection' hijacking attack. Additional security is provided by granting read only access through the web clients.

Further web interfaces can be rapidly deployed by continuing MICE collaborators through use of the existing interfaces as templates.

4.5.4 Security & Stability

The probability of unauthorised access to the database should be minimised, and the effects of any unauthorised access should be isolated and reversible. To this end one must enforce strict access levels and perform timely backup and logging.

As a rule, all processes authorised to write to the database must originate in the MICE Control Room and be under the control of the MICE Operations Manager who may delegate

to those on shift. The control room enjoys a reserved link with the database, isolating it from any unexpected effects of high traffic read operations which are in any case unlikely. No write control is granted to any other database connection.

Web interfaces and client applications interacting with the database are permitted read access only. In addition, their inputs are examined and filtered closely to prevent circumvention of security through SQL injection attacks which can arise from malicious insertion of SQL code within user input which might otherwise be executed within an SQL query.

Client interaction through the database API allows a level of abstraction between the database and the user, allowing additional checks on user input to take place. The API can also be easily mirrored from a backup copy guaranteed to be no more than five minutes old, ensuring access throughout local instabilities such as power outages. Should a security issue arise on the API server at RAL, that server will be isolated and mirroring will allow access to the database during interventions by RAL networking.

4.6 Testing

As part of commissioning it was necessary to demonstrate that the Configuration Database System (CDS) could cope with exceptional levels of traffic without compromising performance. A test was devised whereby 50 Grid processors would send continuous messages to the CDS for a sustained period of one minute, recording and checking the validity of the CDS response. Special care was taken to ensure that each sub-job would begin sending requests at the same time. Every node is guaranteed to be synchronised with all others to millisecond precision. Although it is impossible to ensure that each job is submitted to the Grid and allocated to a resource at exactly the same time, each job was instructed to sleep until a set time when they would all be assumed to be initialised and ready to send requests synchronously. Each request included connection, sending of an XML request to read information from the database, receiving an XML response, performing a simple operation on the response and closing the connection. This represents a realistic instance of client interaction. An average of roughly 600 such requests were made per second, with transaction times given

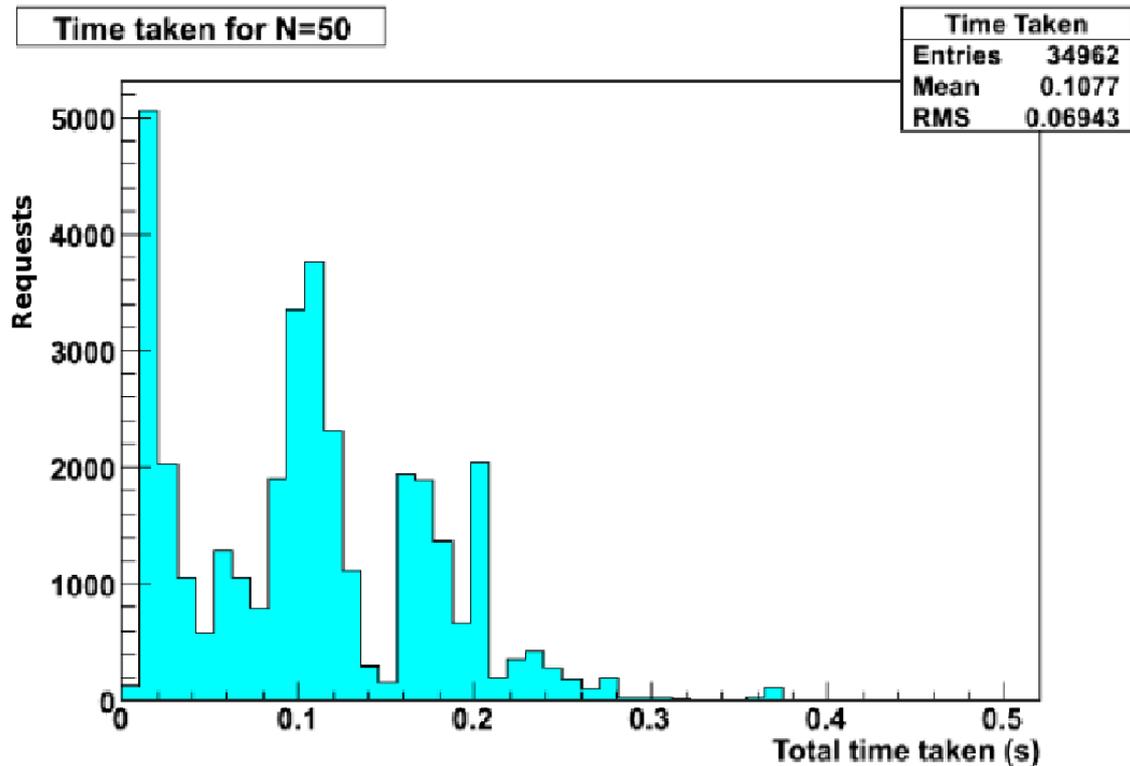


Figure 4.12: 50 processors sent constant synchronous requests to the CDS for a sustained period of one minute. Response times are distributed here. These response times refer to read requests which formed the basis of our background traffic. The success or failure of the test would be determined on the time taken for a write request from a separate CPU simulating control room processes. This was found to be 0.110s on average in the presence of traffic and 0.109s without.

in figure 4.12.

An additional processor simulated control room processes, making write requests during traffic and without traffic. Without traffic the average response time was 0.109s for a write operation. This compares to 0.110s with traffic. Within reasonable expectations given variations in network traffic, there is no evidence of significant slow-down.

It should be noted that the test case far exceeds the likely traffic requirements for MICE which has limited uses for and access to Grid processing in the medium term. Furthermore the very nature of bulk job submission through the Grid to remote sites at least makes it highly probable that jobs will begin asynchronously.

4.7 Client-side Applications

Development of a range of G4MICE and Python [154] based client applications for the CDS is currently underway [155]. Whilst the server side system is fully operational, its full capabilities have yet to be fully exploited. However EPICS-based client applications supporting control room processes have been written by James Leaver (Imperial) and are fully operational. This section will provide a snapshot of each.

Beamline settings are written to the database at the conclusion of the run via a client application built into the EPICS framework. The majority of parameters such as magnet currents and timestamps are automatically recorded, removing the opportunity for human error. The remainder of the configuration parameters are recorded manually by experimenters in a form, shown in figure 4.13. There is no constraint enforced by the CDS for manual entry and it is hoped that as client applications mature further automation will be possible, greatly reducing the number of boxes in figure 4.13. Communication between EPICS and other systems such as the DAQ is required for such automation.

Tagged runs are used at the beginning and end of an experimental shift to automatically invoke a reference run or restore settings to an equilibrium (all magnets off except decay solenoid). The application shown in figure 4.14 allows the user to enter a run number or a preset string identifier corresponding to a tagged setting. These settings are then reinstated. At any point the settings may be reviewed using the web interface described in section 4.5.3.

4.8 Conclusions and Future Work

The analysis, scope, design, implementation and testing of the MICE Configuration Database has been presented. This database is currently running at Rutherford Appleton Laboratory, complementing data taking and control systems with geometry, calibration, cabling, beam-line settings and alarm handler limit information. Having been developed exclusively by the author to date, control of the database system will now revert to the MICE Collaboration for any future enhancement or maintenance required. Training has been provided through

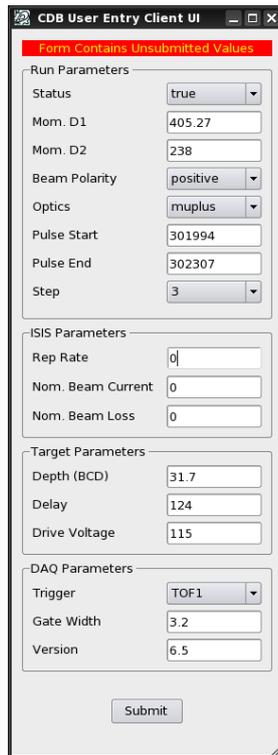


Figure 4.13: Manual interface for entering set values into the CDS. The majority of parameters are automated. The remainder require manual user entry in the current version of the client-side application. The number of parameters available for manual user entry should reduce as a function of time.



Figure 4.14: An EPICS application to reinstate settings from a previous run or associated with a preset string tag.

a workshop attended by roughly 15 collaborators, a number of seminars at collaboration meetings and significant documentation including a MICE Note and audio podcasts of the workshop [156, 157].

Remote access for the database through G4MICE, EPICS and a web interface presently exists. In addition one can run code on the Grid and receive a timely response from the database API. This is achieved with minimal extra effort due to the portability of SOAP and the ubiquity of HTTP.

At present the geometry, calibration, beamline, cabling, set value, tagged set value and alarm handler limit systems meet all their functional requirements and can be measured as complete, although they are also extensible to future potentially unknown needs. Further work is required to fully realise the cabling domain of the database. In addition, although a backup strategy exists and has been described here, further work is required to produce a replication strategy for the database making the configuration systems robust to local downtime. In such a system one envisages that participating MICE institutions could receive a copy of the database every few seconds, and run their own API web service which operates on their local copy of the database. At all times the version in the MICE control room should be considered the master, although in practice replications should have a very short periodicity.

Chapter 5

Statistical Errors and Alignment

Tolerances of Emittance Change

Measurements in MICE

5.1 Introduction

The primary goal of MICE is to measure fractional change in emittance ($\frac{\epsilon_{in}-\epsilon_{out}}{\epsilon_{in}}$) of order 10% to an error of 1%. Attaining an understanding of both the statistical errors and systematic errors involved in the experimental measurement is crucial to achieving this goal.

From Step III of MICE onwards, emittance will be measured by the two scintillating fibre trackers positioned at either end of the MICE channel. Crucially, these detectors measure emittance on roughly the same sample of muons, thereby introducing correlations in the two measurements, complicating a statistical analysis. In order to quantify the statistical errors for a range of different input beam configurations which are likely to be studied in MICE, a large number of simulations with different pseudo-random conditions was therefore required. Grid technology therefore found its first application within MICE [158] in enabling a first study of the statistical errors of the experiment, presented in section 5.3, to determine whether the experimental goals are achievable and to calculate the number of muons required to attain

them.

The same technology was also employed in studying the alignment tolerances of the scintillating fibre trackers. Each tracker is contained within a solenoid producing a 4 T magnetic field. The combination of tracker and solenoid is referred to as a tracker module, both of which measure 2923 mm in length and 1404 mm in outer diameter. Inevitably, misalignments will occur during the lifetime of the experiment due to mechanical tolerances or the effect of magnetic fields. Section 5.4.1 details a simulation-based study into the sensitivity of the tracker modules to misalignment, a key systematic uncertainty.

5.2 Emittance Change Through MICE

In the absence of acceleration, emittance is a conserved quantity. It follows that in the main, transverse emittance change in MICE should occur at the absorbers where the particles are decelerated through energy loss. Figure 5.1 shows an example of ionisation cooling in a MICE Step VI configuration simulated using G4MICE with an initial transverse emittance of 6.2 mm rad and with muon momentum of 200 MeV. The near-step like structure confirms that transverse emittance drops in the absorbers and remains flat elsewhere, with some small deviations due to multiple scattering in the cooling channel materials.

The effect of ionisation cooling in MICE varies as a function of input beam emittance. Figure 5.2 illustrates the emittance change calculated with G4MICE in the same way as figure 5.1 but as a function of input transverse emittance. One observes a greater fractional change in emittance $\left(\frac{\epsilon_{out}-\epsilon_{in}}{\epsilon_{in}}\right)$ in the case of higher emittance input beams. We note that at around 2.5 mm rad, referred to as the 'equilibrium emittance', no cooling is observed. Below this value, we observe a fractional increase in emittance due to the dominant effect of multiple scattering for low emittance beams. Figure 5.2 does not provide statistical errors. This is one of the purposes of the study in this chapter and will be calculated in section 5.3.

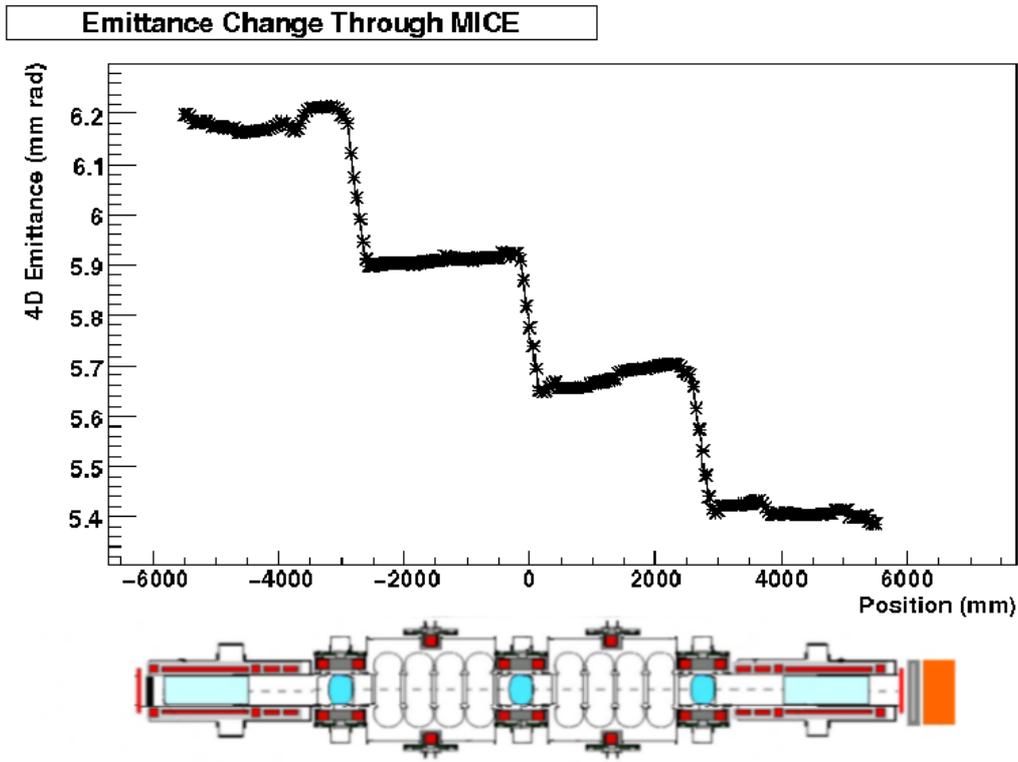


Figure 5.1: Emittance as a function of z -position in MICE Step VI, where z is along the axis of the beam. The central absorber is positioned at $z = 0$. Although deviations exist partly due to multiple scattering and limitations in the simulation, it can be shown that 4D transverse emittance reduces in steps along the MICE channel.

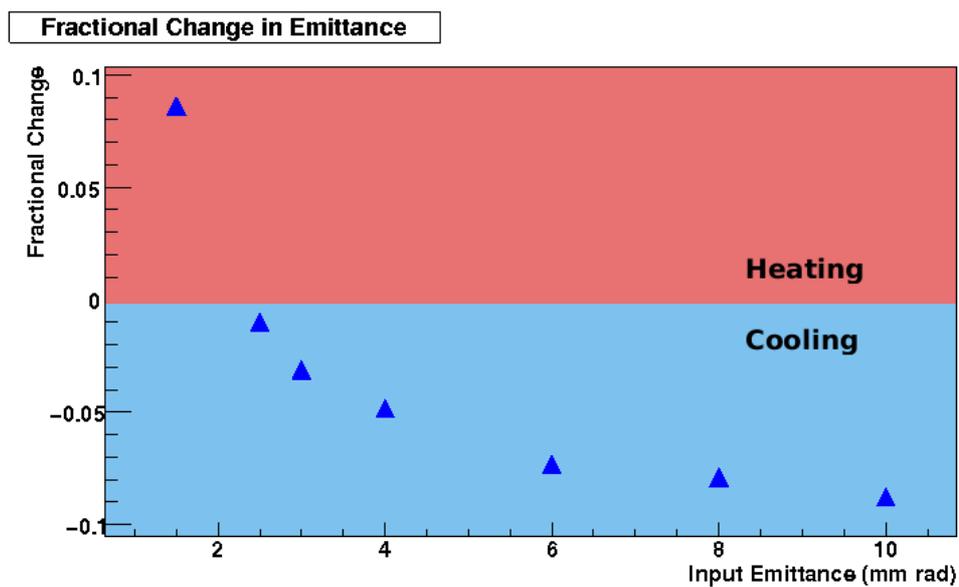


Figure 5.2: Fractional change in emittance as a function of input beam emittance.

5.3 Statistical Errors of Fractional Change in Emittance in MICE

In this section, we will carry out an analysis of the statistical errors associated with the measurement of fractional change in emittance in MICE, taking into account fully the correlations between the initial and final measurements of emittance. This was carried out by performing a large number of G4MICE simulations on the Grid and determining the correlation factors empirically to quantify the statistical errors and any bias in the measurement. This analysis was performed entirely by the author. After these results were published [159, 158] J. Cobb, a MICE collaborator from Oxford University, performed a theoretical analysis of the statistical errors of the fractional change in emittance in MICE that confirmed the results discussed in section 5.3.2.5. The theoretical analysis will be presented first in section 5.3.1 to illustrate the problem facing the calculation of these statistical errors.

5.3.1 Theoretical Analysis of Statistical Errors and Correlations

In this section we will follow the theoretical treatment by Cobb [160]. Let us assume that the standard deviation of fractional change of emittance can be expressed in the following form:

$$\sigma = K \frac{1}{\sqrt{N}}, \quad (5.1)$$

where N is the number of events and K is a constant of proportionality varying as a function of beam configuration. Since the measurement of fractional change in emittance is made at two points on a similar set of muons (allowing for transmission losses and muon decay), we expect that the statistical error should be proportional to $\frac{1}{\sqrt{N}}$ and that the effect of correlations will require K to be less than unity for any typical beam. In the absence of stochastic processes the statistical error is zero, therefore we further expect to find that K is strongly dependent upon scattering effects in the absorbers and that the statistical error does not otherwise vary with position.

Following Cobb [160] and considering the two dimensional case, we suppose that the

emittances ϵ_1 and ϵ_2 , measured before and after the cooling channel are both calculated at the central plane of an absorber. Since the absorbers are thin they may be combined together,

$$m_\mu^2 c^2 \epsilon_1^2 = \sigma_x^2 \sigma_{p_x}^2 \quad (5.2)$$

$$m_\mu^2 c^2 \epsilon_2^2 = \sigma_x^2 \sigma_{q_x}^2 \quad (5.3)$$

with p_x the x momentum before the absorber and q_x the x momentum after the absorber, related as follows:

$$q_x = \alpha p_x + s, \quad (5.4)$$

where $(1 - \alpha)$ is the fractional drop in momentum due to energy loss at a momentum p_x . The variable s is the momentum change due to multiple scattering and is sampled from a distribution centred on 0 with a standard deviation given by equation 5.5.

$$\sigma_s^2 = \frac{(13.6 \frac{\text{MeV}}{c})^2 \Delta X}{\beta_L^2 X_0}. \quad (5.5)$$

Equation 5.5 gives the variance of the momentum due to multiple scattering in a thickness $\frac{\Delta X}{X_0}$ of absorber and β_L is the Lorentz particle velocity factor. The variance of the x momentum is:

$$\sigma_{p_x}^2 = \frac{p_z m_\mu \epsilon_1}{\beta_\perp}, \quad (5.6)$$

where β_\perp is the transverse beta function. Combining equations 5.3 and 5.4, on average we will have:

$$m_\mu^2 c^2 \epsilon_2^2 = \sigma_x^2 (\alpha^2 \sigma_{p_x}^2 + \sigma_s^2). \quad (5.7)$$

We define a quantity R, such that:

$$R = \frac{\epsilon_2^2}{\epsilon_1^2} = \frac{\sigma_{q_x}^2}{\sigma_{p_x}^2} = \alpha^2 + \frac{\sigma_s^2}{\sigma_{p_x}^2} \quad (5.8)$$

Therefore, the fractional change in emittance may now be written as in Equation 5.9:

$$f = 1 - R^{\frac{1}{2}} \quad (5.9)$$

with the standard deviation given by:

$$\sigma_f = \frac{1}{2(1-f)} \sigma_R, \quad (5.10)$$

assuming we know σ_R , the error in R. We will now calculate the measured value of R, which we will denote R_m based on the statistical analysis of a finite sample of N muons, where the measured values of the variances are $\sigma_{p_x}^2$ and $\sigma_{q_x}^2$ and the mean values are $\mu_p = \bar{p}_x$ and $\mu_q = \bar{q}_x$. In this case:

$$\sigma_{p_x}^2 = \overline{p_x^2} - \mu_p^2 \quad (5.11)$$

$$\sigma_{q_x}^2 = \overline{q_x^2} - \mu_q^2 \quad (5.12)$$

Without loss of generality, we can take the mean values as zero so that it follows:

$$\sigma_{p_x}^2 = \frac{1}{N} \sum_{i=1}^N p_i^2 \quad (5.13)$$

$$\sigma_{q_x}^2 = \frac{1}{N} \sum_{i=1}^N q_i^2 \quad (5.14)$$

$$(5.15)$$

and propagating the error in equation 5.4 we obtain:

$$\sigma_{q_x}^2 = \frac{1}{N} \alpha^2 \sum_{i=1}^N p_i^2 + \frac{2\alpha}{N} \sum_{i=1}^N p_i s_i + \sum_{i=1}^N s_i^2 \quad (5.16)$$

The measured value of R is given by:

$$R_m = \frac{\frac{1}{N}\alpha^2 \sum_{i=1}^N p_i^2 + \frac{2\alpha}{N} \sum_{i=1}^N p_i s_i + \sum_{i=1}^N s_i^2}{\frac{1}{N} \sum_{i=1}^N p_i^2} \quad (5.17)$$

$$R_m = \alpha^2 + 2\alpha \frac{\frac{1}{N} \sum_{i=1}^N p_i s_i}{\frac{1}{N} \sum_{i=1}^N p_i^2} + \frac{\frac{1}{N} \sum_{i=1}^N s_i^2}{\frac{1}{N} \sum_{i=1}^N p_i^2} \quad (5.18)$$

The expectation value (but not the variance) of $\frac{1}{N} \sum_{i=1}^N p_i s_i$ averages to zero over many measurements, therefore:

$$\bar{R}_m = \alpha^2 + \frac{\sigma_s^2}{\sigma_{p_x}^2}, \quad (5.19)$$

which compares well with equation 5.8.

We wish to calculate the variance of R_m which follows from the calculation of the variances and covariances of its constituent terms. While $\frac{1}{N} \sum_{i=1}^N p_i s_i$ averages to zero, its estimate of the variance of the covariance term S_{ps} is not zero. Therefore:

$$R_m = \alpha^2 + \frac{2\alpha S_{ps} + S_{ss}}{S_{pp}}, \quad (5.20)$$

where S_{ss} and S_{pp} are estimates of the variances σ_s^2 and $\sigma_{p_x}^2$. The variances of these terms are given as:

$$\text{Var}(S_{ss}) = \frac{2}{N} (\sigma_s^2)^2 \quad (5.21)$$

$$\text{Var}(S_{pp}) = \frac{2}{N} (\sigma_{p_x}^2)^2 \quad (5.22)$$

Cobb [160] provides an estimate of the variance of the covariance S_{ps} in equation 5.23:

$$\text{Var}(S_{ps}) = \frac{1}{N} \sigma_{p_x}^2 \sigma_s^2 \quad (5.23)$$

The variance on R is given by:

$$\sigma_R^2 = \left(\frac{\partial R}{\partial S_{ps}} \right)^2 \text{Var}(S_{ps}) + \left(\frac{\partial R}{\partial S_{ss}} \right)^2 \text{Var}(S_{ss}) + \left(\frac{\partial R}{\partial S_{pp}} \right)^2 \text{Var}(S_{pp}) \quad (5.24)$$

$$= \left(\frac{2\alpha}{S_{pp}} \right)^2 \text{Var}(S_{ps}) + \left(\frac{1}{S_{pp}} \right)^2 \text{Var}(S_{ss}) + \left(\frac{2\alpha S_{ps} + S_{ss}}{S_{pp}^2} \right)^2 \text{Var}(S_{pp}) \quad (5.25)$$

$$= \frac{1}{N} \left[\left(\frac{2\alpha}{\sigma_{p_x}^2} \right)^2 \sigma_{p_x}^2 \sigma_s^2 + 2 \left(\frac{1}{\sigma_{p_x}^2} \right)^2 (\sigma_s^2)^2 + 2 \left(\frac{\sigma_s^2}{(\sigma_{p_x}^2)^2} \right)^2 (\sigma_{p_x}^2)^2 \right] \quad (5.26)$$

$$= \frac{4}{N} \frac{\sigma_s^2}{\sigma_{p_x}^2} \left[\alpha^2 + \frac{\sigma_s^2}{\sigma_{p_x}^2} \right] \quad (5.27)$$

If scattering effects were not present ($\sigma_s = 0$), our statistical error would be zero also.

Combining equations 5.8 and 5.27 we have:

$$\sigma_R^2 = \frac{4}{N} R \left(\frac{\sigma_s^2}{\sigma_{p_x}^2} \right). \quad (5.28)$$

Recalling equation 5.10, this leads us to an equation for the variance in fractional change in emittance.

$$\sigma_f^2 = \frac{\sigma_R^2}{4(1-f)^2} \quad (5.29)$$

$$= \frac{\sigma_R^2}{4R} \quad (5.30)$$

$$= \frac{1}{N} \left(\frac{\sigma_s^2}{\sigma_{p_x}^2} \right) \quad (5.31)$$

There is no α term, and so no energy loss dependency, but a dependency upon the ratio of the variance in scattering with the variance in x -momentum. Given set beam conditions, we have a factor of proportionality, K , which relates the error with the inverse square root of the

number of events as in equation 5.1. Equation 5.31 gives the factor as $K = \frac{\sigma_s}{\sigma_{p_x}}$

In the real experiment, the momenta and positional variables which are used to calculate emittance are highly correlated at the two trackers, as measurement is performed upon a similar sample of muons. This introduces a correlation factor into equation 5.31, given by the ratio of σ_f and the statistical error in the uncorrelated case, $\sigma_{f_{nocorr}}$. If we assume that the error in ϵ_1 and ϵ_2 are uncorrelated, and

$$\frac{\sigma_{\epsilon_1}}{\epsilon_1} = \frac{\sigma_{\epsilon_2}}{\epsilon_2} = \frac{1}{N} \quad (5.32)$$

then:

$$\sigma_{f_{nocorr}}^2 = \frac{2}{N} \frac{\epsilon_2^2}{\epsilon_1^2} \quad (5.33)$$

$$= \frac{2R}{N} \quad (5.34)$$

Therefore, we can define a correlation factor k_c :

$$k_c = \frac{\sigma_f}{\sigma_{f_{nocorr}}} \quad (5.35)$$

$$= \frac{1}{\sqrt{2R}} \left(\frac{\sigma_s}{\sigma_{p_x}} \right) \quad (5.36)$$

$$= \frac{1}{\sqrt{2 \left(\alpha^2 + \frac{\sigma_s^2}{\sigma_{p_x}^2} \right)}} \left(\frac{\sigma_s}{\sigma_{p_x}} \right). \quad (5.37)$$

Where scattering dominates, as in the case for small beams, $\frac{\sigma_s^2}{\sigma_{p_x}^2} \gg \alpha$ and k_c approaches $\frac{1}{\sqrt{2}}$. In the case of minimal scattering k_c approaches 0.

We may easily extend equation 5.31 to four dimensions, where four dimensional emittance is the average of x and y emittances. Measurements are made where x and y motions are decoupled with respect to one another. It follows that the error in the fractional change in 4D emittance is the error in fractional change in 2D emittance divided by $\sqrt{2}$, that is:

$$\sigma_f^2|_{4D} = \frac{1}{2N} \left(\frac{\sigma_s^2}{\sigma_{p_x}^2} \right). \quad (5.38)$$

This can also be written as a function of f_m , the measured fractional emittance decrease, by using equation 5.8:

$$\sigma_f^2|_{4D} = \frac{1}{2N} ((1 - f_m)^2 - \alpha^2). \quad (5.39)$$

We have now arrived at an equation which is suitable for the purposes of quantifying the statistical error in MICE and therefore predicting the number of muons required to meet the goals of the experiment.

This analytical treatment by Cobb [160] was validated by the empirical study presented in section 5.3.2 of this chapter but it was carried out after the conclusions from that study.

5.3.2 Empirical Study of Statistical Errors in MICE

5.3.2.1 Use of the Grid

Empirical verification of the statistical errors in the emittance calculation in MICE was carried out by performing a vast number of G4MICE simulations. This was made possible through the processing power of the Grid and aided considerably by the archiving and data management potential of Grid storage resources. In addition, the use of Ganga (section 5.3.2.2), a Python based interface for job submission, greatly simplified job management. Ganga provided a level of abstraction for the roughly forty thousand simulations involved in this study, by virtue of composite objects called bulk jobs.

The work represented the first use of the Grid for the MICE experiment, and as such was a useful learning experience [161] which can now be drawn upon by the wider collaboration for other uses, as well as attracting interest from the e-Science community [158] as an application of the Grid to a very large computational problem.

5.3.2.2 Ganga

Much benefit was derived from the use of the Ganga tool [162], an optional job submission management software program. By default one manages job submission manually through a suite of `glite-wms` workload management system commands for distributed computing [163] and job description language files (`jdl` files) [164] which introduce a learning overhead and substantial opportunities for error. `jdl` files in particular are prone to user error and use technical distributed computing terms as part of attribute-key pairs, which the user must correctly manipulate to define the job.

Ganga allows one to easily parameterise jobs, in Python, with a level of abstraction from the `glite-wms` commands and `jdl` files associated with conventional job submission. This allows the study to be easily understood, repeated and modified by members of the collaboration who presently lack Grid experience.

Ganga provides book-keeping facilities for status polling and archiving of job input, scripting and output, providing a record of running conditions for each simulation, proving that they are statistically independent from one another. The usefulness of this functionality is particularly apparent when running tens of thousands of simulations.

Ganga also provides 'splitters' which facilitate the ready definition of composite jobs. An argument `splitter` allows composite jobs to be defined based on one job script. Sub-jobs could be discriminated from one another by the randomisation seed, number of events and beam conditions provided as arguments. Batch job submission, interfaced exclusively via Python, together with built-in book-keeping functionality, greatly simplified the challenge of management with such a large number of jobs (see figure 5.3).

Parameterisation allowed these large statistics studies to run from a single well-understood simulation script provided by the analysis user. The statistical study exploited a randomisation seed parameter defined in the script to ensure that the random seeds were independent. In the misalignment study of section 5.4.1, details of the geometry changes to be constructed were passed through a single string tag. The statistics and alignment studies are therefore directly comparable and a reliable combined error can be calculated.

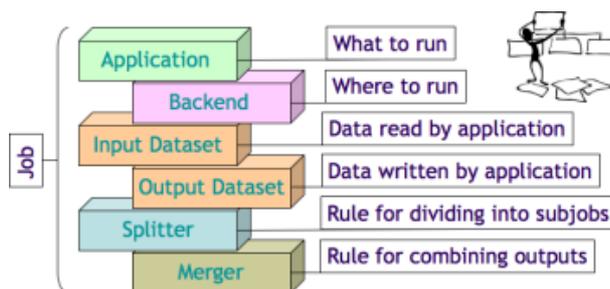


Figure 5.3: The structure of a job in the Ganga Job Submission software. Abstraction over application and backend afford the system generality over many different applications and systems (including Grid and batch systems). Splitters and mergers allow for the ready definition of composite jobs. Image taken from the Ganga project page [165].

5.3.2.3 G4MICE Simulations

This study represented the first serious examination through simulation of the statistical error in MICE. The primary aim was to determine the constant of proportionality K in equation 5.1 that relates the error in the fractional change in emittance and $\frac{1}{\sqrt{N}}$, where N is the number of muons in an experiment and to verify empirically the scaling law that governs this error as a function of N for a number of different beam emittance configurations. This was achieved by carrying out a large number of simulations where the beam parameters and number of events were varied. By quantifying the scale factor K , one can determine the number of particles required to meet the statistical error requirements of the MICE experiment, as in [159].

A range of beam conditions will be studied during the course of the MICE experiment. In this study, over 1,500 simulations were run for each beam configuration with at least 500 each for runs of $N = 1,000, 2,000$ and $10,000$ particle events. There were eight beam configurations which were considered for this study, with initial emittances of 0.2, 1.5, 2.5, 3.0, 4.0, 6.0, 8.0 and 10.0 mm rad. At least 12,000 simulations were required, although not all were selected for the final analysis. Additionally, retuning and subsequent repeating of the simulations meant that around 50,000 simulations were run and analysed. This study would have taken years, given the computing resources locally available to MICE for this purpose, but instead was able to be carried out in a reasonable time using high performance Grid computing discussed in section 5.3.2.1.

All simulations were performed using the final Step VI configuration of the MICE cooling channel in figures 3.3 and 3.4.

5.3.2.4 Results

For each beam configuration, sets of roughly 500 different pseudo-experiments were run with 1,000, 2,000 and 10,000 muons simulated by G4MICE through the MICE Step VI cooling channel. For each experiment, the fractional change in emittance $f = \frac{\epsilon_f - \epsilon_i}{\epsilon_i}$ was calculated and the distribution of all values of f for all simulated experiments at a fixed number of muons was plotted. A Gaussian fit was applied to each of these distributions and the mean fractional change in emittance and standard deviation σ_f were calculated for the three sets of runs with $N = 1,000, 2,000$ and $10,000$ muons.

Figure 5.4 shows the three distributions, one for each value of N , and their fit for a beam of initial emittance 0.2 mm rad. One can see how the σ_f diminishes as the number of muons in each simulation increases (figure 5.4d). Figures 5.5-5.11 show the same plots but for an initial beam of emittance 1.5, 2.5, 3.0, 4.0, 6.0, 8.0 and 10.0 mm rad.

Figure 5.12 shows a summary of all results, in which σ_f is plotted against $\frac{1}{\sqrt{N}}$ for each of the beam configurations. A linear dependence for each graph was confirmed and the linear fit provides the constant K of equation 5.1.

5.3.2.5 Discussion

The constant of proportionality, K , was found to be different for each beam configuration but crucially less than 1 for larger beam emittance configurations more pertinent to the goals of MICE. Most K values, shown in figure 5.13 were around 0.3 for higher emittance beams. This was due to the effect of correlations, previously described, which act to reduce the statistical error. This study quantified that effect for the first time. The measured standard deviation includes the correlation factor k_c defined in equation 5.37:

$$\sigma_{sim} = k_c \sigma_{nocorr} \tag{5.40}$$

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

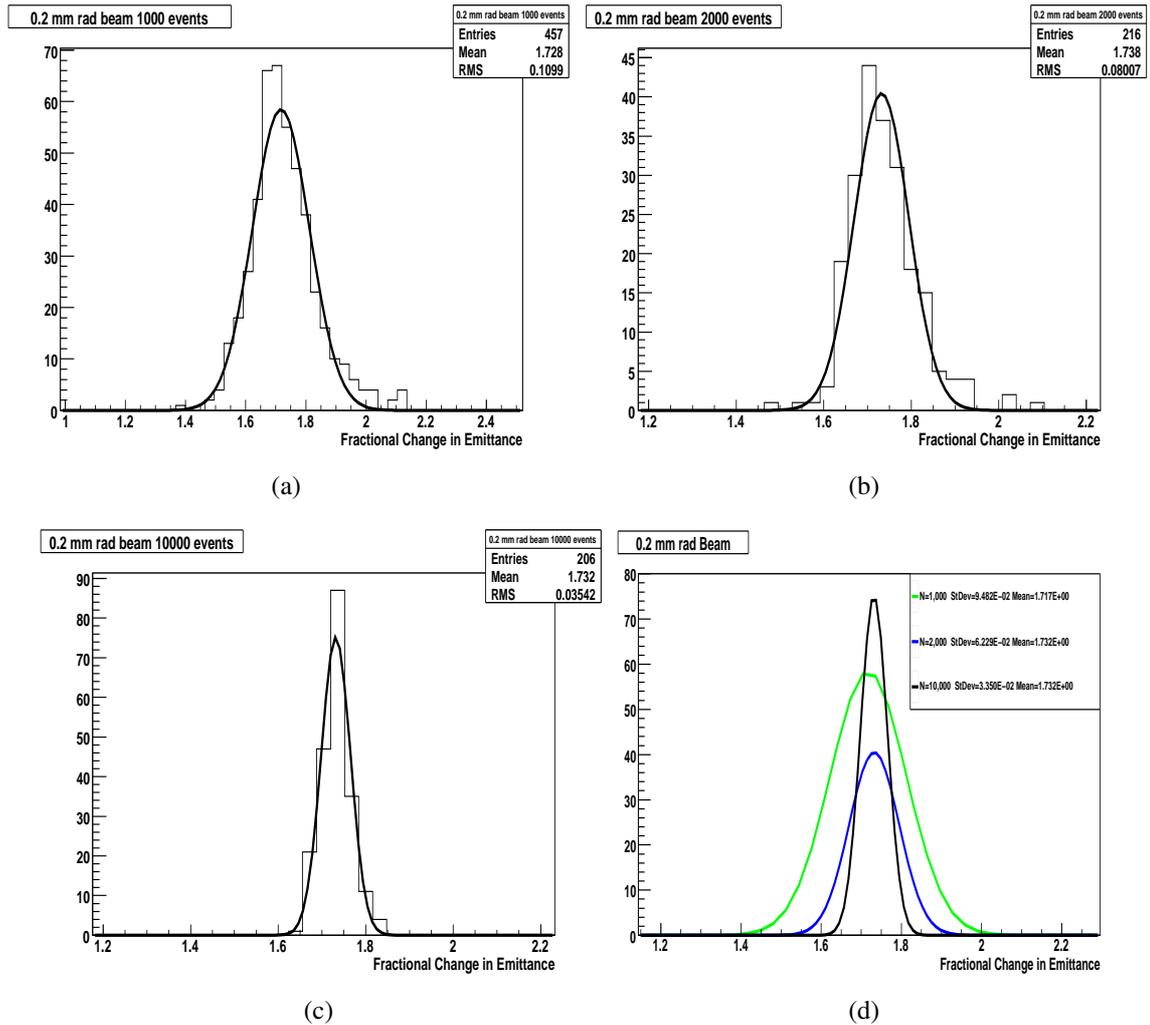


Figure 5.4: Distributions of fractional change in emittance for a 0.2 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

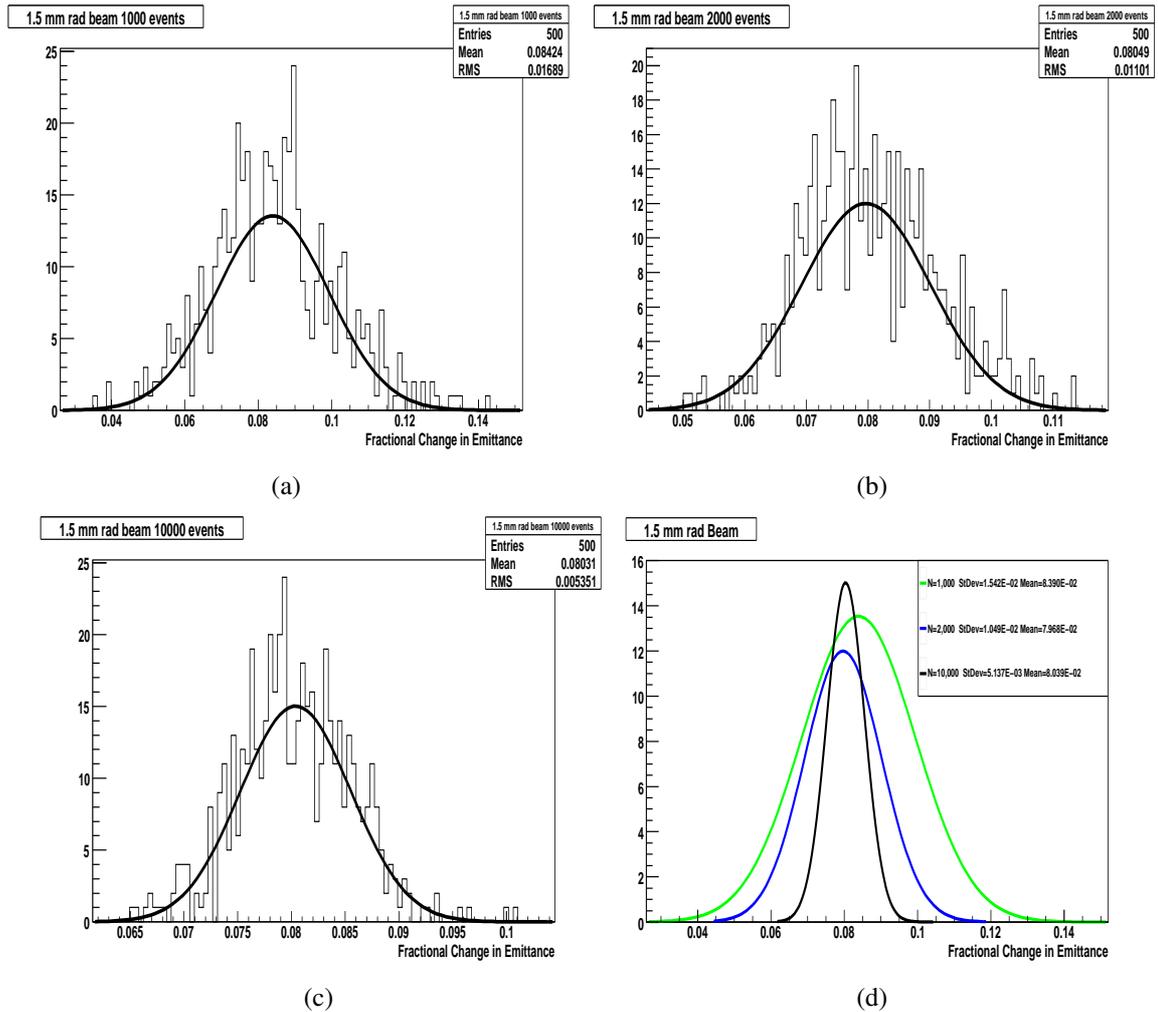


Figure 5.5: Distributions of fractional change in emittance for a 1.5 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

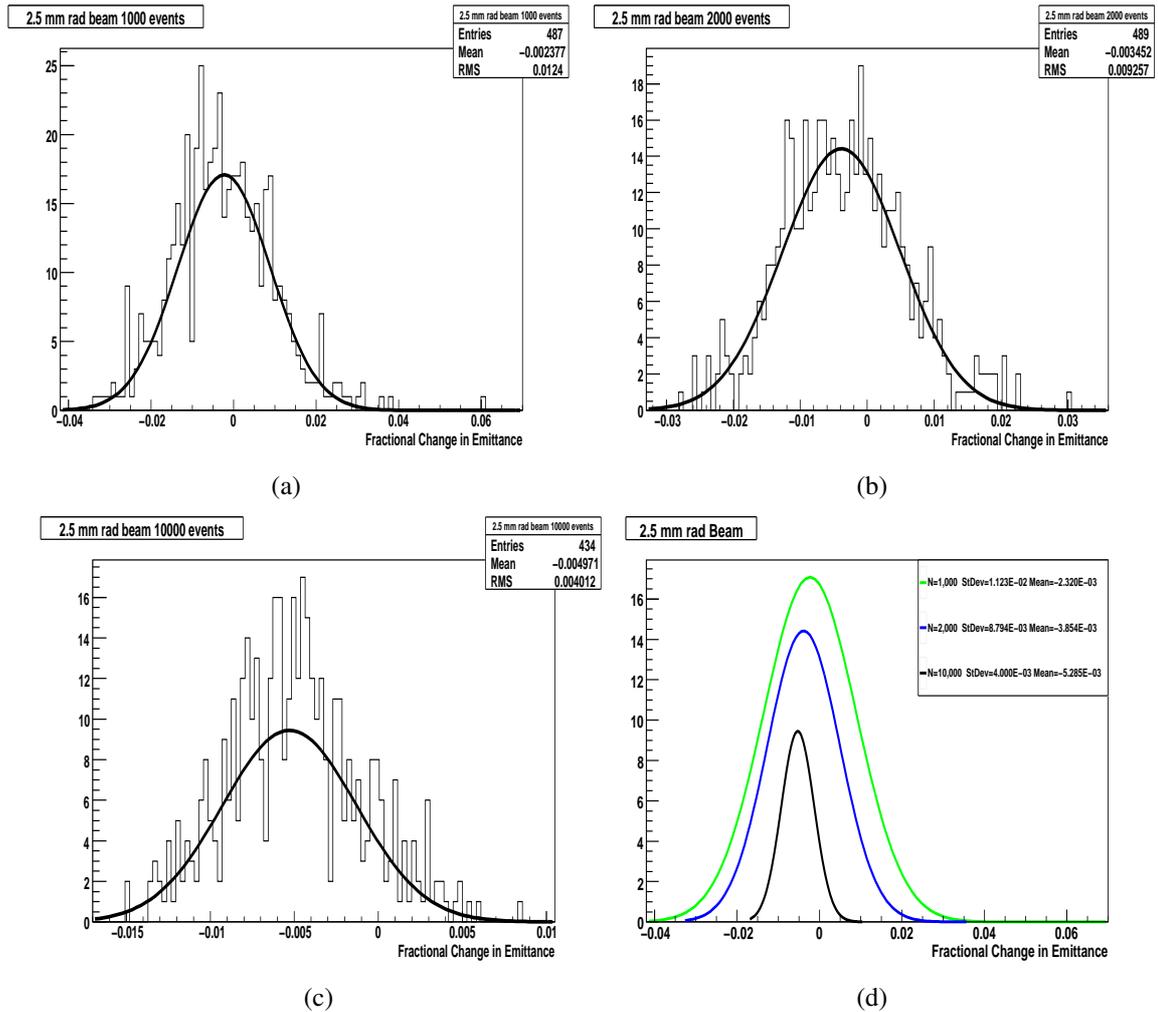


Figure 5.6: Distributions of fractional change in emittance for a 2.5 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

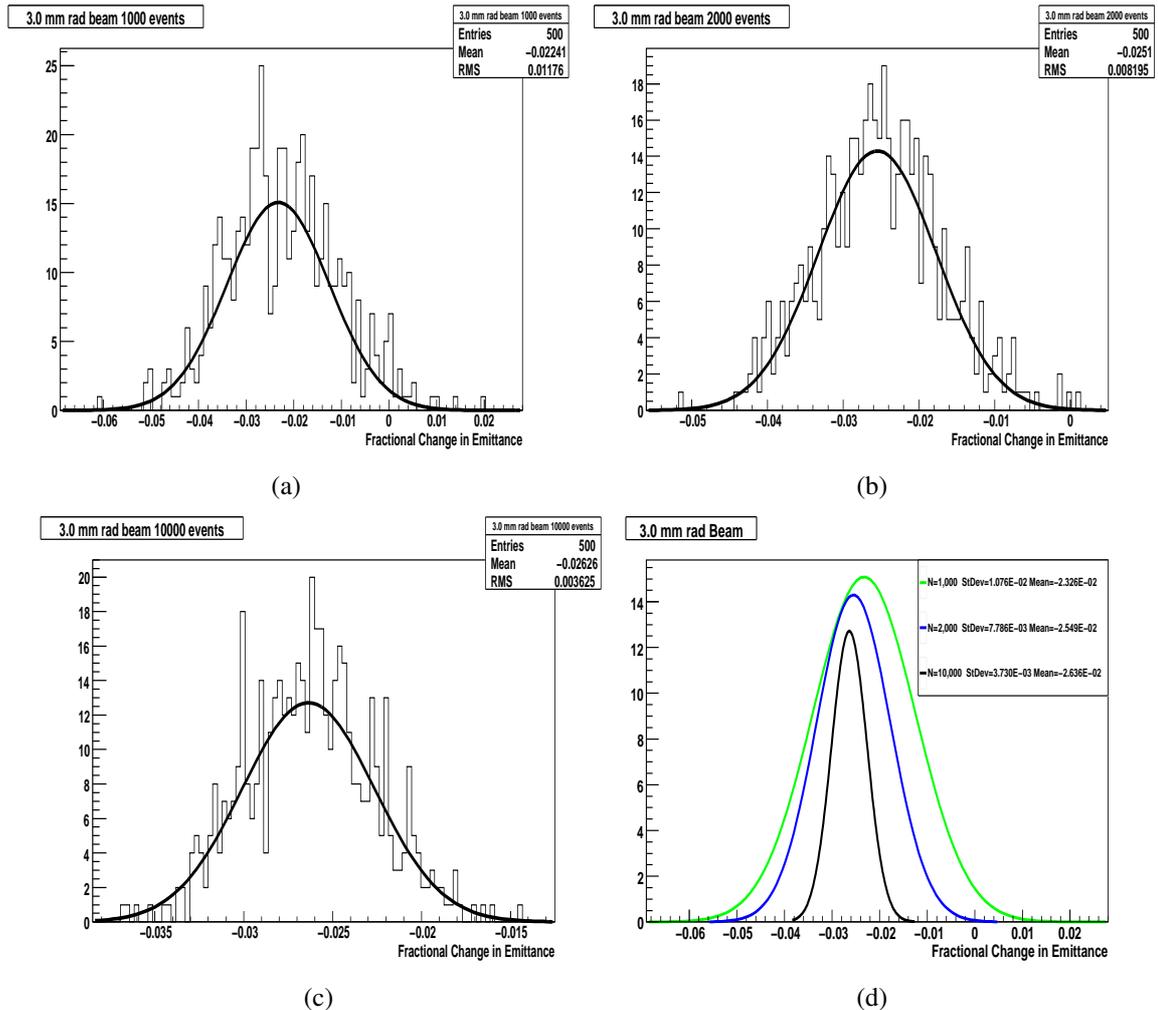


Figure 5.7: Distributions of fractional change in emittance for a 3.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

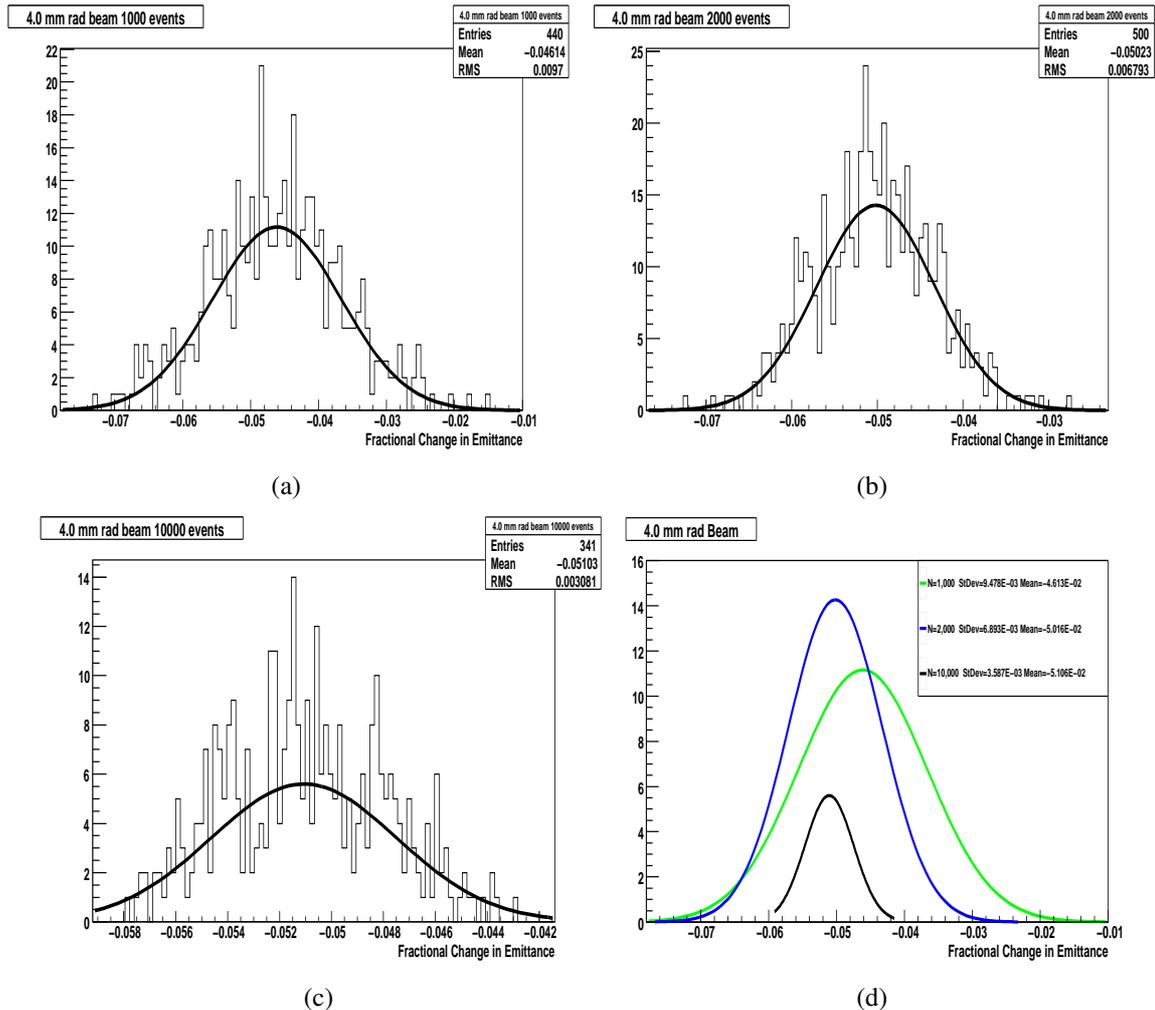


Figure 5.8: Distributions of fractional change in emittance for a 4.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

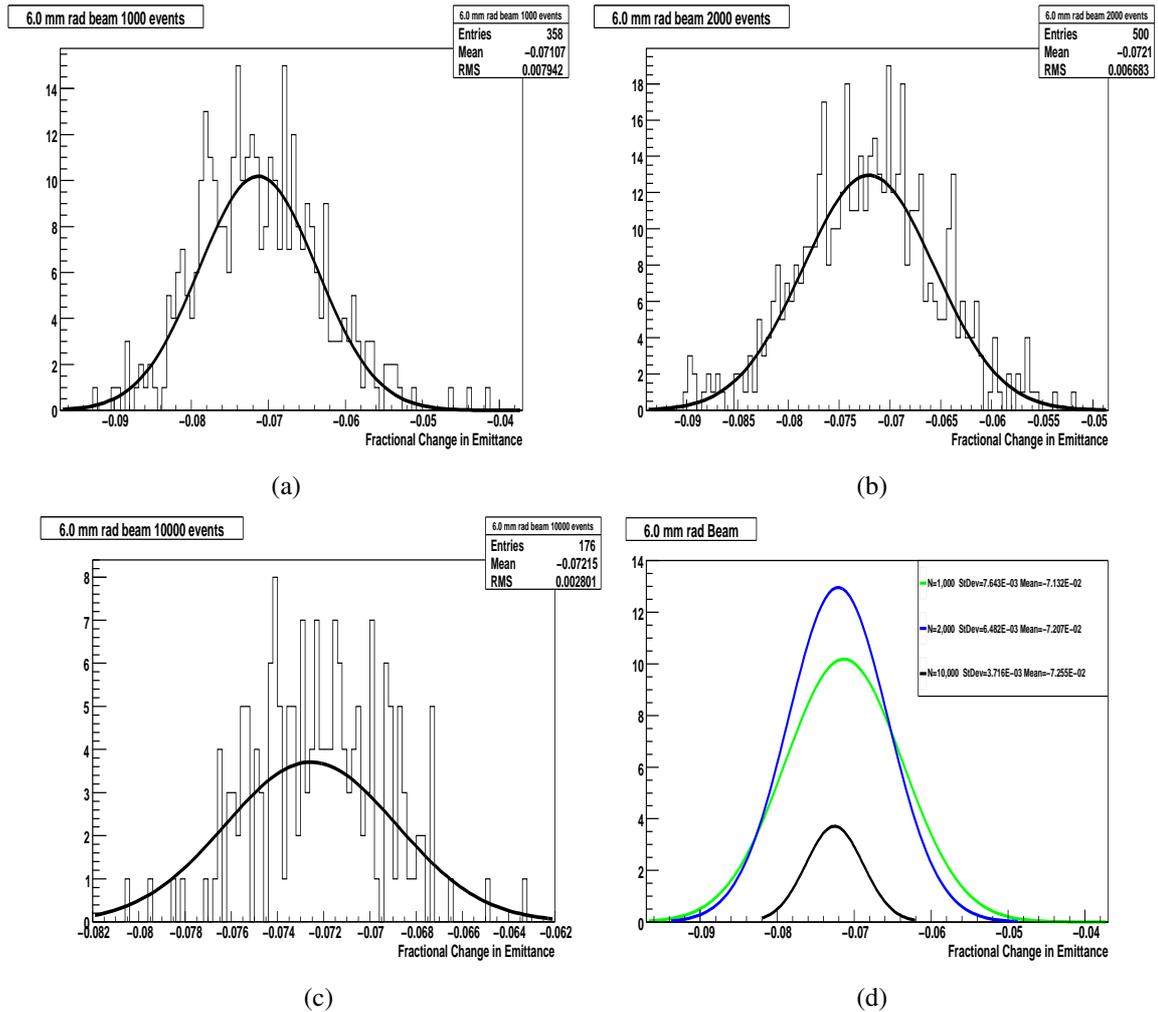


Figure 5.9: Distributions of fractional change in emittance for a 6.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

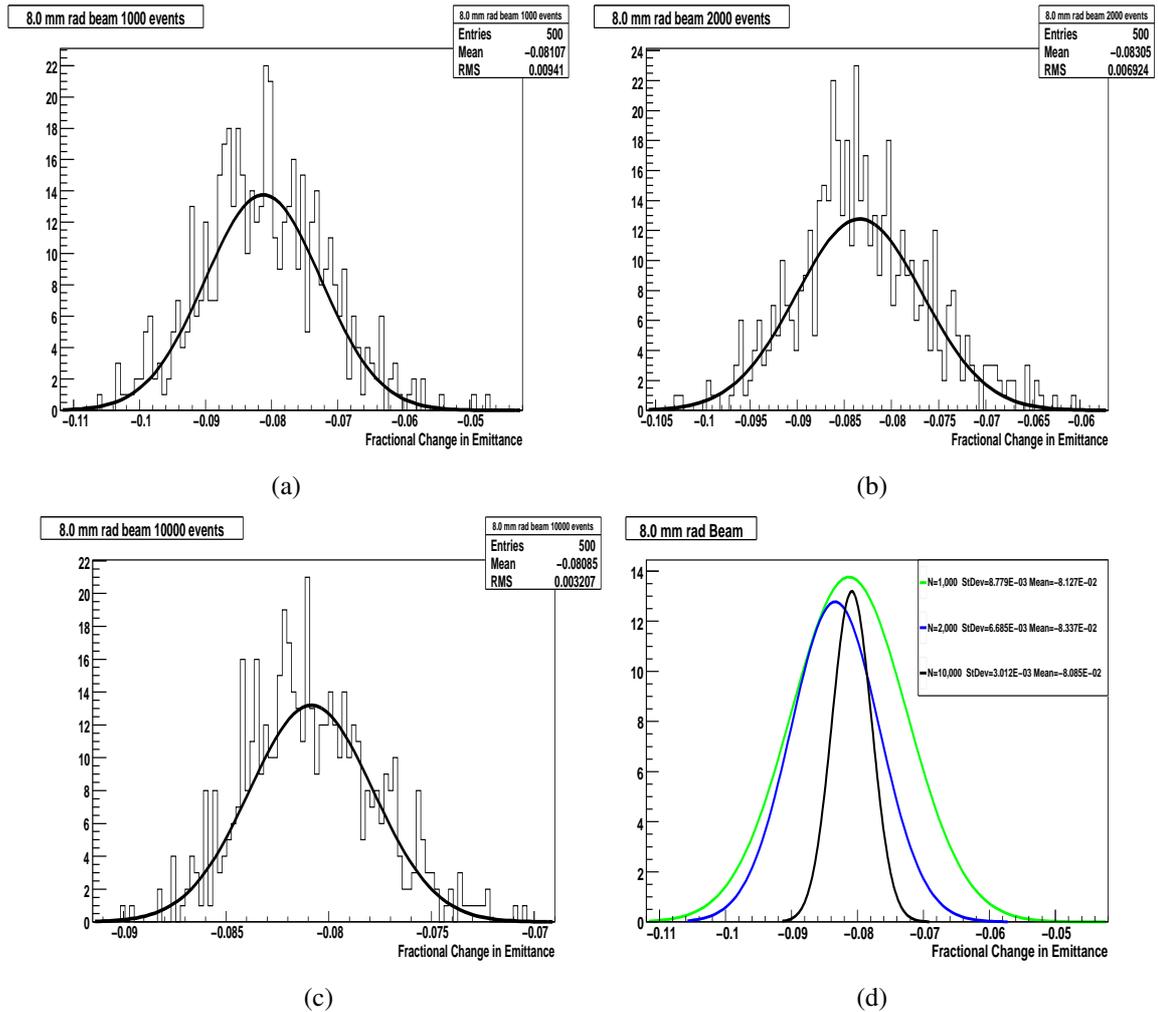


Figure 5.10: Distributions of fractional change in emittance for a 8.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

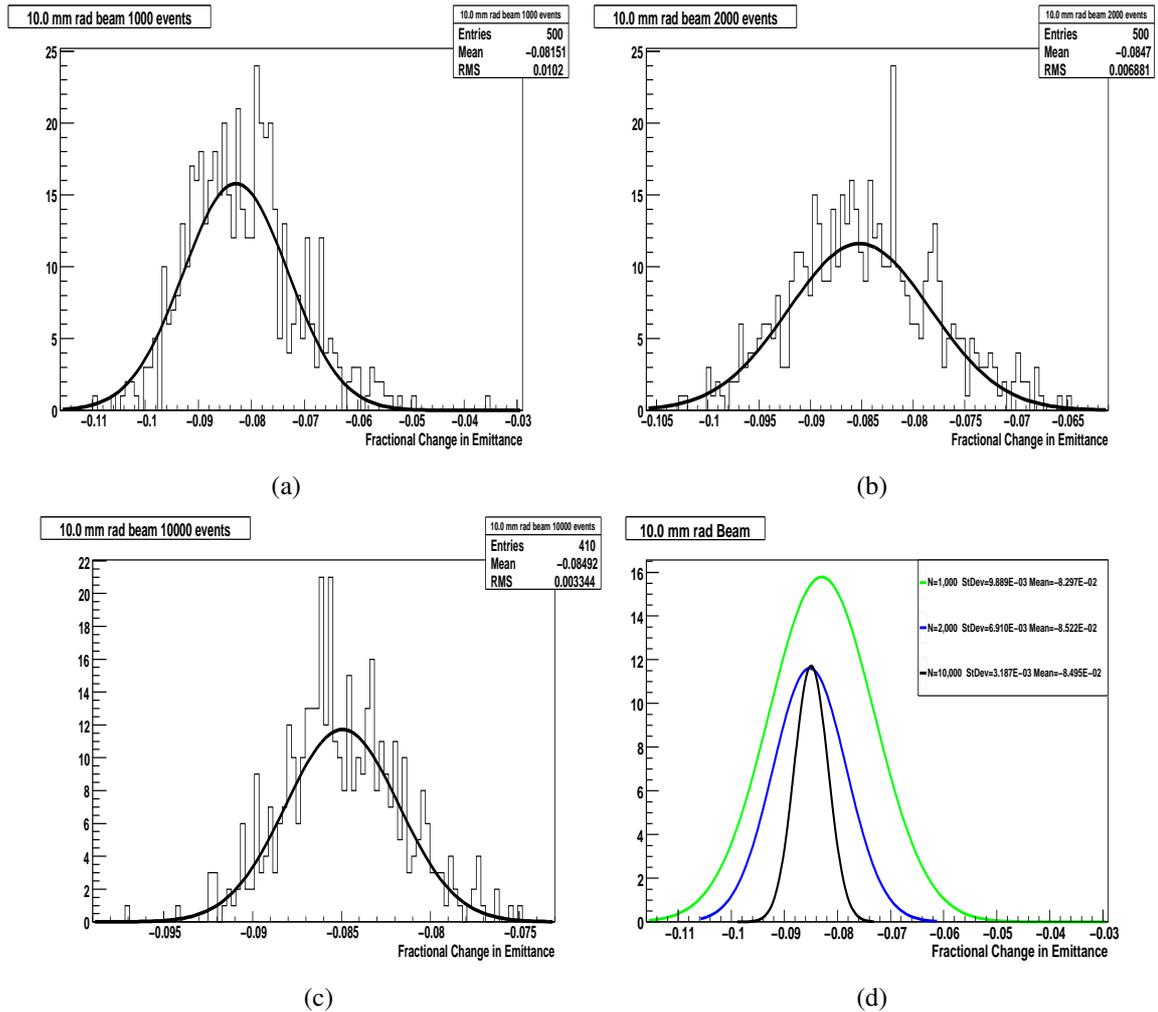


Figure 5.11: Distributions of fractional change in emittance for a 10.0 mm rad beam for (a) 1,000 (b) 2,000 and (c) 10,000 events (muons) per simulation. (d) shows an overlay of the fits for each distribution with the standard deviation narrowing as the number of events per simulation increases.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF EMITTANCE CHANGE MEASUREMENTS IN MICE

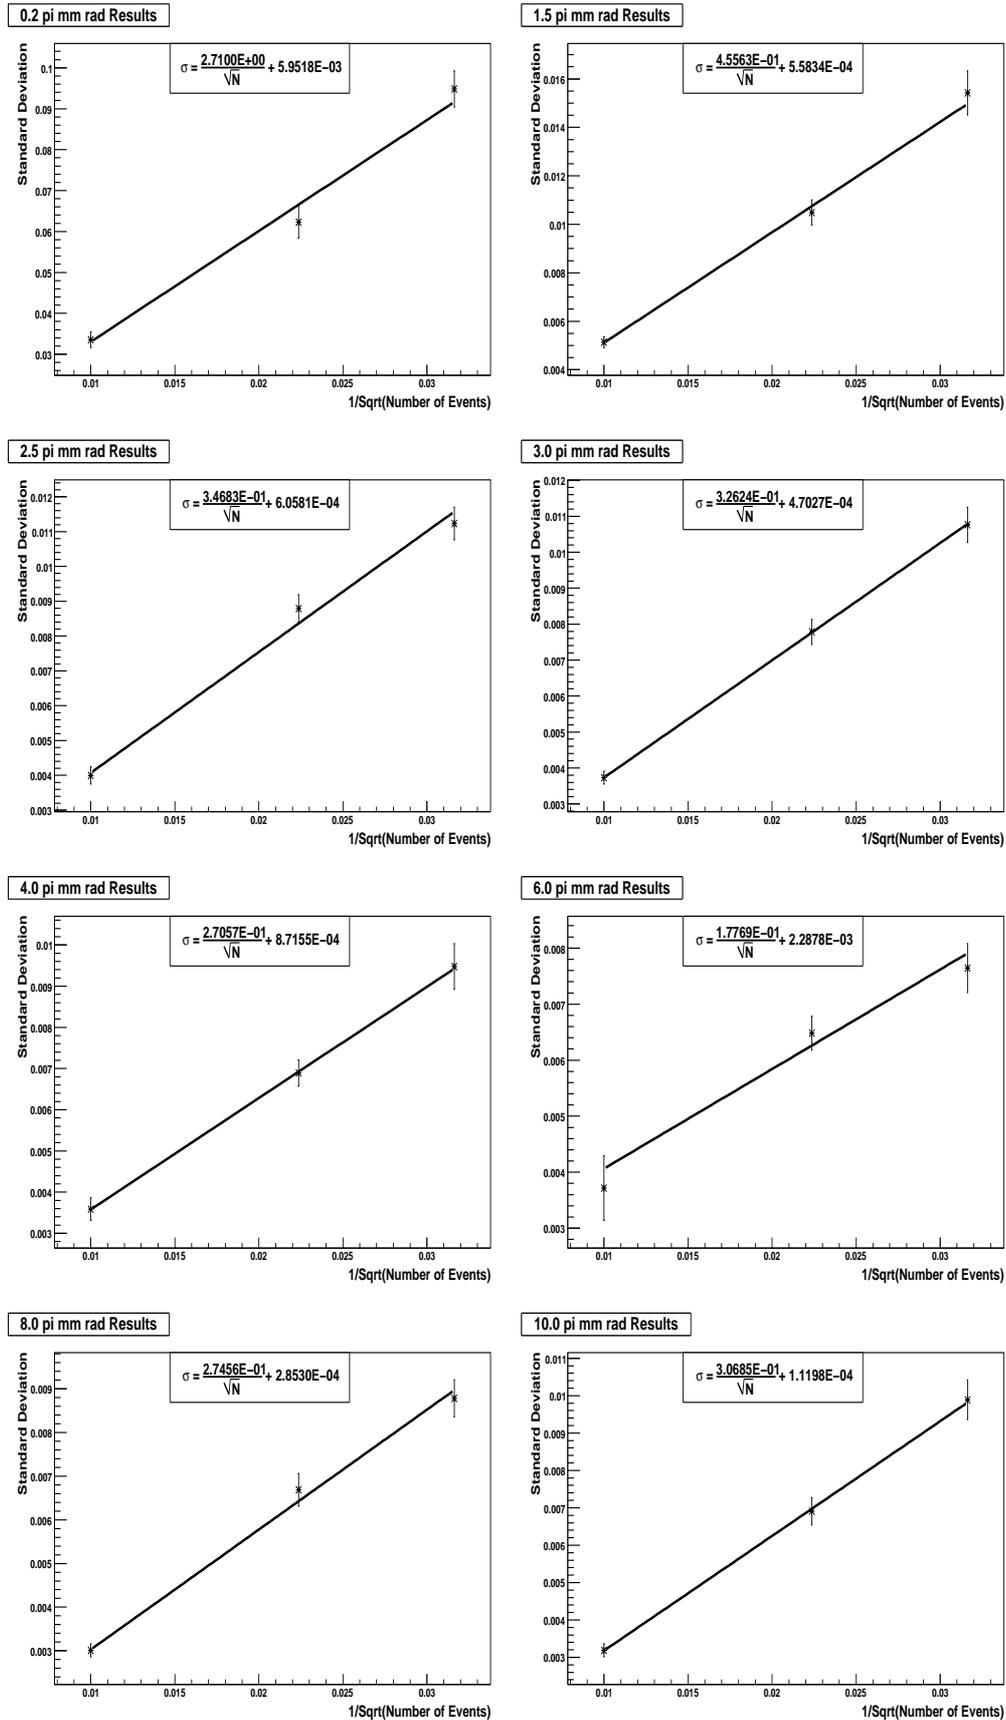


Figure 5.12: For each beam, we plot σ_f (taken as σ in the plots) against $\frac{1}{\sqrt{N}}$ to find a constant of proportionality, K .

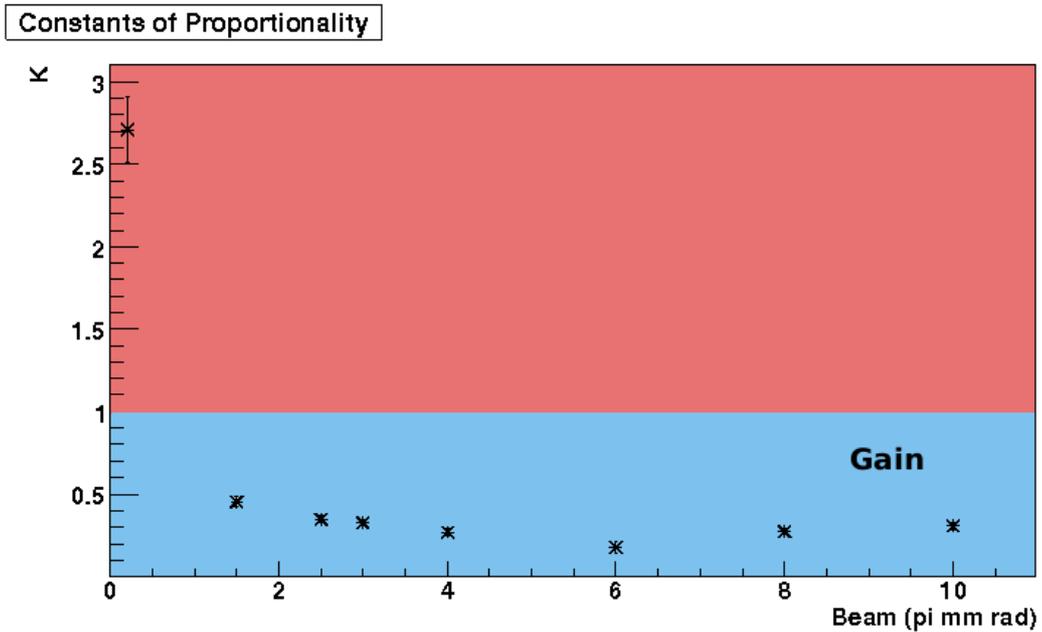


Figure 5.13: The constant of proportionality K between σ and $\frac{1}{\sqrt{N}}$ as a function of input beam size. The region of $K < 1$, where we make gains due to the effect of correlations, is highlighted in blue.

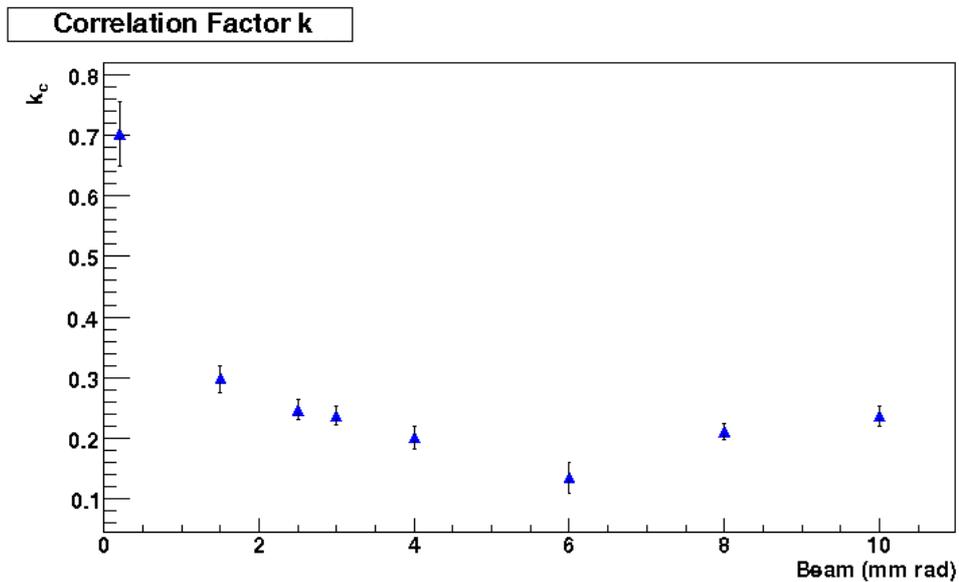


Figure 5.14: The correlation factor k_c due to measuring a similar sample of muons in two places. This correlation factor is included in the constant of proportionality K .

Recalling equation 5.34:

$$\sigma_{sim} = k_c \sqrt{\frac{2R}{N}} \quad (5.41)$$

$$\sigma_{sim} = k_c \sqrt{2(1+f)} \frac{1}{\sqrt{N}}. \quad (5.42)$$

It follows that:

$$K = k_c \sqrt{2(1+f)} \quad (5.43)$$

For most beams where cooling is present and f is positive, the factor $\sqrt{2(1+f)}$ will exceed unity. The effect of the correlation factor k_c , which is plotted in figure 5.14 as a function of input beam, is to reduce the total value of K by an order of magnitude. Since K is around 0.3, it means that MICE is required to run with 10 times less muons than if the constant was equal to 1 (i.e. with uncorrelated errors).

The stated MICE goal is to measure fractional change in emittance with an error of less than 1%. Since the total error is influenced by systematic errors also, it is desirable for the statistical error to contribute no more than 10% to the total error, calculated by summing all errors in quadrature. This requirement allows us to calculate the number of particles required to meet the goals of MICE for a range of beam conditions. The quantification of the effect of correlations upon the measurement means that an order of magnitude less particles are required than if the effect of correlations was not understood. As MICE is a single particle experiment, this result could have the effect of reducing the necessary running time significantly.

The number of muons required to achieve a measurement error of 0.1% will be inferred from our results for the constant of proportionality K . K has been shown to be less than unity for a range of realistic beam conditions due to the effect of correlations in position and momentum measurements at the two scintillating fibre trackers.

By requiring that statistical error should contribute only 10% of the total error, that is, a 0.1% absolute error, we calculate the number of muons directly from equation 5.1:

$$N = \frac{K^2}{0.001^2} \quad (5.44)$$

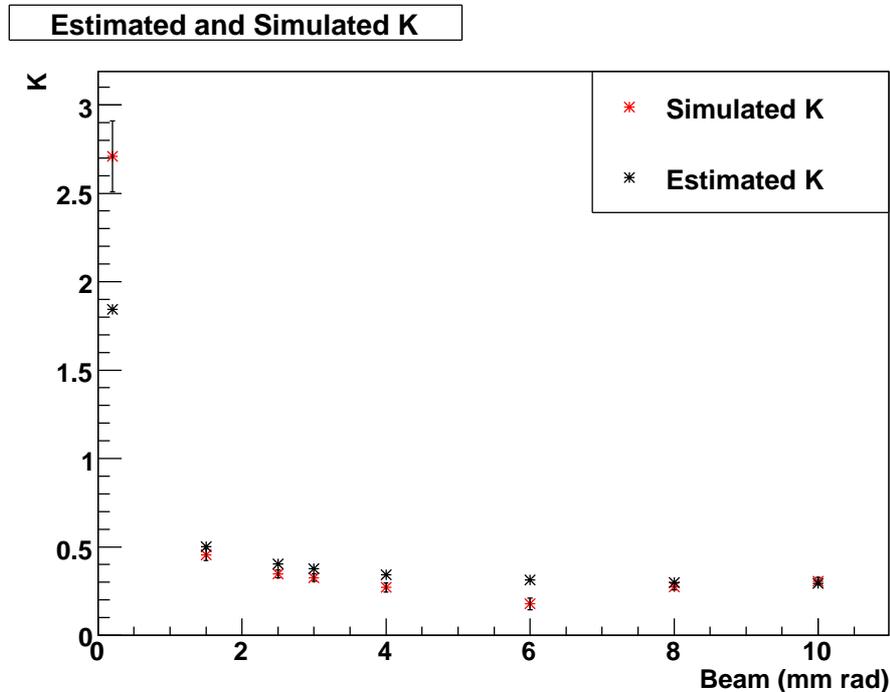


Figure 5.15: The constant of proportionality relating the simulated error in fractional change in emittance to the inverse square root of the number of events is plotted and compared to an analytical solution provided by Cobb[160]. Agreement is good except in the case of the very small 0.2 mm rad beam.

Table 5.1 presents the number of muons necessary to meet the goals of MICE for each beam considered in this study. The case of the 0.2 mm rad is of little experimental interest since it is not practical. For higher emittance beams of 4 mm rad and above, due to the effect of correlations, an order of magnitude less muons are required than might have otherwise been assumed given the uncorrelated case, greatly reducing the necessary running time of the experiment.

Figure 5.15 compares the simulated study with the analytical solution, presented in section 5.3.1, taking the fractional change in emittance from the results of the simulations. There is good agreement, except in the case of the pencil beam. This is likely due to deficiencies in the beam matching in the simulation combined with stronger scattering effects for the lowest emittance beams. A full comparison is provided in table 5.2.

Beam (π mm rad)	K	δK	With Correlations (10^5 Muons)	Without Correlations (10^6 Muons)
0.2	2.710	0.020	74	15
1.5	0.456	0.033	2.1	2.3
2.5	0.347	0.023	1.2	2.0
3.0	0.326	0.020	1.1	1.9
4.0	0.270	0.025	0.8	1.8
6.0	0.178	0.033	0.4	1.7
8.0	0.274	0.018	0.8	1.7
10.0	0.307	0.021	1.0	1.7

Table 5.1: Estimated number of muons required with and without correlations for a range of beam conditions. Correlations in our measured sample of muons have the effect of reducing the total number of muons required by an order of magnitude.

Input Emittance (mm rad)	f	α	K (simulated)	$K = \sqrt{\frac{1}{2}((1-f)^2 - \alpha^2)}$
0.2	1.68	0.0162	2.71	1.84
1.5	0.0860	0.0195	0.456	0.501
2.5	-0.00999	0.0200	0.347	0.403
3.0	-0.0311	0.0203	0.326	0.376
4.0	-0.0483	0.0213	0.271	0.343
6.0	-0.0734	0.0246	0.178	0.313
8.0	-0.0791	0.0290	0.275	0.299
10.0	-0.0877	0.0331	0.307	0.293

Table 5.2: Comparison between simulated and analytical results.

5.4 MICE Tracker Alignment Tolerances

5.4.1 Motivation of Alignment Study

In the light of the results presented in section 5.3.2.5, it may be assumed that MICE will run with enough events to ensure that the statistical errors on fractional emittance change are small. The total error will therefore be dominated by systematics. One crucial source of systematic error to quantify is the alignment of the scintillating fibre trackers.

For the purposes of this study [166] it is assumed that in order to meet the requirements of MICE, misalignment of the trackers should contribute no more than 10% of the total error. That is, an error of 0.1% on fractional change in emittance to be added in quadrature with other sources of error.

At the time of writing, both trackers are yet to be installed in the MICE channel, due to ongoing issues surrounding the solenoid magnets. Engineering work must be completed on the solenoids before the full assembly of either module. In the absence of data, simulation allows us to compare directly the misaligned case and the case where the trackers are aligned perfectly. During experimental running, alignment will be calculated using reconstructed tracks, to improve the original mechanical alignment.

5.4.2 Simulation

Let us assume that the z -axis is aligned along the direction of the beam. Twelve misalignment settings were studied: ± 10 mm translation along the x axis, ± 10 mm translation along the y axis, ± 3 mrad rotation in the x - z plane, ± 3 mrad rotation in the y - z plane, the combination of both ± 10 mm translation in x and ± 3 mrad rotation in the x - z plane and the combination of both ± 10 mm translation in y and ± 3 mrad rotation in the y - z plane. These were applied to the downstream tracker only. Since we are interested in measuring the emittance change before and after the cooling channel, it is the relative misalignment between the two trackers that is of interest. Therefore, it is only necessary to misalign the downstream tracker to determine the effect upon the emittance change measurement.

Monte Carlo simulations of beams with emittance between 0.2 mm rad and 10 mm rad were run using the final configuration of MICE shown in figure 3.4. For each beam, twelve separate simulations were generated using each of the misalignment settings in addition to a correctly aligned simulation for comparison purposes. Each simulation featured 10,000 muon tracks passing through MICE and was generated using the same generated input beam files and G4MICE release (1.9.5) as the statistics study detailed previously in this chapter (section 5.3).

Every simulation was given identical initial pseudo-random conditions in order to separate the error due to misalignment from the statistical error. The only cut applied was the exclusion of particles which did not reach the second tracker.

5.4.3 Results

For each combination of beam and misalignment setting, the fractional change in emittance was compared to the non-misaligned case, with examples given in figures 5.16-5.18 for 2.5, 6.0 and 8.0 mm rad emittance beams. There are five points plotted in each graph, corresponding to the centre of the first tracker, the centres of the three absorbers and the centre of the second tracker. For computational reasons it is not feasible to present near-continuous graphs of emittance of the form shown in figure 5.1 however the same step-like function should be assumed.

Each graph contains the correctly aligned case and the cases of both positive and negative misalignment in some. The positive and negative misalignment both cause a different fractional change relative to the correctly aligned case. This difference is a source of systematic error. We take the maximum difference, an absolute emittance, and convert it to a percentage of the correct emittance change. These percentages are summarised in table 5.3.

The total error must be 1% or less. Since the total error is influenced by statistical and other systematic errors also, it is desirable for the misalignment error to contribute no more than 0.1% error on fractional change in emittance. The results in table 5.3 range from 0.01 to 0.81%. Although the simulations are based on the same pseudo-random seed, some statistical fluctuation exists due to small differences in transmission for each misalignment setting. To provide a general and robust alignment tolerance we therefore divide these misalignments by a factor of ten. This produces a tolerance requirement of 1 mm in x or y translation and 0.30 mrad in rotation. These figures suggest that the minimisation of the alignment error is achievable.

5.5 Conclusions

We have stated that the main goal of MICE is to measure a fractional drop in emittance of order 10% to an accuracy of 1% for a range of beams. An understanding of both the statistical errors and the alignment tolerances involved in this measurement is essential to

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

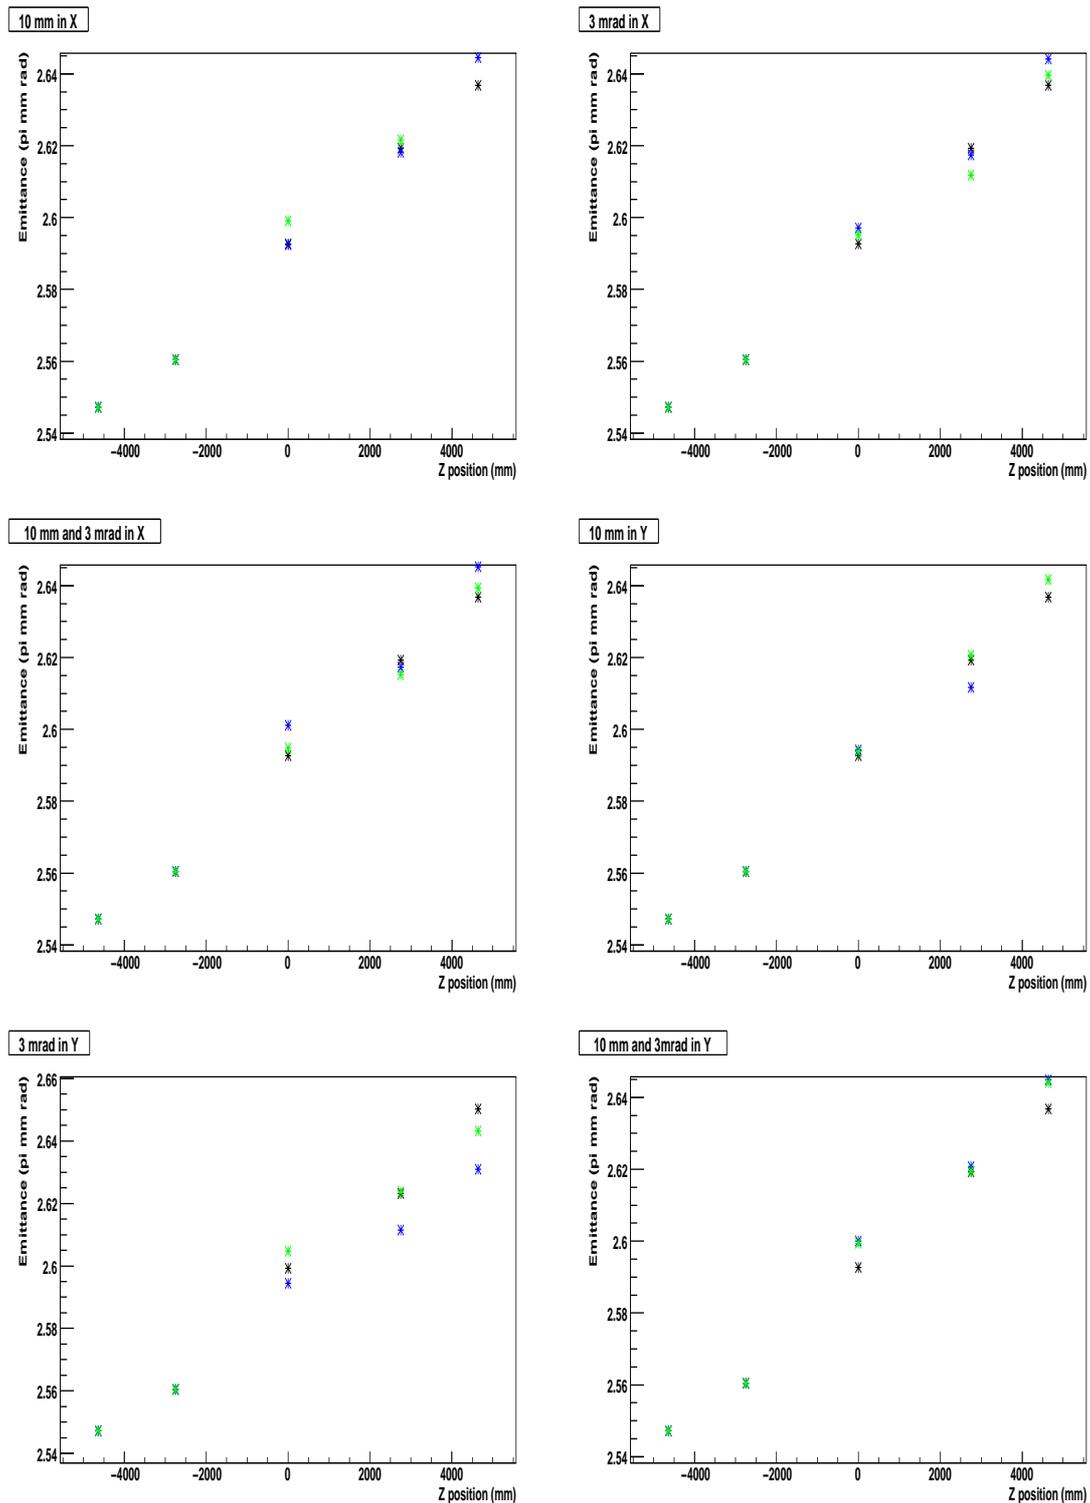


Figure 5.16: Plots shown apply to misalignments using the 2.5 mm rad beam. The input beam is sufficiently small that multiple scattering dominates and the beam experiences fractional increase in emittance, or ‘heating’. Green markers indicate positive misalignment, blue markers indicate negative misalignment and black markers indicate nominal emittance with no misalignment.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

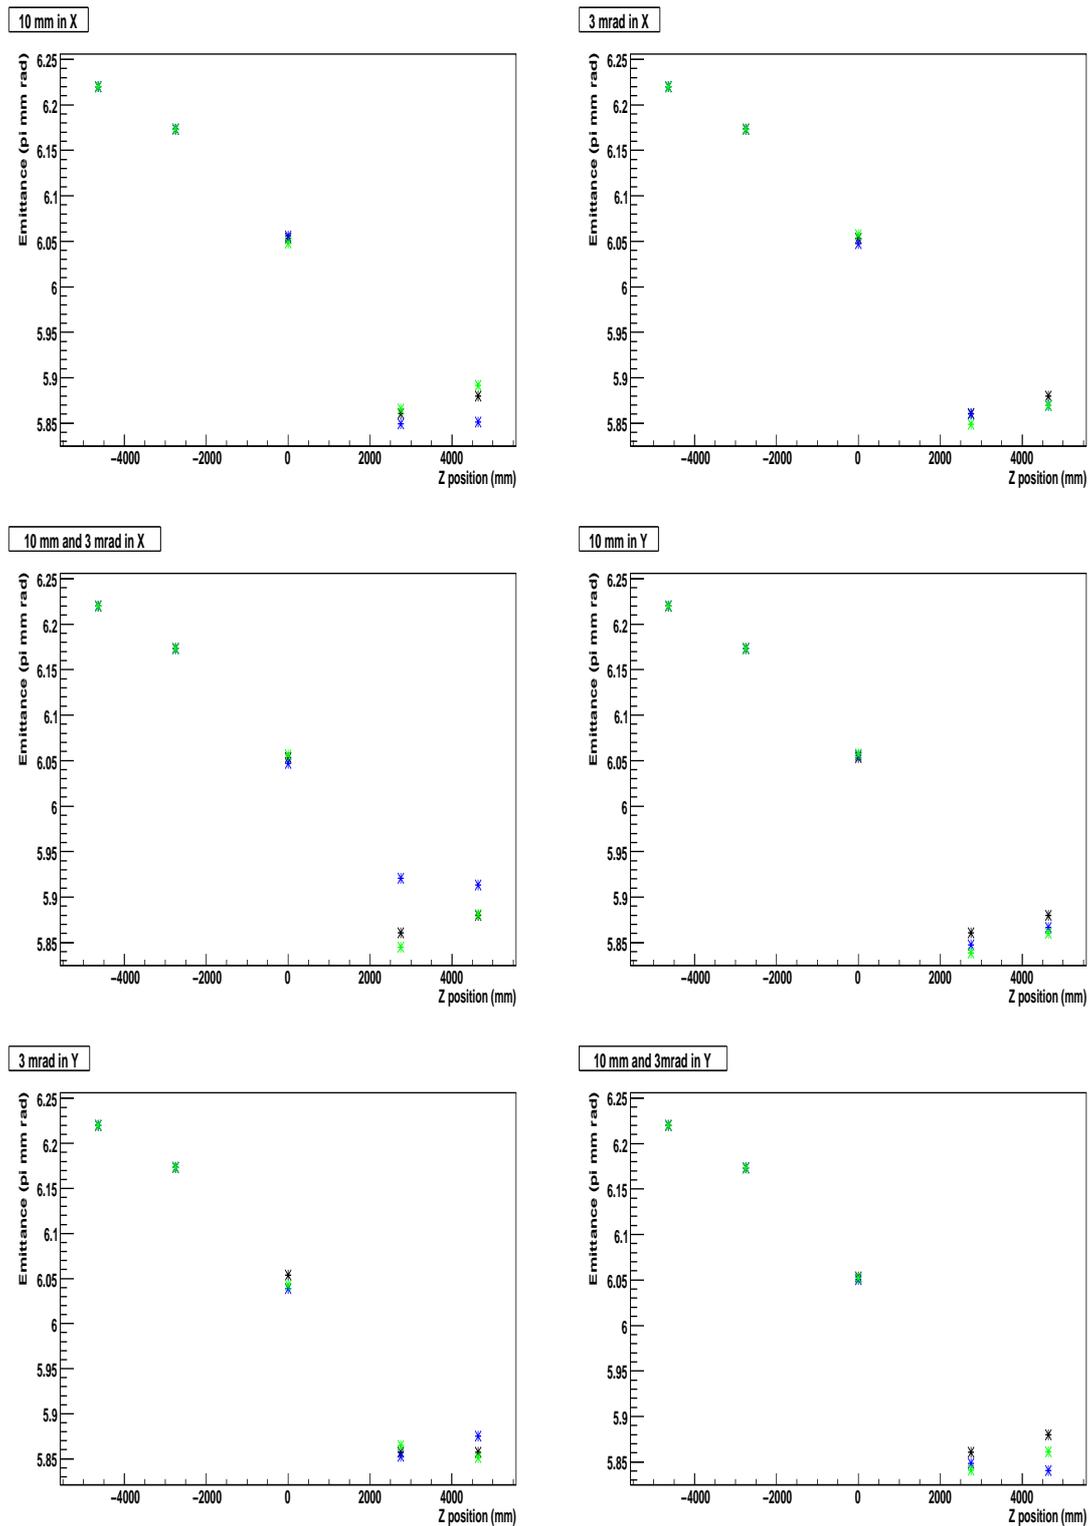


Figure 5.17: Plots shown apply to misalignments using the 6 mm rad beam. Green markers indicate positive misalignment, blue markers indicate negative misalignment and black markers indicate nominal emittance with no misalignment. In this beam we see significant cooling.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

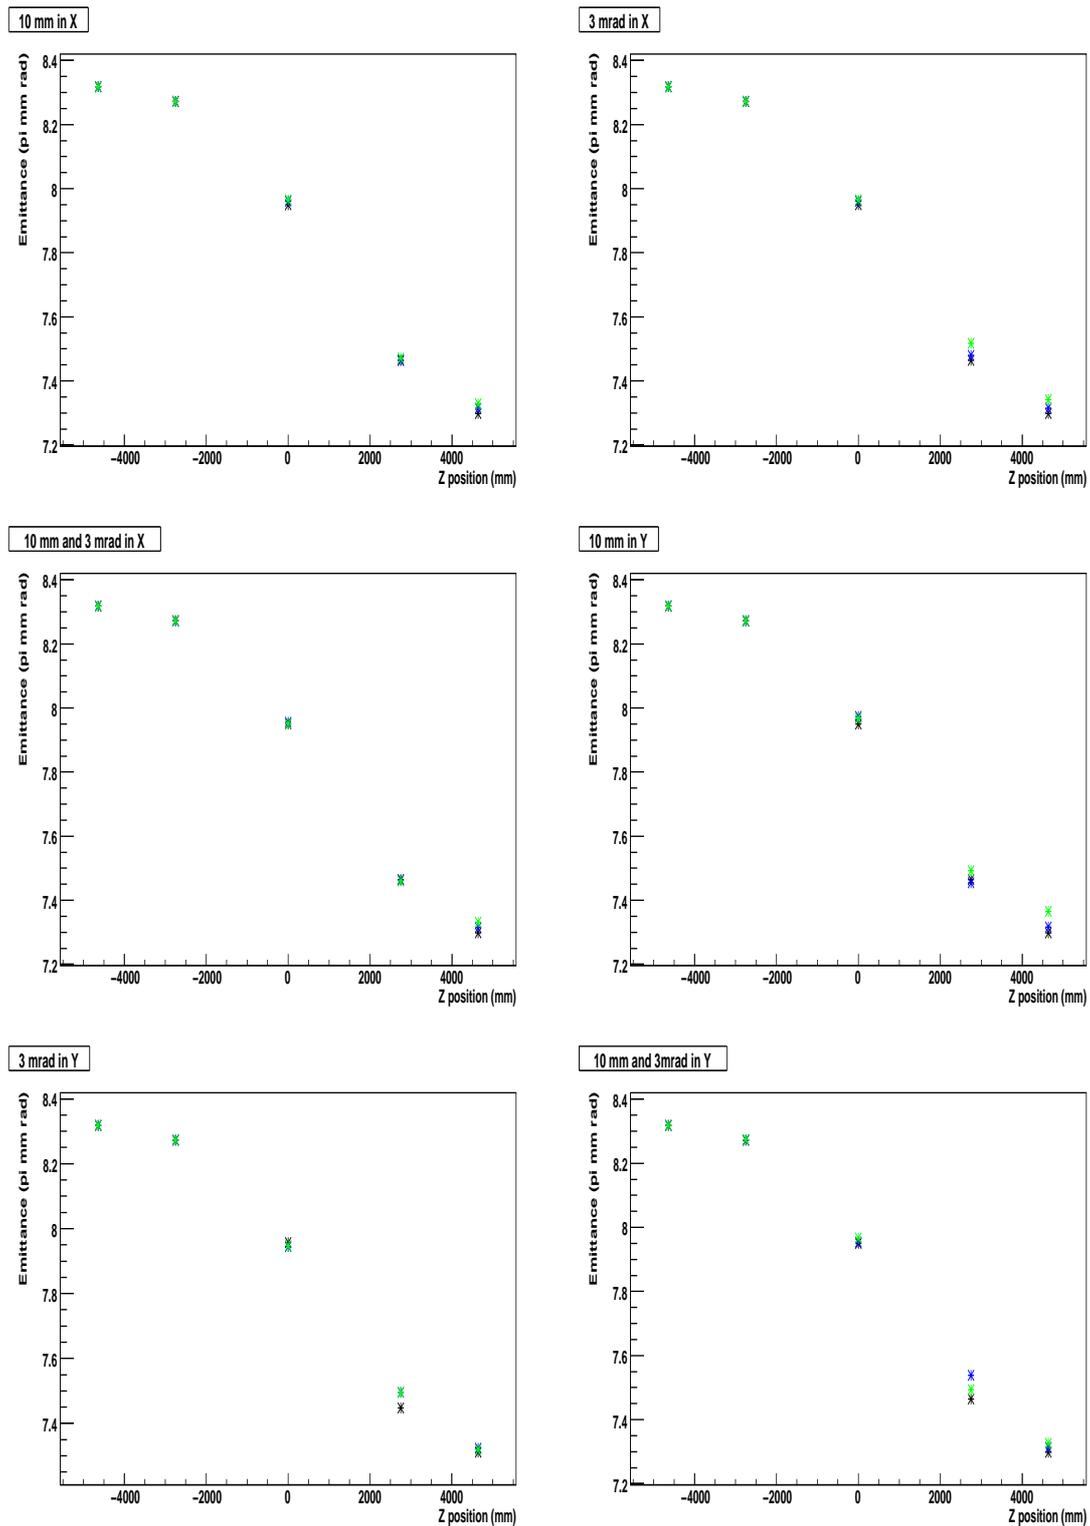


Figure 5.18: Plots shown apply to misalignments using the 8 mm rad beam. Green markers indicate positive misalignment, blue markers indicate negative misalignment and black markers indicate nominal emittance with no misalignment. In this beam we see significant cooling.

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF
EMITTANCE CHANGE MEASUREMENTS IN MICE

Misalignment	Emittance (mm rad)						
	1.5	2.5	3.0	4.0	6.0	8.0	10.0
x	Change caused in f (%)						
10 mm	0.144	0.499	0.126	0.253	0.454	0.374	0.058
3 mrad	0.012	0.110	0.120	0.277	0.164	0.527	0.176
10 mm 3 mrad	0.170	0.099	0.046	0.540	0.016	0.406	0.285
y							
10 mm	0.406	0.191	0.284	0.192	0.311	0.808	0.271
3 mrad	0.577	0.757	0.223	0.274	0.099	0.097	0.602
10 mm 3 mrad	0.374	0.288	0.358	0.511	0.632	0.331	0.531

Table 5.3: Percentage changes in emittance due to misalignment of the downstream scintillating fibre tracker.

any assessment of its feasibility. In both regards, this study has illustrated that the goal of MICE is achievable.

The results of the statistics study demonstrate that a constant of proportionality K exists between the standard deviation of fractional emittance change and the inverse square root of the number of events. In the absence of correlations, one can assume a constant of proportionality of unity. However, due to the effect of correlations, the constant of proportionality is reduced to approximately 0.3 in the case of larger, more useful beams. It has been shown that an order of magnitude less muons are required than in the case where $K = 1$ to achieve the statistical uncertainty goal of emittance measurement in MICE. MICE is a single particle experiment; this therefore has a direct and significant effect on the necessary running time for the experiment.

More generally, this study has been the first MICE related study which has been run on the Grid, providing useful learning artifacts for the collaboration as a whole [161]. It should be noted that without the use of the Grid or high performance computing to run tens of thousands of simulations, such a statistical analysis would not have been possible.

Such technology was also applied in quantifying the alignment tolerances of the scintillating fibre trackers positioned before and after the MICE cooling channel. In order to achieve the required precision, tracker misalignment should contribute an error of no greater than 0.1% on fractional emittance change, to be combined with other errors including statistical errors into the total error through addition in quadrature. Given this requirement the

CHAPTER 5. STATISTICAL ERRORS AND ALIGNMENT TOLERANCES OF EMITTANCE CHANGE MEASUREMENTS IN MICE

misalignment study has recommended a tolerance of 1 mm in translation and 0.30 mrad in rotation which suggests that MICE can meet its requirements through tight mechanical specifications of the tracker installation and careful alignment guided by software carrying out an analysis of reconstructed tracks.

MICE is currently running and taking data, evolving in a stepwise fashion to its final Step VI configuration. This study will prove useful in the final statement of the measurement of fractional change in emittance, a necessary step in the development of the cooling channel in a future Neutrino Factory facility.

Chapter 6

The MICE Luminosity Monitor

6.1 Purpose

We have described the MICE beamline in some detail (chapter 3) and thus far focused primarily on quantifying the alignment tolerances and statistical errors (chapter 5) associated with the measurement of fractional change in emittance over the length of the full MICE cooling channel. However, in order to characterise the beamline and analyse data produced from it, it is essential that we are able to both measure and predict particle rate.

The luminosity monitor (LM) provides a means of measuring particle rate close to the MICE target independently from the ionisation chamber beamloss monitors positioned around the ISIS synchrotron. A mapping between luminosity monitor rate and protons on target (POT) as a function of target depth affords normalisation for all detectors in the MICE beamline and cooling cell. This has clear applications to both beamline commissioning and data analysis. Additionally, the monitor allows us to assess the validity of simulation codes such as G4Beamline, within the limits of the accuracy of our model of the MICE geometry both in terms of particle production from the target and rate through the MICE channel, normalised to counts in the LM.

6.2 Design

The LM consists of two pairs of Hamamatsu H5783P photomultiplier tubes [167] attached to blocks of scintillator at either side of a slab of polyethylene, as shown in figure 6.1. Optical tape is used to ensure the system is light tight. The photo-multipliers feature low noise, 0.8 ns rise time, 10^6 gain and require an operating voltage of less than 15V, which is supplied through cables running from the MICE control room. The polyethylene block acts as a filter against low energy protons (below ~ 500 MeV/c) and pions (below ~ 150 MeV/c). Neutrons may be detected in the event that they are captured in the polyethylene, producing a proton.

Three scaler readout channels are provided to the DAQ system and recorded in the DAQ data files from the LM front end electronics as shown in figure 6.2. Readout is performed on a spill by spill basis, with spills coinciding with pulses of the target into the ISIS beam. These channels provide for three sets of coincidences on a spill, each counted by a scaler unit: PMTs 1 and 2 (Channel C12), 3 and 4 (Channel C34) and all four PMTs together (Channel C1234). Counting occurs within the experimental trigger gate. This data is available both for offline analysis and also for online monitoring from the control room. Singles rates are not recorded in the present configuration.

The design was created by Dr Paul Soler at the University of Glasgow. Installation has been completed with the detector now positioned in the ISIS vault at Rutherford Lab, as shown in figure 6.3. A distance of 10 m was chosen in order to ensure that the particle rate would be manageable within the limits of the vault geometry. We compare LM readings to beamloss measurements from ISIS in section 6.4. We follow a convention of quoting beamloss as the integral of the output voltage from the beamloss monitors of ISIS over the gate time of the measurement (typically during the last 3 ms of the ISIS acceleration cycle) in units of Volt milliseconds (noting however that beamloss signals from ISIS are provided as negative values of Volt milliseconds).

Different LM coincidence gates were studied and the lowest practicable width was found to be 10 ns, which is now the default and has been used in all data analysis presented here. During Step I, the detector has operated successfully without external shielding for beamloss

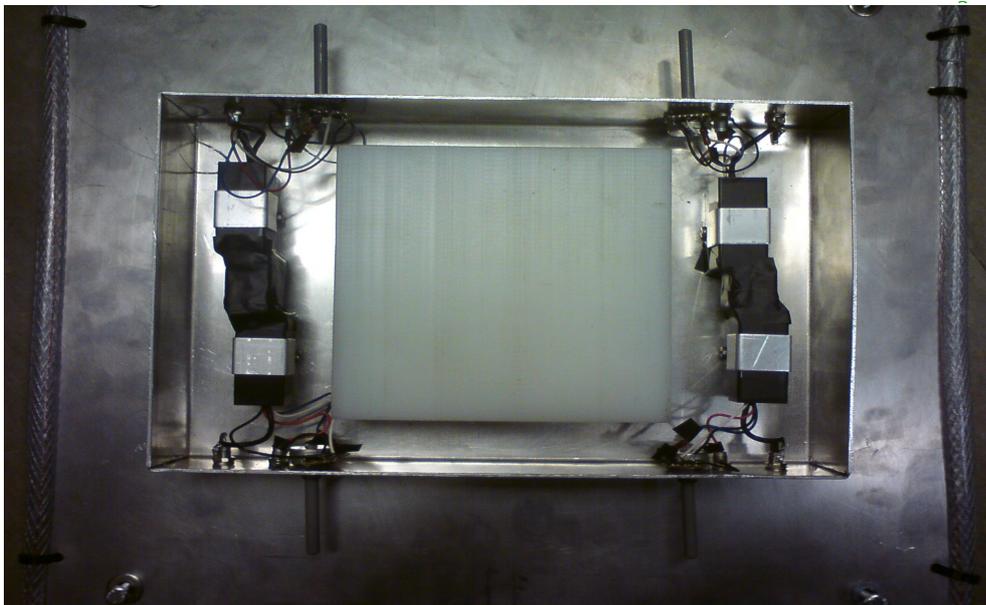
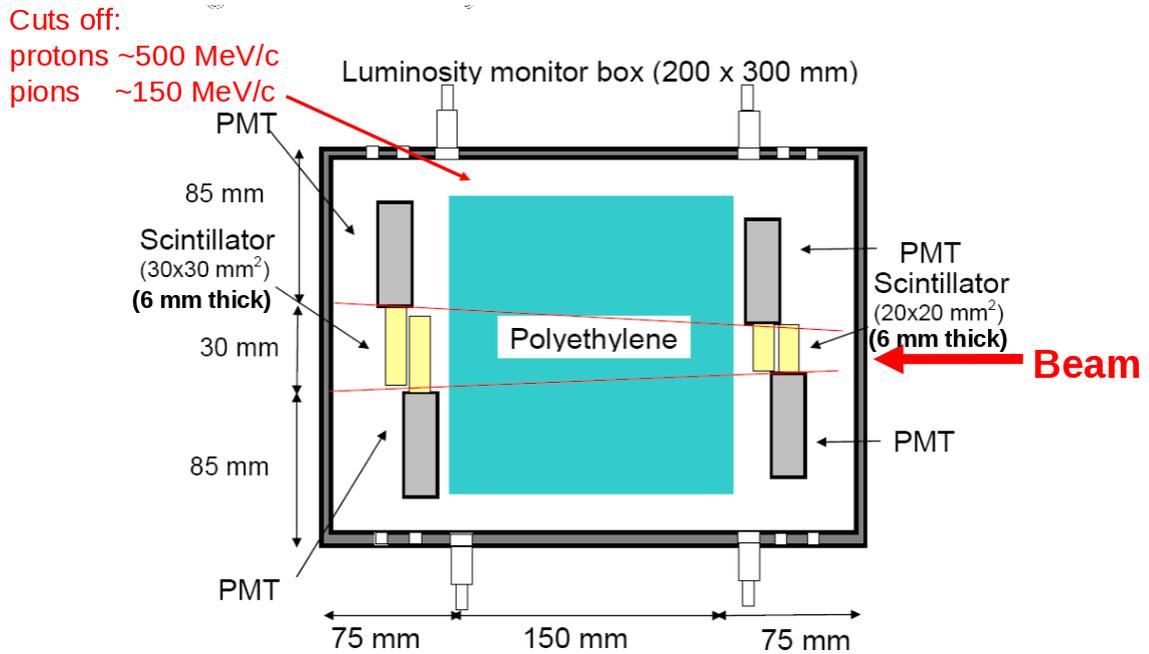


Figure 6.1: The design of the luminosity monitor features four low noise photo multiplier tubes in sets of two. We measure counts from coincidences of each pair of PMTs, and all four together. A block of polyethylene provides a filter for protons below 500 MeV/c and pions below 150 MeV/c.

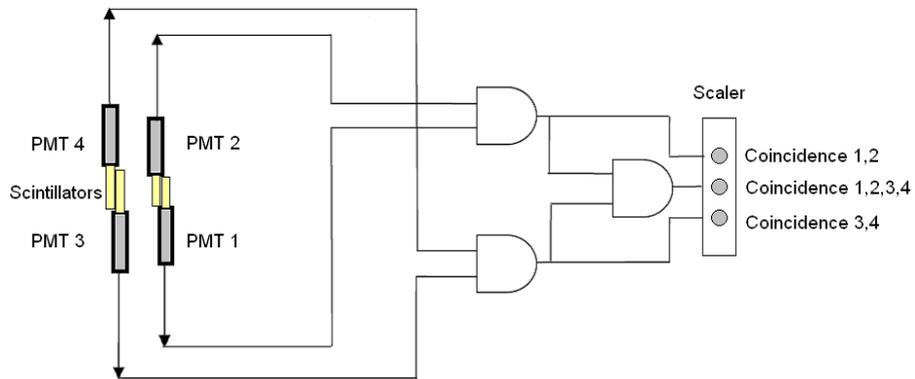


Figure 6.2: Scaler readout channels available to the luminosity monitor. The data from each run includes scaler counts of the coincidences of PMT sets 1 and 2 (C12), 3 and 4 (C34) and the fourfold coincidence of PMTs 1,2,3 and 4 (C1234).

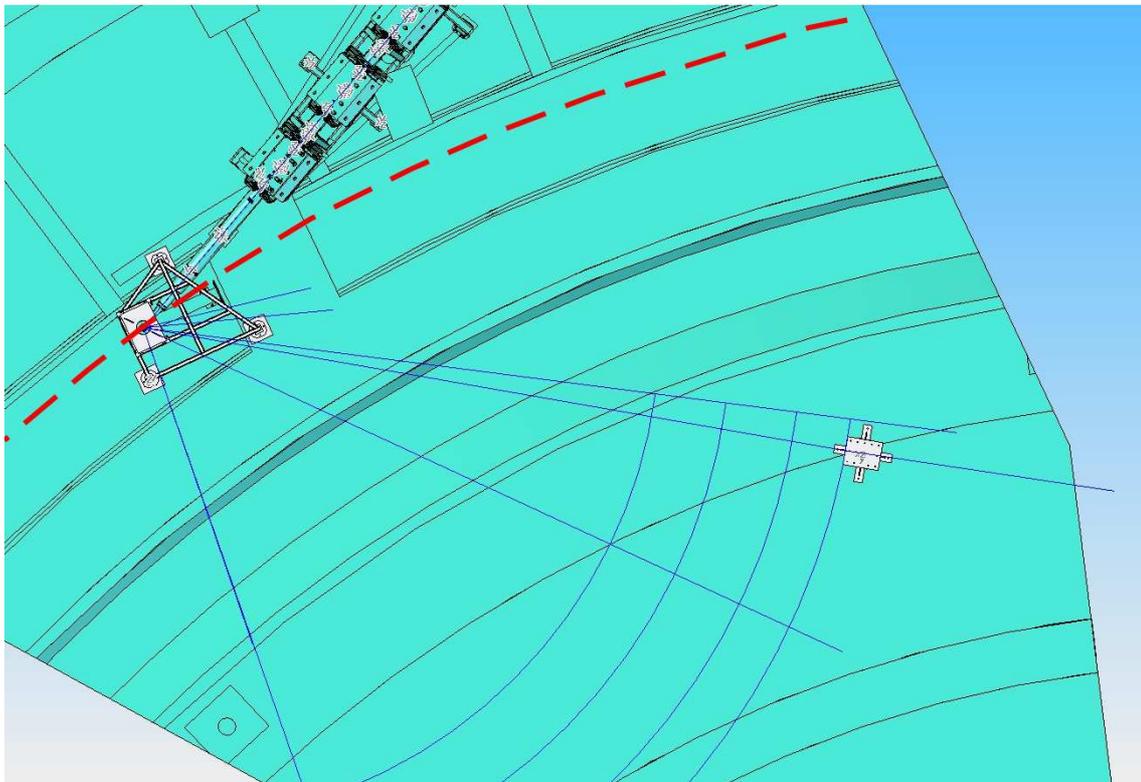


Figure 6.3: The luminosity monitor is positioned ten metres from the target within the ISIS vault. The red line follows the path of the ISIS ring. Original figure provided courtesy of Geoff Barber (RAL).

between 0 and 4 Vms, showing signs of saturation beyond this region, as will be shown in section 6.6.

6.3 Beamloss

We have described in section 3.6 how the dipping of the MICE target produces beamloss in ISIS. ISIS measures beamloss through the use of a series of 3 and 4 metre argon gas ionisation tubes horizontally placed along the inside track of the synchrotron [168], approximately 2-3 m from the beam axis. A distribution of beamloss is shown in figure 6.4. An integral of the MICE target pulse over the trigger width of 3.2 ms (from 6.8 ms to 10 ms of the ISIS acceleration cycle) is calculated at the end of the ISIS spill. Full background subtraction does not occur outside the region where the target is dipping into the beam, resulting in poor estimates for low beamloss runs. However such runs are of little practical interest to MICE. It should be noted that although an integral method is employed, by convention the loss signals from ISIS are negative. Beamloss values have been provided by Adam Dobbs (Imperial College) and manually synchronised with MICE data as part of this study.

Pre-existing conversions (table 6.1) relating beamloss to protons on target provide a basis from which we interpret the response of the beamloss monitors, however the precision of these conversions is not accurately known. Calibrations provided are estimated, by ISIS staff, through comparison of measurements from all beamloss monitors when beamloss occurs over the whole ring to the reduction of current in the beam toroids [169]. We cannot know the calibration error associated with this process and have therefore estimated an error of 30% [170].

6.4 Commissioning

The luminosity monitor counts must be converted to protons on target. During commissioning (7th February 2010), the counts were plotted against beamloss in sector 7 of ISIS (where the MICE target is located) demonstrating a linear relationship as shown in figure 6.5. As

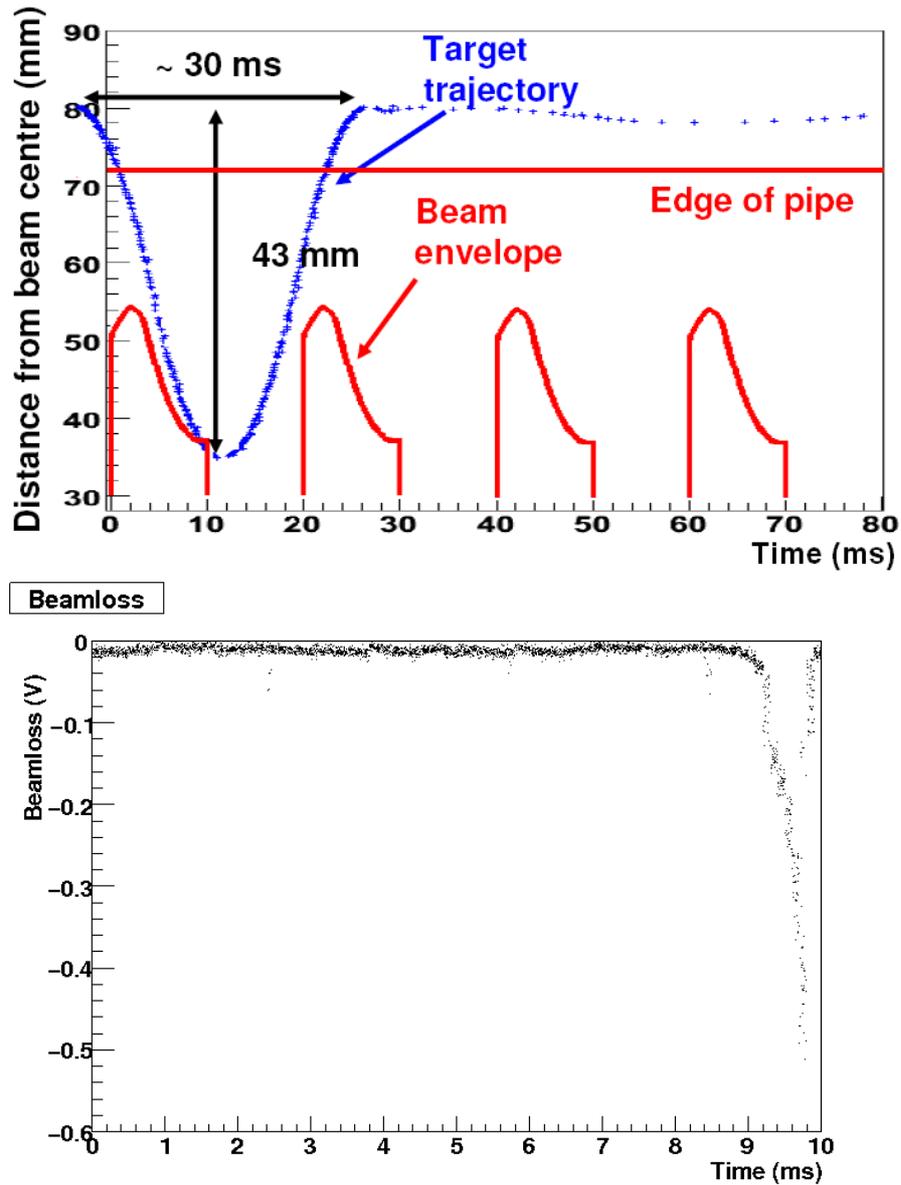


Figure 6.4: Top: Taken from [171]. The target intersects the ISIS beam at the last millisecond of its acceleration cycle. Bottom: Beamloss as a function of time due to target oscillations, measured by ionisation chambers in ISIS. The large peak is caused by the dipping of the MICE Target.

one would expect, beamloss and luminosity are proportional to one another.

The beamloss calibration at 9 ms from table 6.1 is used to convert beamloss to protons on target and by extension also interpret the response of the luminosity monitor, as shown in equation 6.1.

$$\frac{\text{LM Counts}}{\text{POT cm}^2} = \frac{\text{Gradient} \times \text{Sensitivity}}{\text{Area}}, \quad (6.1)$$

Time (ms)	Energy (MeV)	Sensitivity (Vs/POT)
0	70	2.2E-16
3	172	2.6E-16
5	374	4.3E-15
7	617	1.6E-14
9	780	3.5E-14
10	800	3.8E-14

Table 6.1: Estimates of sensitivity at different points in the ISIS machine cycle. The MICE target dips into the beam at around 9 milliseconds. Table provided by Dean Adams (STFC).

Coincidence	Gradient	Normalised Yield $\left(\frac{10^{-8}}{\text{pot} \times \text{cm}^2}\right)$
1, 2	1967.9 ± 2.0	1.72 ± 0.52
3, 4	2111.8 ± 2.0	0.82 ± 0.25
1, 2, 3, 4	899.50 ± 1.2	0.78 ± 0.24

Table 6.2: Gradients of fitted lines from the graphs contained in figure 6.5, which provide conversion factors between luminosity counts and beamloss. The region of low beamloss < 0.5 Vms is excluded from the fit due to imprecision in the beamloss calculation. Data was taken during experimental runs 1447-1457.

where Gradient is derived from figure 6.5, sensitivity is from table 6.1 and area is the area of the scintillators which is $(4 \text{ cm}^2$ for scintillators 1 and 2 and 9 cm^2 for scintillators 3 and 4; the C1234 channel assumes an area of 4 cm^2).

The gradients of each of the graphs in figure 6.5 were used to provide conversion factors between luminosity counts and beamloss and to calculate the response of each of the channels as detailed in table 6.2. Note that the conversion is not valid in the region of low beamloss (below 0.5 Vms), due to inherent inaccuracies in the method for calculating beamloss.

6.5 Simulations

6.5.1 Transmission

Two principal sets of simulations were produced with a view to accurately predicting the performance of the luminosity monitor. The first featured a cylindrical target, steel pipe and an approximation of the luminosity monitor structure. Here the luminosity monitor featured

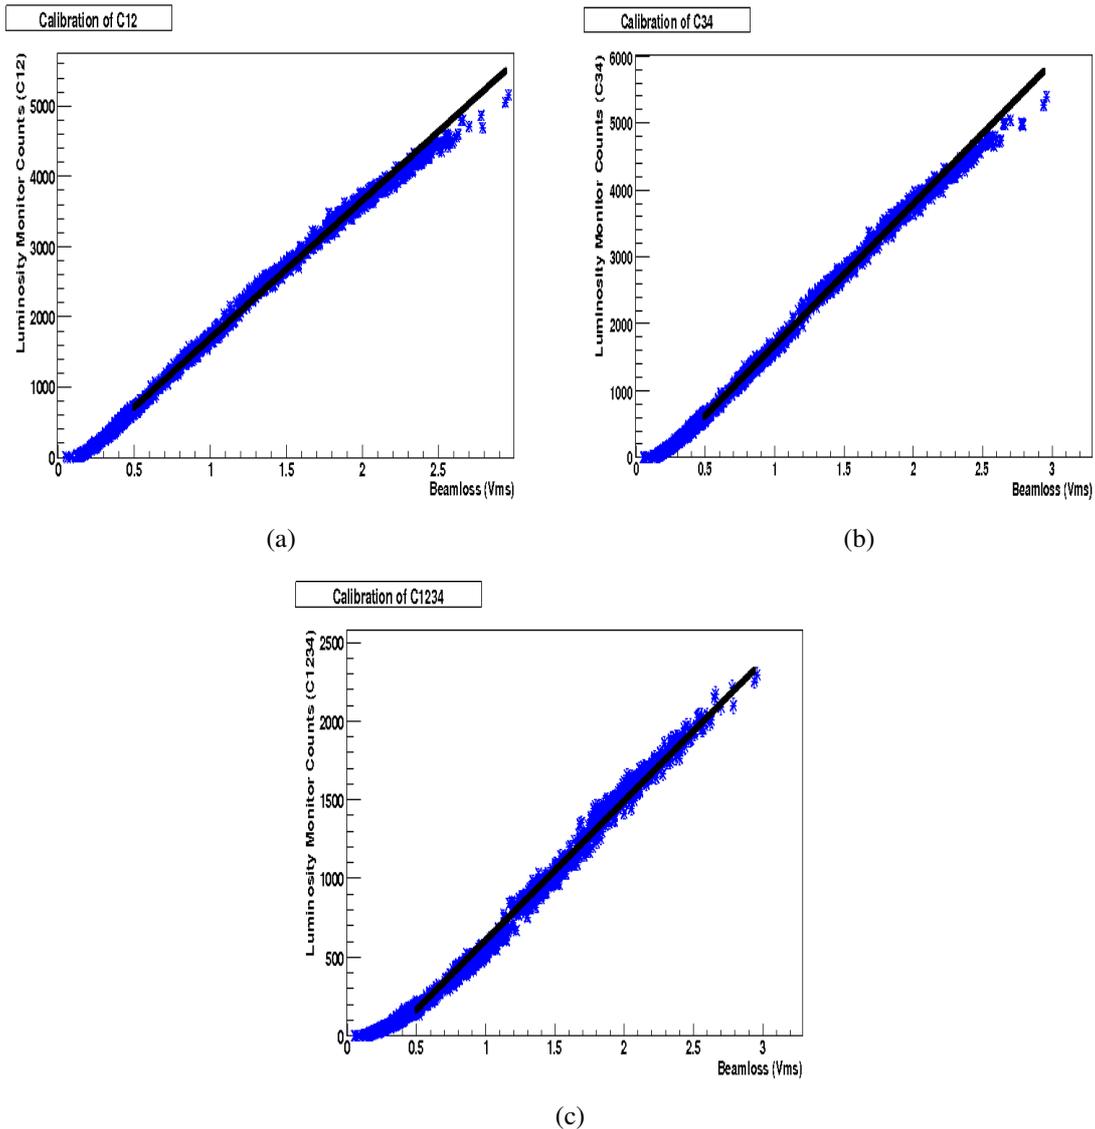


Figure 6.5: First data from the luminosity monitor showing a direct relationship between luminosity and beamloss on all three scaler output channels. Graphs show counts for coincidences of signals for PMTs (a) 1,2 (b) 3,4 and (c) 1,2,3 and 4. The straight lines indicate a linear fit to the data between 0.5 Vms and 3 Vms, since the method for calculating beamloss is unreliable below 0.5 Vms.

two pairs of scintillator detector planes at either side of the plastic slabs, placed 10 metres and 25 degrees from the target in accordance with figures 6.3 and 6.1.

Figure 6.6 shows the active simulation geometry. The target is contained within a (red) cylindrical steel casing with 6 mm thickness that mimics the beam pipe. The LM is shown left of the target and has a sandwich-like structure with a polyethylene block in the centre, braced by two pairs of scintillator slabs (shown in blue and pink), with 3 mm aluminium

sheets representing the LM casing at the outermost layer. The LM is positioned 10 m and 25 degrees from the target.

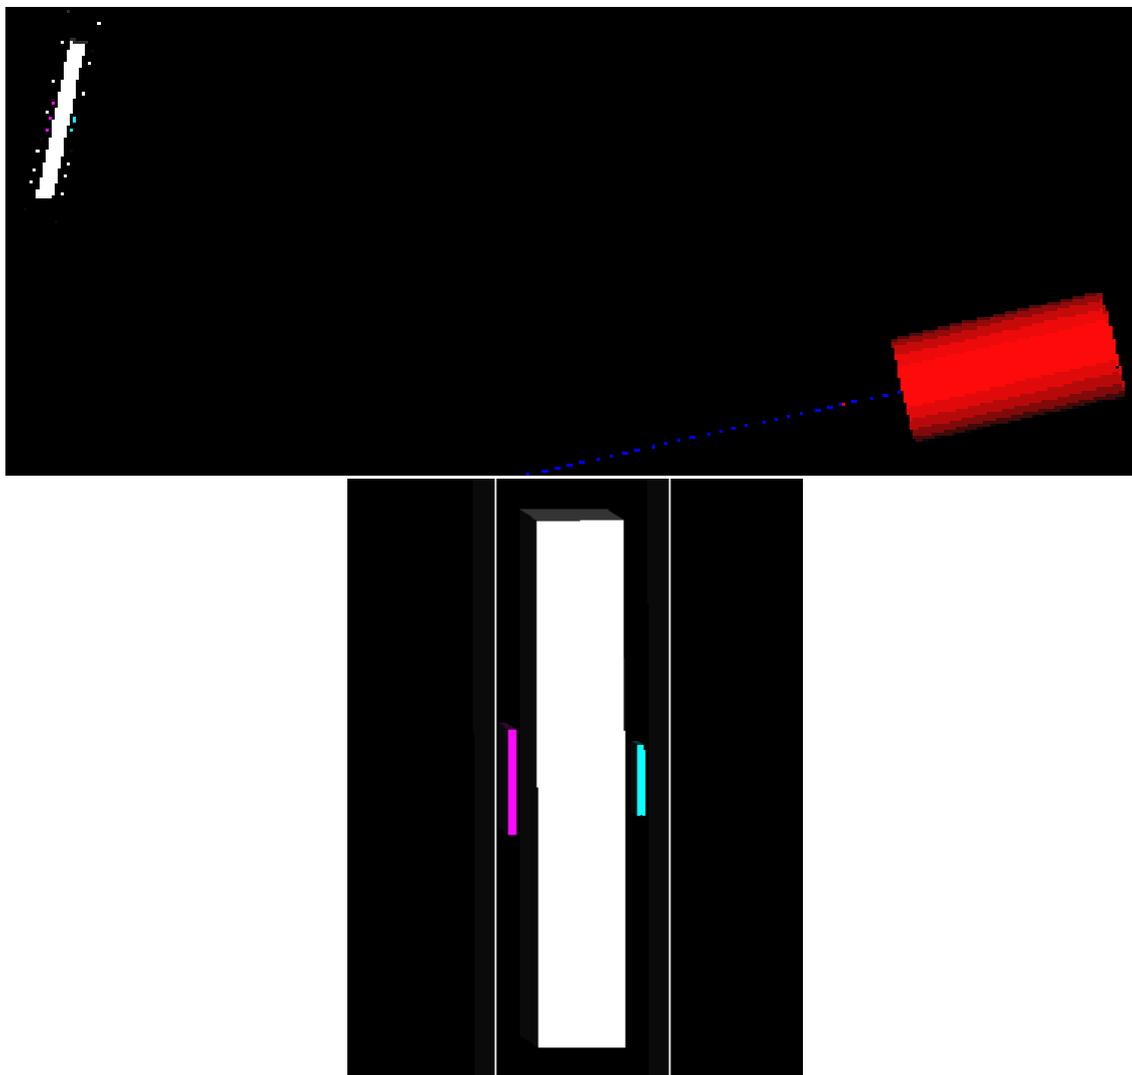


Figure 6.6: Above: A top-down view of the active simulation geometry with the LM model positioned 10 m and 25 degrees from the target. Below: A magnification of the LM model.

In order to increase statistical accuracy, the area of the box face, detector planes and plastic were increased by a factor of one hundred compared to the experimental geometry, whilst the thickness of each remained to scale. The area of the first pair of scintillators was therefore $20 \times 20 \text{ cm}^2$ and the second pair $30 \times 30 \text{ cm}^2$. The remaining sides of the aluminium box were not included since their effect would not scale with area. As in the live experiment, a count was recorded for a set of two scintillators if and only if both scintillators were traversed by a charged particle. There were 200 simulations run per physics model considered, each

with 10^7 protons fired at the target. In the remainder of our discussion we will refer to this set of simulations as ‘Simulation A’.

Simulation A should accurately predict the rate of C12, when normalised to area. A distribution of particle x positions (defined as the dimension perpendicular to the plane of the target and the LM) at the first scintillator, shown in figure 6.7 is seen to be flat, demonstrating that the density of particles detected scales with area. However it cannot produce a valid result for the second set of scintillators or the four-fold coincidence of all scintillators. The area of the detector planes has increased by a factor of 100 but the distance between them has not changed. As such, any event which produces a hit in the first pair of scintillators has a reduced probability, with respect to the experiment, of missing the second pair through scattering in the plastic. Nevertheless the result for the first set of scintillators should scale with area.

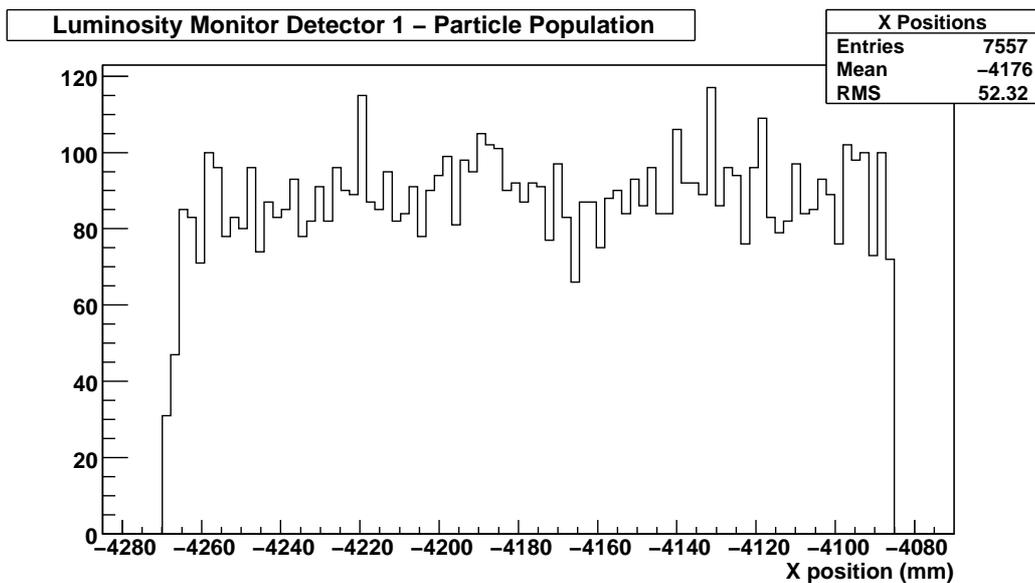


Figure 6.7: The x positions of particles at the first scintillator, which is at an angle of 25 degrees to the x -axis. The area of the first scintillator has been enlarged by a factor of 100 however the density of particle hits remains constant.

To resolve this issue, a second set of simulations, which we will call Simulation B, was produced with a luminosity monitor model of the correct dimensions inside the full aluminium box. In this instance, a snapshot of the particles prior to the luminosity monitor face was taken from Simulation A, scaled to the size of the real luminosity monitor and fired

Model	Explanation
QGSP_BIC	Quark Gluon String model for high energy hadronic interactions, using Binary cascade below approximately 10 GeV
QGSP_BERT	Quark Gluon String model for high energy hadronic interactions, using Bertini cascade below approximately 10 GeV
QGSC_BERT	QGSP for initial reaction, chiral invariant phase-space (CHIPS) decay (multi-quasmon fragmentation) for system fragmentation, Bertini cascade for nucleon and pion reactions.
QGSC_CHIPS	As QGSP but applies CHIPS modelling for nuclear de-excitation.
LHEP	Parametrised modelling of hadronic interactions

Table 6.3: Explanation of GEANT4[172] hadronic physics models used.

through Simulation B. This simulation set was used purely to establish the transmission, ie the ratios of C34 and C1234 to C12. The population of inbound particles is generated from the recorded flux across the enlarged geometry of the first scintillator plane in Simulation A, scaled to the realistic geometry of Simulation B. Therefore, the absolute counts in all scintillator pairs in Simulation B are grossly inflated. For this reason, the numbers produced in simulation B are used to calculate the relative transmission only.

Significant variation was observed as a function of physics model. Five models were studied, and these are described in table 6.3 with results broken down according to physics model in table 6.4.

Physics Model	$\frac{C34}{C12}$ (%)	$\frac{C1234}{C12}$ (%)
QGSP_BIC	113 ± 3	63.4 ± 2.2
QGSP_BERT	102 ± 3	58.2 ± 2.0
LHEP	56.0 ± 2.6	28.8 ± 1.7
QGSC_CHIPS	113 ± 5	66.7 ± 3.4
QGSC_BERT	105 ± 3	58.8 ± 2.0

Table 6.4: The percentage transmission between C12 and both C34 and C1234, calculated in simulation B. Note that the transmission may exceed 100% in C34 due to the larger area of the second set of scintillators (9 cm² in C34, compared with 4 cm² in C12).

6.5.2 Geometry Factor

Extensive simulation work, using G4Beamline [133], has been completed studying the performance of the luminosity monitor in concert with both the current target geometry and that of the previous MICE target in use prior to 2009. MICE had previously employed a titanium fin of cross-section $1 \times 1 \text{ mm}^2$ and length 10 mm. Simulations were performed using this target and their results, described in table 6.5, are similar to those obtained by Walaron [169] for an earlier study using a simplified luminosity monitor. Comparisons for coincidence channels other than C12 are not possible since these were not included in the earlier simulation. This study has shown, using 188 simulations and 1.88×10^9 POT, a performance of $(1.67 \pm 0.01) \times 10^{-8} \frac{\text{counts}}{\text{POT} \times \text{cm}^2}$ for C12 compared with $(1.70 \pm 0.1) \times 10^{-8} \frac{\text{counts}}{\text{POT} \times \text{cm}^2}$ in the study by Walaron for simulations featuring the previous target geometry. These results compare favourably to data taken previously. Further data taking in this configuration is no longer possible, since the titanium fin target is permanently beyond use.

The current target geometry is that of a hollow titanium cylinder, 10 mm in height with inner radius of 2.3 mm and outer radius of 3 mm as shown in the rightmost picture in figure 6.8. When normalising simulated detector counts using the new target we must specify the number of protons interacting with the target. Let us assume that the probability of a proton interacting with a titanium target sufficiently to leave the ISIS ring is the product of some efficiency and the thickness of the target. In reality this efficiency will depend upon the interaction length of protons in Titanium and the momentum spread (or ‘momentum bite’) over which ISIS can maintain their trajectory in the accelerator. Characterising the momentum bite of ISIS reliably is a significant undertaking beyond the scope of this thesis and is being studied by other members of MICE using software that tracks particles along ISIS. However, we may assume that the current conditions in ISIS have remained constant since the round of data taking involving the original MICE fin shaped target. The thickness of the target places a dependence upon the target geometry. In the case of the cylindrical target, there is variable thickness, contributing to an average thickness¹ of 1.945 mm compared to

¹The average thickness is calculated from a hypothetical rectangular target with a width the size of the subtended face and equivalent material volume as the cylindrical target.

Quantity	Simulated Yield ($\frac{\text{counts}}{\text{POT} \times \text{cm}^2} \times 10^{-8}$)	Experimental Yield ($\frac{\text{counts}}{\text{POT} \times \text{cm}^2} \times 10^{-8}$)
C12	1.67 ± 0.01	n/a
C34	1.07 ± 0.10	n/a
C12 (Walaron)	1.70 ± 0.10	1.59 ± 0.24

Table 6.5: Comparison of results for fin target geometry between this study and that of Walaron [169]. This study benefits from greater statistics, with 1.88×10^9 POT compared to 10^7 .

the thickness of 10 mm present in the original target.

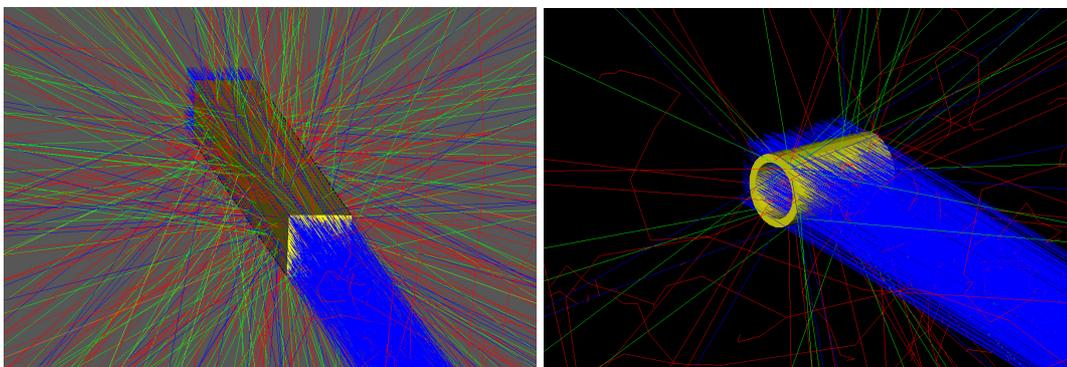


Figure 6.8: A proton beam containing 10^7 particles inbound upon a model of the previous (left) and current (right) MICE target. The previous, thicker target produces a greater number of proton interactions than the new target. Here protons are coloured blue, neutrons red and pions green. Figure produced using G4Beamline.

Therefore, when considering the number of protons interacting with the current target we would expect the correction factor to be $\eta = 0.1945$. Performing the identical simulation but replacing the old target geometry with the new target geometry yields the results for C12 of $0.283 \pm 0.002 \times 10^{-8} \frac{\text{counts}}{\text{POT cm}^2}$. This differs from the case of the old target by a factor of $\eta = 0.169 \pm 0.001$, close to our correction factor. This shows that the geometrical efficiency scales with the amount of material in the target. However, the probability of interaction and the mechanism by which primary protons are removed from the ISIS beam are more complex, so we will use this factor as our correction factor which takes into account the new geometry of the target and the probability that a proton interacts with this new geometry. In discussing the body of simulations featuring the new target, this correction factor η will be used to approximately scale the protons inbound on the target in the simulation to the protons which are considered to have interacted with the target. As such, it is necessary to include it

when calculating the total LM yield.

6.5.3 Particle Population

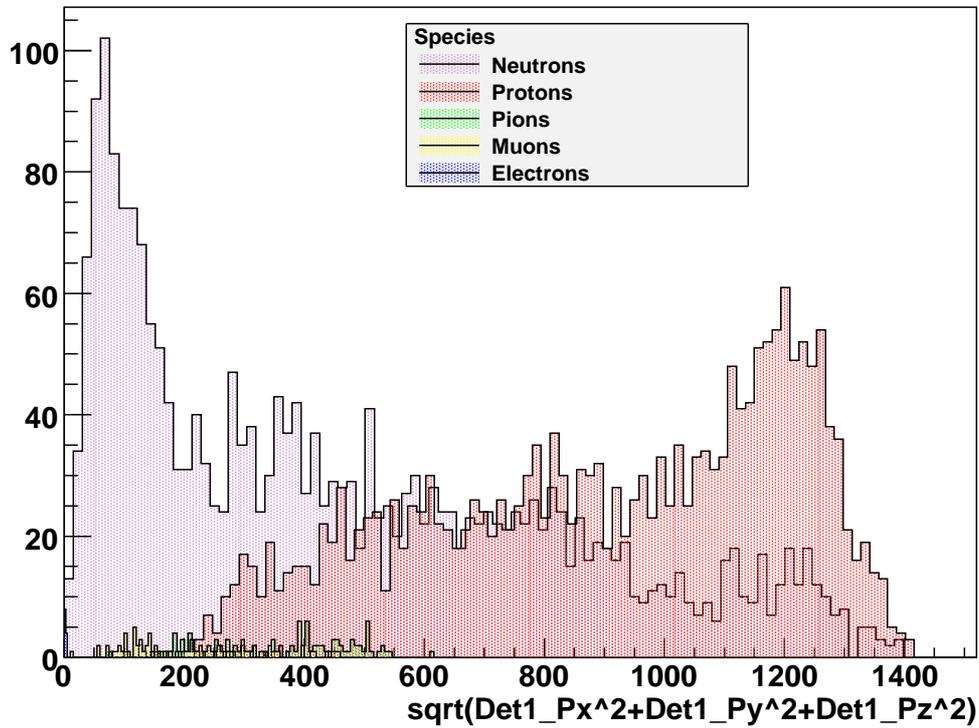
A comparison has been made between QGSP_BIC, the official hadronic model for use in MICE and a selection of other available hadronic models. The models studied are explained in table 6.3. In the energy regime of MICE, QGSP_BERT and QGSC_BERT provide near-identical results, with small differences on account of statistical variations. This is because they are using an identical hadronic model below 10 GeV. Significant variation was observed in the remaining models, with LHEP performing particularly poorly, exaggerating the quasi-elastic peak close to 1400 MeV/c and effectively failing to model protons above 800 MeV/c. This model should not be used for simulations featuring the luminosity monitor.

The momentum distributions before and after the polyethylene filter, for the model QGSP_BIC are shown in figure 6.9. Figure 6.10 shows the momentum at the first set of scintillators for those particles which create events in all four scintillators (Channel C1234). Using this model, we see clear evidence of a proton momentum filter due to the the polyethylene block at around 570 MeV, as opposed to the design prediction of 500 MeV. Momentum distributions for the remaining hadronic models at the first and second sets of scintillators are shown in figures 6.11-6.14.

Clearly the dominant particle species are protons and neutrons with small numbers of electrons, pions and muons also observed. Tables 6.6-6.8 catalogue the number of particles observed in the simulations for each of the remaining hadronic models according to species.

Although the dominant particle species, neutrons may only be detected in the event that they produce a proton in the polyethylene. The efficiency of neutron detection is summarised in table 6.9. These results show that although neutrons are the dominant particle, the rate at which they produce a charged particle hit in both scintillators in coincidence is sufficiently low that they may be safely neglected. In the discussion of the LM yield which follows, neutron efficiency has been included.

Momentum Distribution: QGSP_BIC Channel: C12



Momentum Distribution: QGSP_BIC Channel: C34

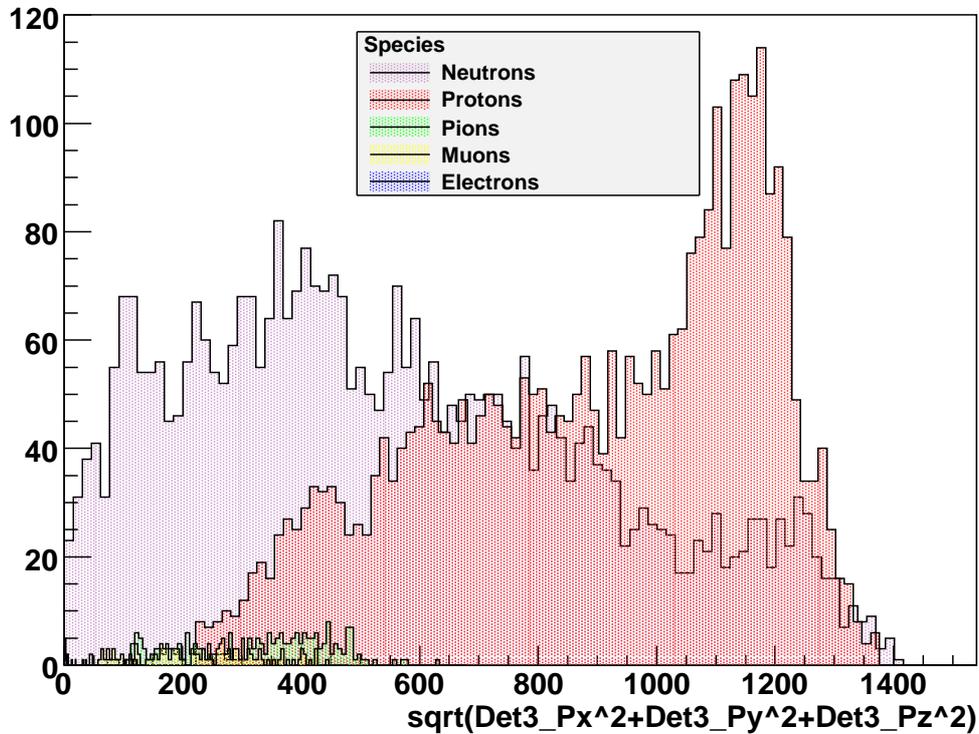


Figure 6.9: The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the QGSP_BIC hadronic model.

Momentum Distribution: QGSP_BIC Channel: C1234

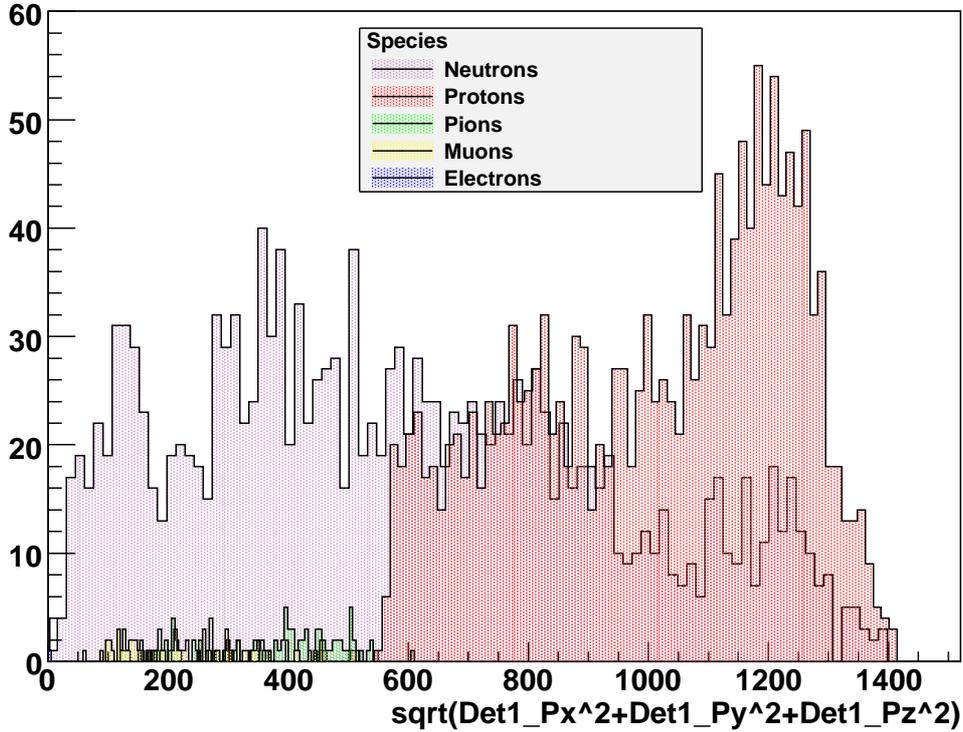


Figure 6.10: The momentum distributions at the first set of scintillators for those events which create hits in all four scintillators (Channel C1234). Particle species is recorded at the first set of scintillators. Note the effect of the polyethylene in filtering the protons below approximately 570 MeV.

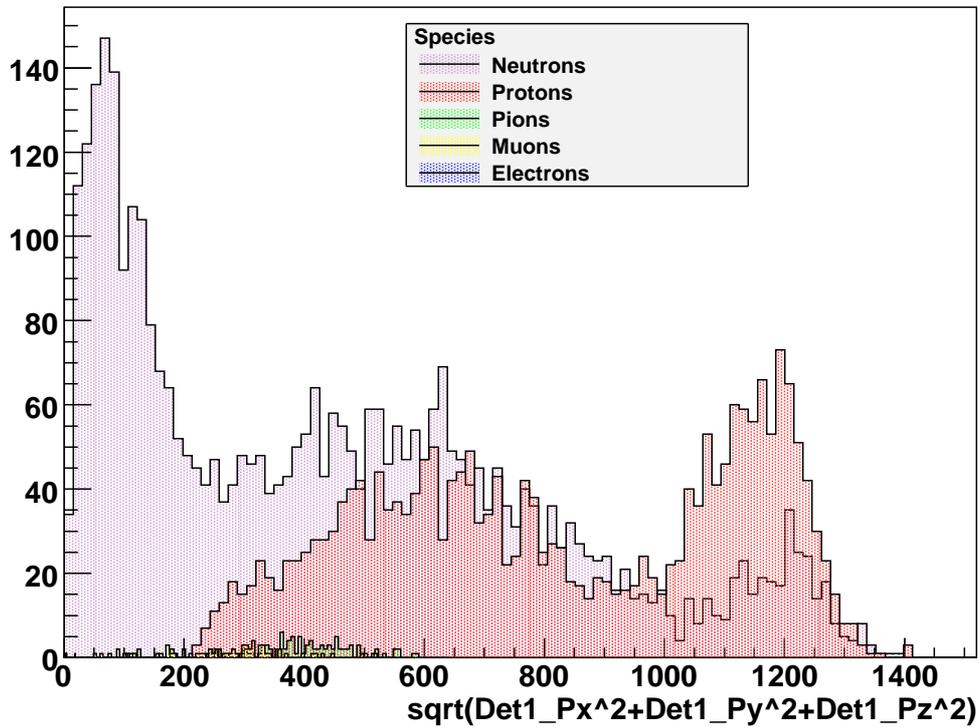
Species	QGSP_BIC				QGSP_BERT			
	C12	C12(%)	C34	C34(%)	C12	C12(%)	C34	C34(%)
n	2445	52	3937	50	3838	61	5688	60
p	2061	43	3469	44	2279	36	3341	35
π^+	92	1.9	173	2.2	90	1.4	187	2.0
π^-	19	0.40	47	0.60	12	0.19	25	0.27
μ^+	80	1.7	134	1.7	59	0.94	128	1.4
μ^-	30	0.63	43	0.55	9	0.14	17	0.18
e^-	14	0.29	15	0.19	6	0.095	11	0.12

Table 6.6: Particle populations in simulation A for the QGSP_BIC and QGSP_BERT hadronic models.

6.5.4 Yield

We will now calculate the simulated and experimental response, or yield, of the LM channels with respect to the charged particles listed in tables 6.6-6.8.

Momentum Distribution: QGSP_BERT Channel: C12



Momentum Distribution: QGSP_BERT Channel: C34

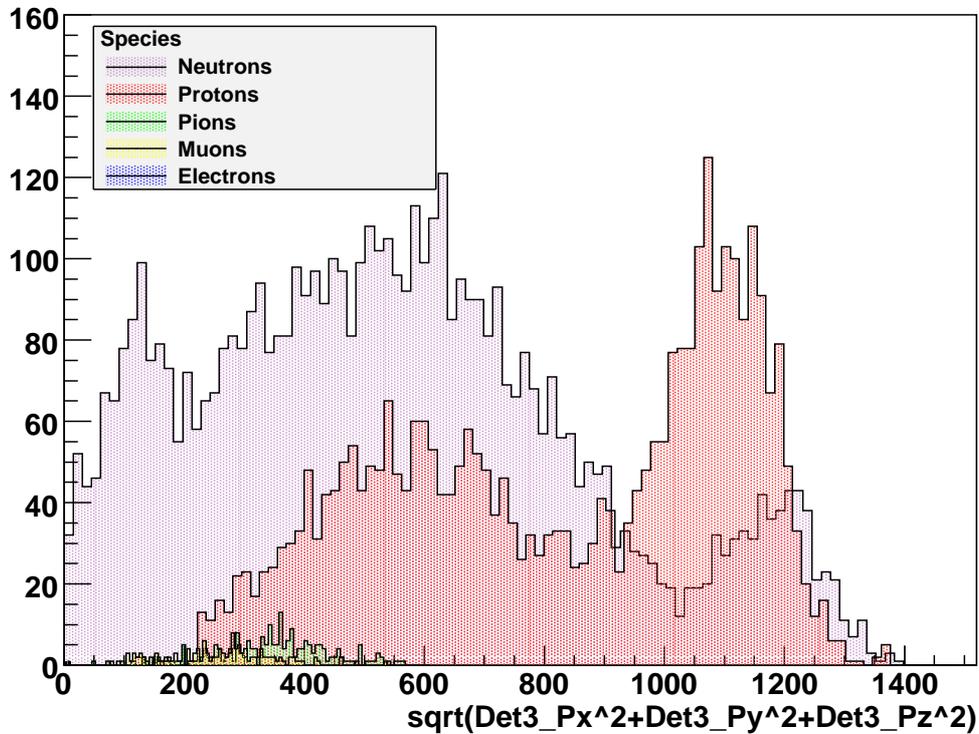


Figure 6.11: The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the QGSP_BERT hadronic model.

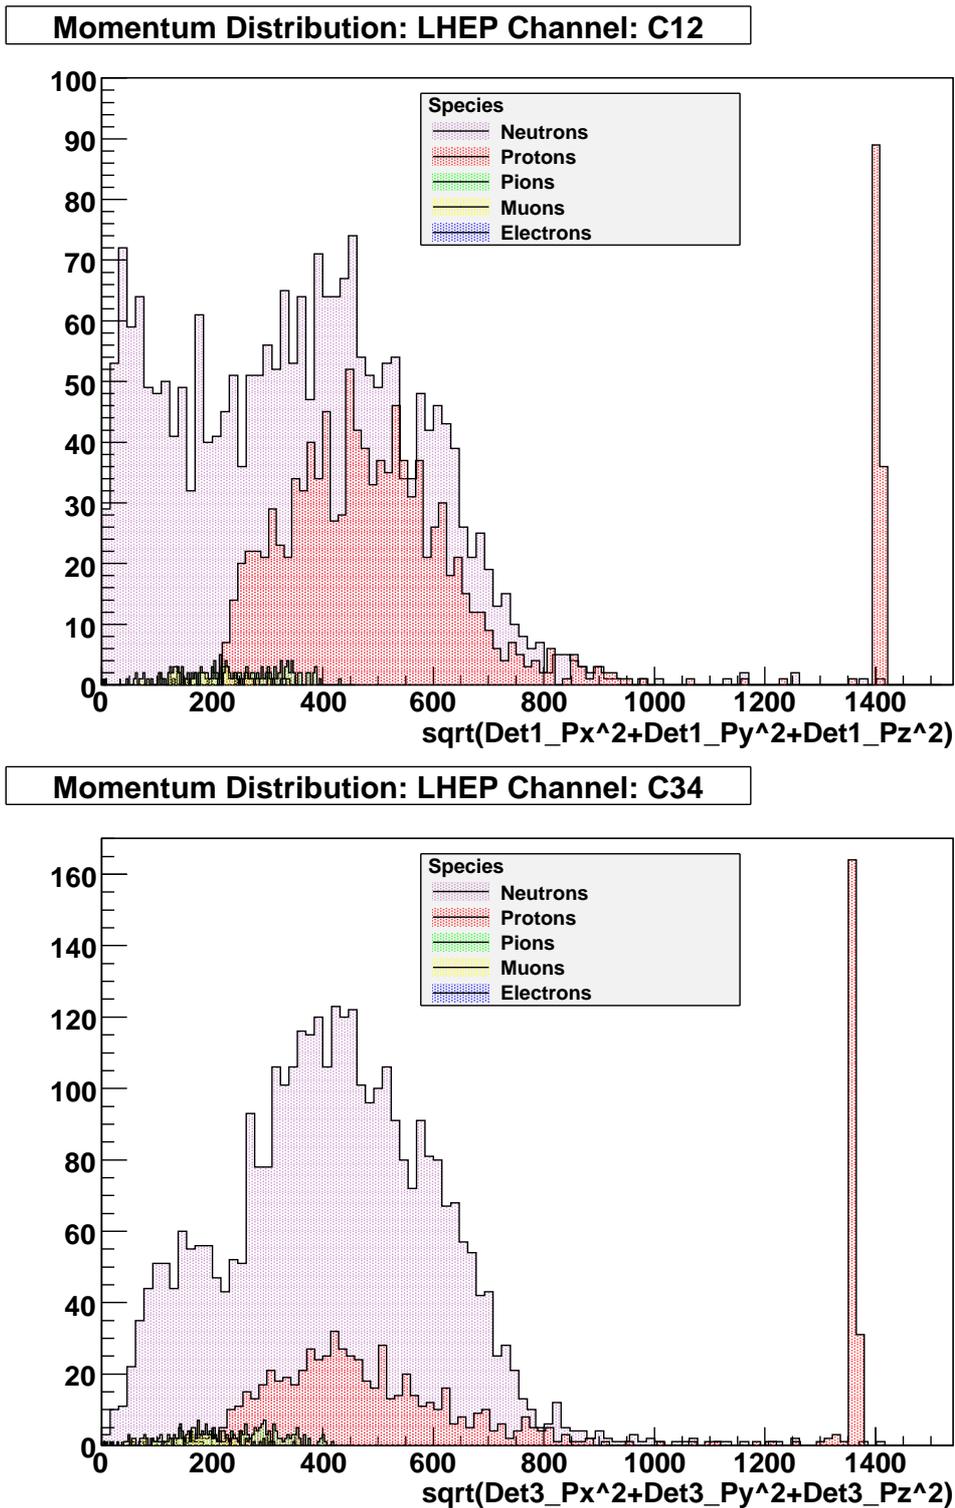
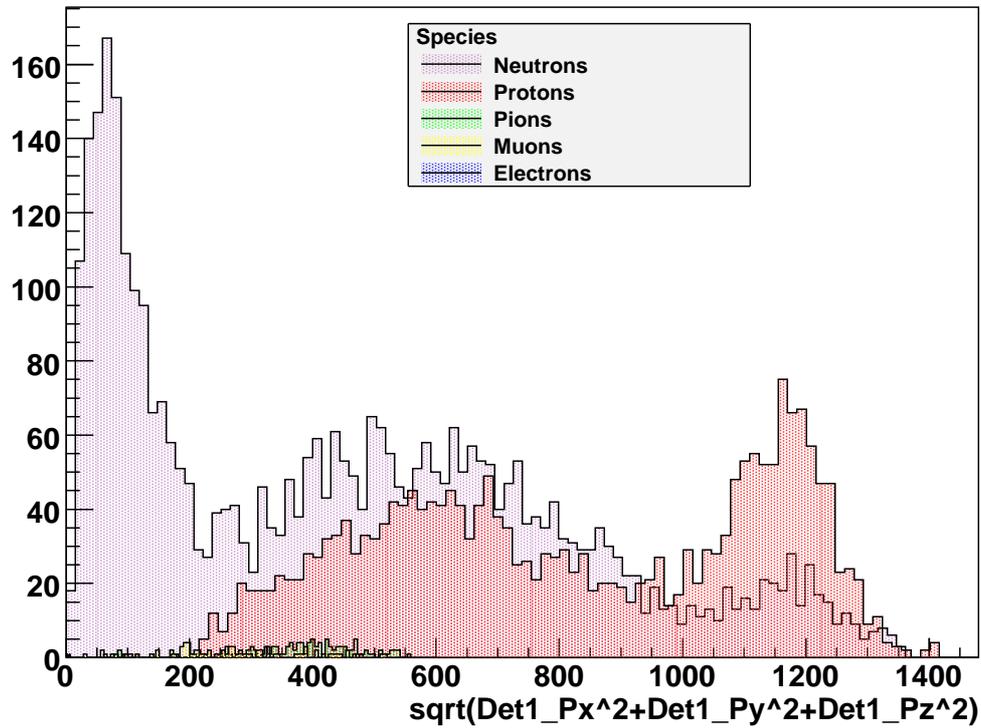


Figure 6.12: The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the LHEP hadronic model. The momentum distributions provided by this parametrised model appear to handle higher momentum protons particularly poorly and predict an overall reduced rate inconsistent with data.

Momentum Distribution: QGSC_BERT Channel: C12



Momentum Distribution: QGSC_BERT Channel: C34

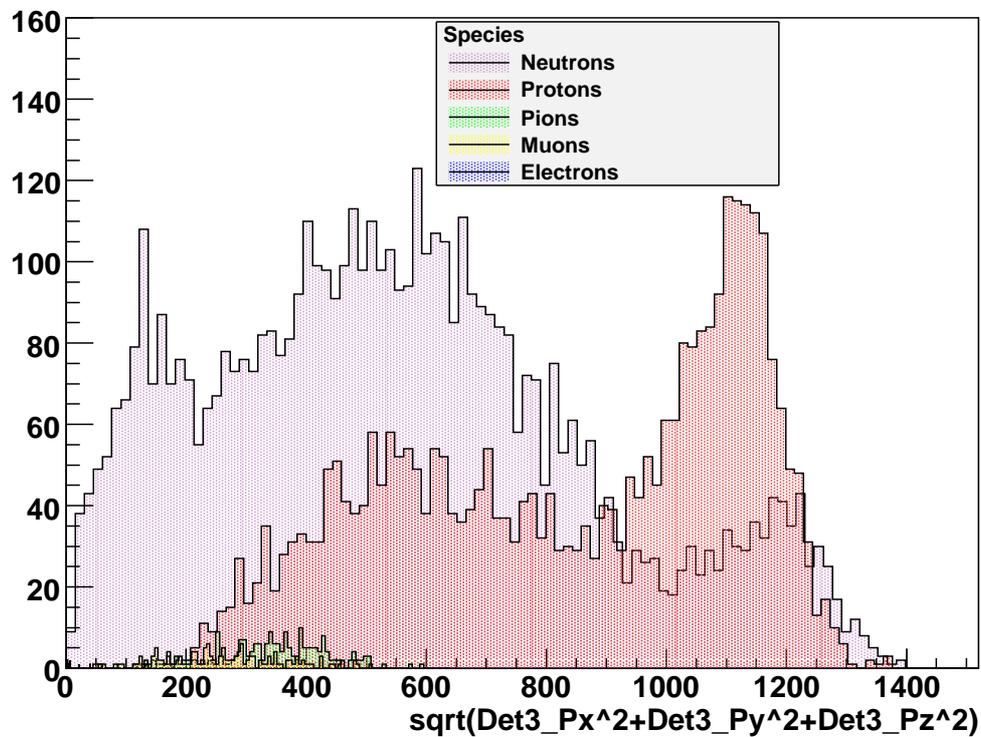
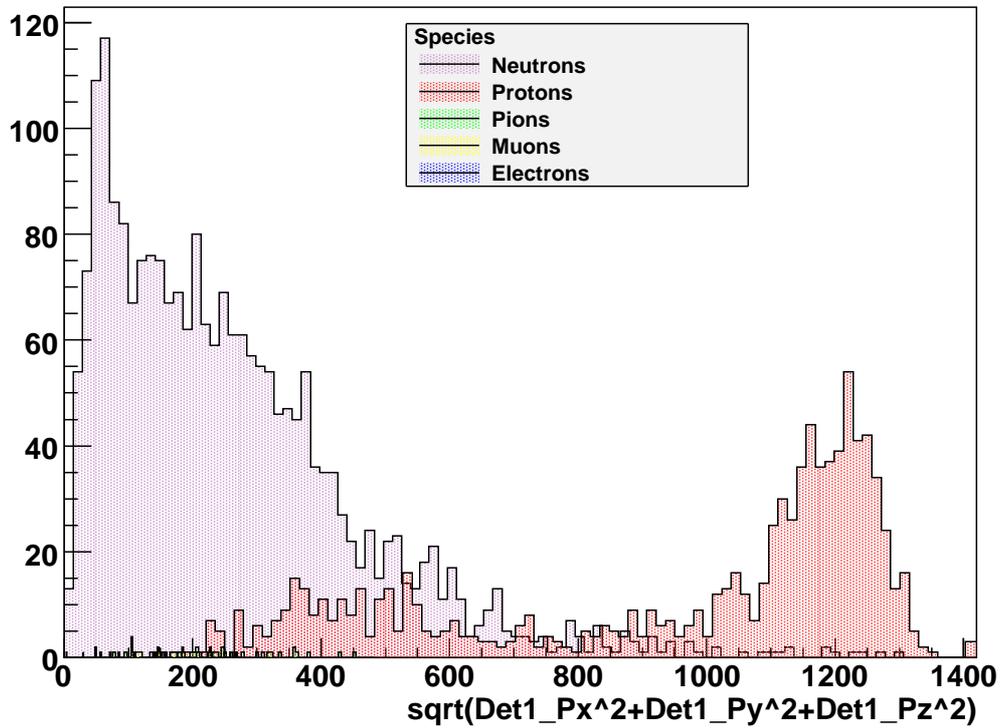


Figure 6.13: The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the QGSC_BERT hadronic model.

Momentum Distribution: QGSC_CHIPS Channel: C12



Momentum Distribution: QGSC_CHIPS Channel: C34

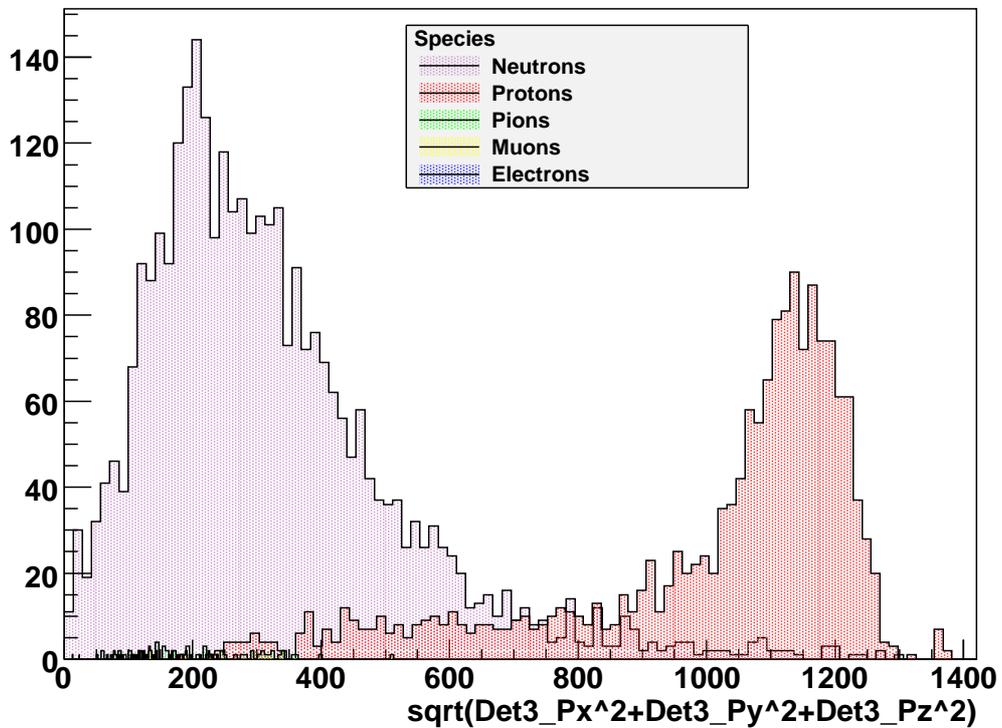


Figure 6.14: The momentum distributions at the first (top plot) and second (bottom plot) set of scintillators using the QGSC_CHIPS hadronic model.

Species	QGSC_CHIPS				QGSC_BERT			
	C12	C12(%)	C34	C34(%)	C12	C12(%)	C34	C34(%)
n	2274	69	3216	65	3851	61	5492	59
p	927	28	1606	32	2252	36	3352	36
π^+	18	0.55	36	0.73	92	1.4	173	1.9
π^-	15	0.45	19	0.38	26	0.41	36	0.39
μ^+	138	1.1	42	0.85	76	1.2	138	1.5
μ^-	24	0.58	19	0.38	14	0.22	24	0.26
e^-	15	0.091	2	0.040	10	0.16	15	0.16

Table 6.7: Particle populations in simulation A for the QGSC_CHIPS and QGSC_BERT hadronic models.

Species	LHEP			
	C12	C12(%)	C34	C34(%)
n	3542	71	3509	74
p	1158	23	820	17
π^+	82	1.7	135	2.9
π^-	40	0.80	53	1.1
μ^+	94	1.9	142	3.0
μ^-	41	0.82	63	1.3
e^-	8	0.16	13	0.27

Table 6.8: Particle populations in simulation A for the LHEP hadronic model.

The C12 yields quoted in table 6.10 are calculated from the C12 counts in Simulation A. The yields for C34 and C1234 are calculated by multiplying the yield for C12 by the transmission ratios calculated from simulation B and listed in table 6.4. These are normalised in the following way:

$$Y_{12} = \frac{C12}{POT \times 2 \times 2 \text{ cm}^2} \times \frac{1}{\eta} \quad (6.2)$$

$$Y_{34} = \frac{C12 \times T(34)}{POT \times 3 \times 3 \text{ cm}^2} \times \frac{1}{\eta} \quad (6.3)$$

$$Y_{1234} = \frac{C12 \times T(1234)}{POT \times 2 \times 2 \text{ cm}^2} \times \frac{1}{\eta}, \quad (6.4)$$

where Y_{12} , Y_{34} and Y_{1234} each denote the yield of a given coincidence channel; ‘POT’ is the number of protons fired at the target; T is the transmission and denotes the ratio of particles from the first set of scintillators producing a coincidence in the second set of scintillators

Model	Efficiency C12 (%)	Efficiency C34 (%)
QGSP_BIC	0.18 ± 0.09	0.63 ± 0.24
QGSP_BERT	0.088 ± 0.051	1.2 ± 0.3
QGSC_CHIPS	0.098 ± 0.069	0.60 ± 0.27
QGSC_BERT	0.088 ± 0.051	0.59 ± 0.18
LHEP	0.22 ± 0.08	0.55 ± 0.22

Table 6.9: The efficiency of neutron detection. For C12, this is the efficiency with which neutrons knocks off a proton in the first scintillator which then creates a coincidence in both scintillators. For C34, this is the efficiency with which a neutron knocks off a proton in either of the first two scintillators or the plastic shielding which subsequently creates a coincidence in C34. In both cases we consider only those events which cross all four scintillators.

(or a fourfold coincidence of 1, 2, 3 and 4) and η is the efficiency with which a proton passing through the target is lost in ISIS.

Model	Yield ($10^{-8} \frac{\text{counts}}{\text{POT cm}^2}$)		
	C12	C34	C1234
Data	1.72 ± 0.52	0.82 ± 0.25	0.79 ± 0.24
QGSP_BIC	1.67 ± 0.01	0.84 ± 0.03	1.06 ± 0.04
QGSP_BERT	1.80 ± 0.02	0.81 ± 0.03	1.05 ± 0.04
LHEP	1.03 ± 0.01	0.25 ± 0.01	0.29 ± 0.02
QGSC_CHIPS	0.74 ± 0.01	0.37 ± 0.02	0.49 ± 0.03
QGSC_BERT	1.80 ± 0.02	0.84 ± 0.03	1.06 ± 0.04

Table 6.10: Response of the luminosity monitor as a function of hadronic model. Each simulation has an identical geometry description and is given the same set of functional parameters, including 2×10^9 protons on target. Errors in data are dominated by the 30% error attributed to the beamloss measurements whilst errors in the simulation are statistical.

Again, it should be noted that QGSP_BERT and QGSC_BERT produce near-identical results. These models provide the results closest to the experimental yield, followed by QGSP_BIC. LHEP is clearly incapable of correctly modelling the LM response. QGSC_CHIPS is also not favoured.

Within errors, QGSP_BERT, QGSC_BERT and QGSP_BIC match the experimental results. Nevertheless due to the significant error attributed to the conversion of beamloss to protons on target it is difficult to draw definitive conclusions regarding which model performs best from these numbers alone.

6.6 Saturation Limits

It is not yet known what the beamloss constraints set by ISIS will be in the latter stages of the MICE experiment, with higher beamloss increasing the number of muons in the cooling channel and therefore reducing the necessary running time of the experiment. Since the luminosity monitor is, at present, unshielded and positioned close to the target then it is likely that high beamloss could result in saturation of the PMTs. During a set of high beamloss runs the saturation point was studied and results are shown in figure 6.15. The trigger width was set at 0.5 ms rather than the normal 3.2 ms for lower beamloss runs, in order to avoid placing stress upon the data acquisition system.

Saturation occurred from roughly 4 Vms onwards. Beyond this point the PMTs are no longer able to operate quickly enough to distinguish the end of one event from the beginning of another. As beamloss increases, so too does this effect, gradually reducing the count rate to some limit close to zero because the signals from the PMTs are constantly firing and never reset. The individual rate of each of the PMTs is not recorded in data by the data acquisition system. However, we may infer the singles rate from these results as follows:

$$R_{ab} = 2N_a N_b \tau, \quad (6.5)$$

where R_{12} indicates the random coincidence rate through PMTs a and b; N indicates the rate for a single PMT and τ is the gate width of the PMT. By inspection of figure 6.15 we may infer that saturation begins at approximately 1,000 counts in channel C12 for a trigger width of 0.5 ms (a rate of $\sim 2 \times 10^6$ Hz). Assuming that $N_1 \approx N_2$ (the rate of each PMT is approximately the same) and that the signal gate width is $\tau = 10$ ns, then:

$$N_1 = \sqrt{\frac{R_{12}}{2\tau}} \quad (6.6)$$

$$= \sqrt{\frac{2 \times 10^6}{2 \times 10 \times 10^{-9}}} \quad (6.7)$$

$$= 10^7 \text{ Hz} \quad (6.8)$$

The estimated rate in a single PMT at saturation is therefore 10 MHz. The onset of saturation will therefore commence when single PMT count rates are ~ 10 MHz.

A separate study undertaken by Drummond [173] demonstrated that this saturation point can be increased to 8 Vms with 30.1 cm of polyethylene shielding before the first set of scintillators. For experimental running from 2011, 5 cm of polyethylene will completely surround the LM. It is expected that the LM should be able to operate at a rate 12.5% higher than without shielding, i.e. below 4.5 Vms, with 5 cm of shielding.

6.7 Conclusions

A luminosity monitor has been designed for the MICE Experiment. Installation and commissioning at ISIS has been completed. The luminosity monitor is well placed for studying rate close to the target, normalising rate throughout the MICE channel and validating simulations. One may calculate the protons on target from luminosity counts independently to the readings from the ISIS beamloss monitors.

A suite of simulations has been performed providing predictions of the ratio of luminosity counts to protons on target. Given a realistic estimate of protons interacting in the target, simulations using QGSP_BERT, QGSC_BERT and QGSP_BIC provide close descriptions of the LM response, whereas the parametrised model LHEP appears to poorly describe the transmission as a function of momentum for the dominant species of particles. Neither QGSC_CHIPS nor LHEP provide a realistic prediction of the rate at the luminosity monitor and as such are disfavoured.

The luminosity monitor has shown strong performance throughout Step I without evidence of saturation except in cases of exceptional beamloss > 4 Vms. Shielding will be installed at a later date in order to raise the saturation point further and enable MICE to increase its operational limits.

The LM is currently operational and will continue to provide particle rate data throughout future data taking runs. A first full analysis of the MICE beamline exploiting the LM will now follow in chapter 7.

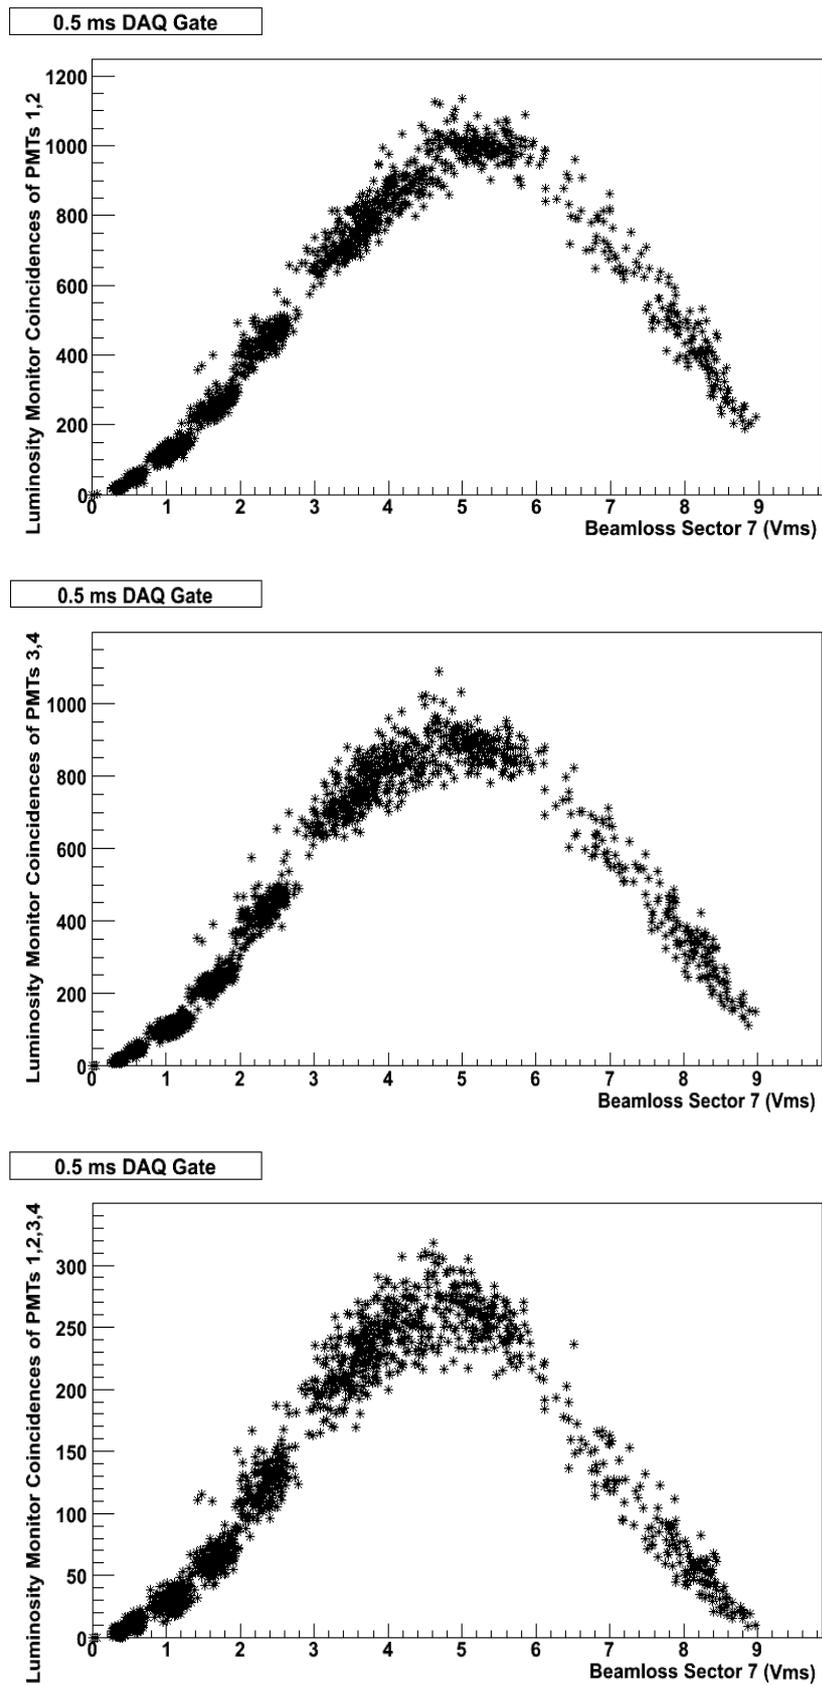


Figure 6.15: Luminosity counts against beamloss for a set of high beamloss runs with DAQ gate width of 0.5 ms. Saturation is clearly evident from around 4 Vms.

Chapter 7

Step I Analysis of the MICE Beamline

7.1 Introduction

MICE has recently undergone Step I of its program; constructing, characterising and commissioning a muon beamline. This chapter presents the first analysis using the previously described luminosity monitor (chapter 6) to provide relative normalisation of particle rate between the LM and each of the detectors in MICE. Of principal importance in Step I are the time of flight (TOF) detectors TOF0 and TOF1. These detectors enable both momentum reconstruction and particle identification. Their response as a function of luminosity will be analysed and compared to simulation. An additional comparison will be made to quantify the effectiveness of the simulation in predicting the beam profile widths governed primarily by multiple scattering and magnetic effects.

The reference settings of the MICE beamline are characterised by the emittance-momentum (ϵ, p) matrix shown in figure 7.1.

As a matter of MICE terminology, a beam is said to have a momentum of 200 MeV/c if it is designed to provide muons of 200 MeV/c in the central absorber of Step VI, however, it is expected that the actual momentum will vary along the beamline and MICE channel due to particle decay and material present. For example, a 200 MeV/c beam may have momentum of 450 MeV/c at the beginning of the first quadrupole triplet and 254 MeV/c by the end of the second quadrupole triplet [175].

		p_z (MeV/c)		
		140	200	240
ϵ_N (mm·rad)	3	t=0.0 mm $P_{diff}=151$	t=0.0 mm $P_{diff}=207$	t=0.0 mm $P_{diff}=245$
		$\alpha=0.2$	$\alpha=0.1$	$\alpha=0.1$
		$\beta=56$ cm	$\beta=36$ cm	$\beta=42$ cm
	6	t=5.0 mm $P_{diff}=148$	t=7.5 mm $P_{diff}=215$	t=7.5 mm $P_{diff}=256$
		$\alpha=0.3$	$\alpha=0.2$	$\alpha=0.2$
		$\beta=113$ cm	$\beta=78$ cm	$\beta=80$ cm
	10	t=10.0 mm $P_{diff}=164$	t=15.5 mm $P_{diff}=229$	t=15.5 mm $P_{diff}=267$
		$\alpha=0.6$	$\alpha=0.4$	$\alpha=0.3$
		$\beta=198$ cm	$\beta=131$ cm	$\beta=129$ cm

Figure 7.1: The emittance-momentum (ϵ, p) matrix. For each emittance and momentum, a set of beamline optics has been defined. α and β are the familiar Twiss parameters and t defines the thickness of lead diffuser, used to inflate the emittance. P_{diff} is the momentum immediately before the diffuser [174].

Each element of the matrix corresponds to a set of pre-defined optics developed by Apollonio [176]. Step 1 data taking has focused on populating this matrix with experimental data corresponding to each matrix element. This provides a basis for a comparative study between the results of the official MICE simulations and data for defined reference configurations, which will be the subject of this chapter.

7.2 Step I Data Taking

For each element of the emittance-momentum (EM) matrix, a set of experimental runs exist which were recorded during Step 1 of MICE data taking. Each matrix element corresponds to a set of magnet currents, or beamline ‘optics’, described in [176]. These currents configure the beamline to select particles of a desired momentum and transmit them to the MICE channel.

The beamline may transport either negative or positively charged particles, depending on the polarities of the magnets. It is therefore possible to collect both negative and positive

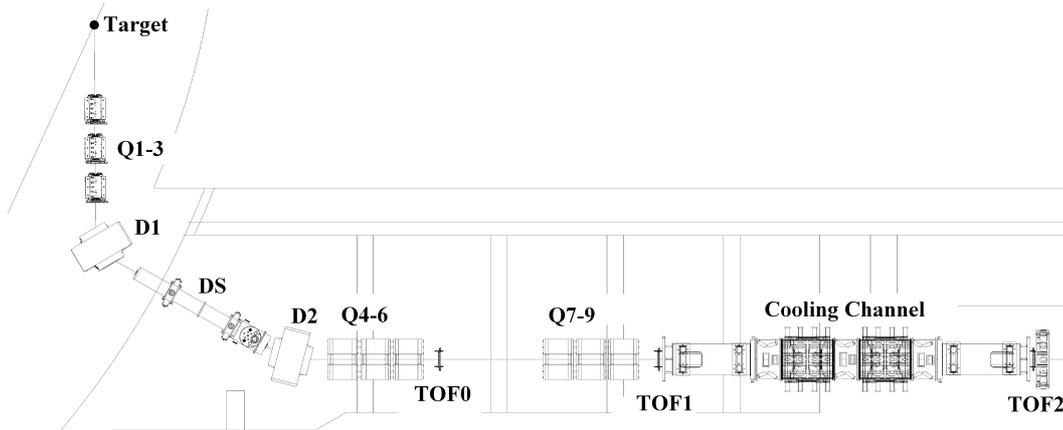


Figure 7.2: Schematic of the beamline showing the positions of each of the magnets and particle detectors relative to ISIS.

polarity data for each EM matrix element.

The data from these runs provide scaler readout for the particle counters (GVA1, BPM1, BPM2), TOF0 and TOF1 and all three luminosity monitor (LM) coincidence channels. These detectors are defined in section 3.9. A schematic of the beamline can be seen in figure 7.2. TDC readout is also available for each of the above detectors, excluding the luminosity monitor which provides scaler readout only.

We note that the response of the LM is independent of the magnetic lattice, since it is not located along the beamline (see section 3.9.1). Unless otherwise stated, the studies presented in this chapter make use of a selection of MICE runs¹, which were taken between the 8th and 30th of July 2010.

7.3 Step 1 Simulations

7.3.1 MICE Beamline

A full G4Beamline description of the MICE beamline geometry, terminating before the MICE Cooling channel, has been provided by Marco Apollonio (Imperial College). This geometry includes the magnetic lattice and detectors. However, the high statistics required

¹MICE run numbers used for the analysis work presented in this chapter: 2297, 2299, 2301, 2305, 2386, 2402, 2431, 2452, 2461, 2477, 2505, 2507, 2512, 2518, 2520, 2537, 2545 and 2564

at the target to produce substantial numbers of hits downstream at TOF1 precludes a full simulation of hadronic interactions in the target. Apollonio has therefore provided the input to the simulation by creating a ‘target box’ that produces particles with a largely uniform momentum distribution, weighting their species according to previous target simulation work, greatly reducing necessary running time for simulations. In his model, 1.97×10^9 protons yield 79,000 protons, 6,500 positively charged pions, 1,900 negatively charged pions, 49,000 electrons, 49,000 positrons and 7,200 photons. One may define a beam by specifying a number of protons which is a multiple of 1.97×10^9 .

Appropriate momentum distributions are secured by the magnetic system which selects a sample from the momentum distribution according to a given configuration of beamline optics. Figure 7.3 illustrates how this affects the pion momentum distribution. Pions will decay into muons which provide the basis of our emittance measurement. The figure shows the raw pion momentum distribution immediately after the target box. By excluding pions that do not create events in the MICE detectors we can see the momentum selection effect of the magnetic system very clearly.

These simulations are key to testing our understanding of data produced in Step I of MICE. They have not been fundamentally modified by the author except to correct detector geometries in light of the measurements provided in chapter 3 and adapt them for batch processing on the local computing cluster and Grid environments.

The configuration database (chapter 4) is used to provide a reliable record of the magnetic currents used for a given run. These currents are supplied to the simulation to recreate in simulation the experimental configuration for each element of the EM matrix.

The MICE Beamline simulations use the QGSP_BIC hadronic model, however in the absence of a rigorous high density target simulation relatively little modelling of hadronic interactions is required.

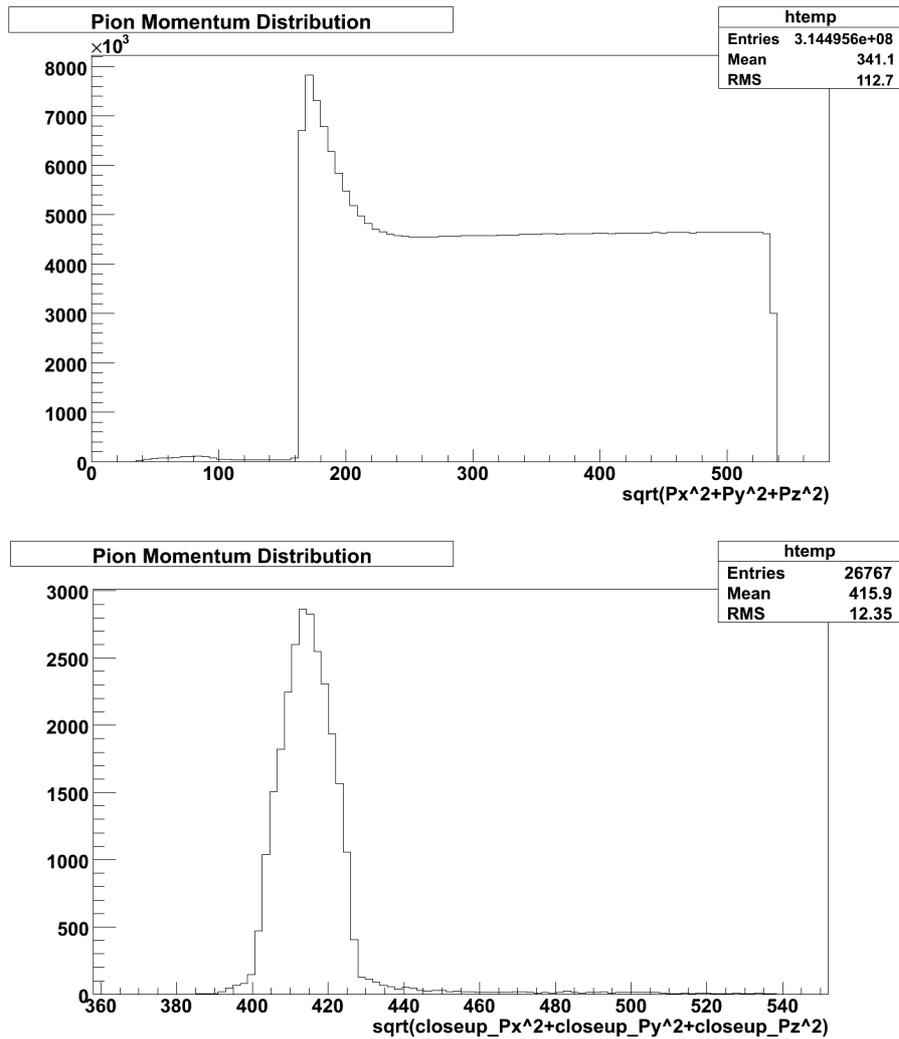


Figure 7.3: Top: The momentum spectrum of positively and negatively charged pions leaving the target box. Bottom: The same plot, excluding particles which do not create events in TOF1.

7.3.2 Luminosity Monitor

As stated previously, the MICE Beamline simulations cannot accurately model the response of the LM since they do not produce realistic momentum spectra in the area of the luminosity monitor. The QGSP_BIC simulations presented in chapter 6 are used again here. The number of hits in each LM channel may be compared with detector hits in the beamline simulation by normalising both sets of simulations to protons on target. The LM results are also downscaled to allow for the factor of 100 increase in area.

7.4 Detector Efficiency

The efficiency for each detector was calculated by constructing from real data a sub-sample consisting only of those particles which created a hit in the immediate neighbour detector(s). From here one may take the ratio of the number of events causing a hit in the detector of interest and the total number of events in the sample.

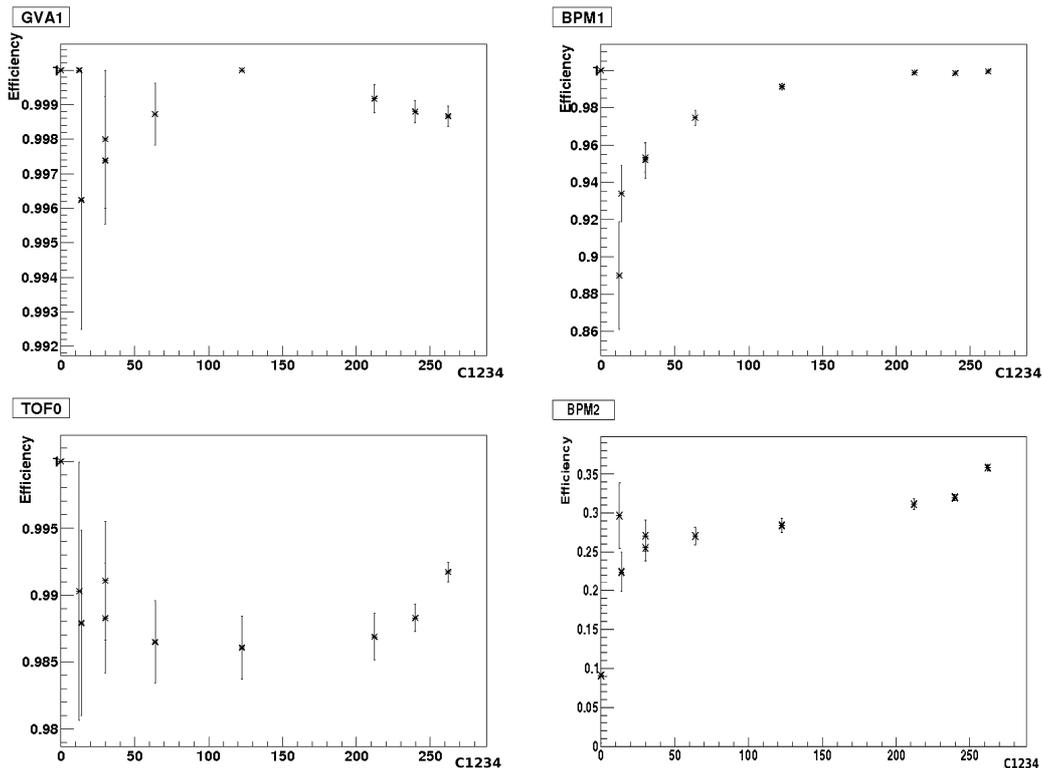


Figure 7.4: Detector efficiency as a function of average C1234 counts. For each run, an average luminosity (C1234) is calculated from each spill. The graph shows efficiency remains flat as a function of luminosity for each detector considered.

Repeating this procedure for many runs with varying luminosity shows that efficiency remains reasonably stable. Figure 7.4 provides an illustrative example during a period of stable running conditions (Runs 2875-2884 taken on 14/8/10 as part of a beamloss study). In practice one may calculate an efficiency for each set of runs used. The calculated efficiencies are listed in table 7.1, which shows that the efficiencies found for runs 2875-2884 are close to 100% in the case of GVA1, BPM1 and TOF0. Since no detector data exists downstream of TOF1, our method is less reliable here and yields a 100 % efficiency. In the studies which follow the efficiency of TOF1 will be assumed to be equal to that of TOF0.

A significantly lower efficiency has been found for BPM2. Further investigation has shown that although the vertical plane operates with a similar efficiency to BPM1, the horizontal plane reduces the overall efficiency to 32.60%. This effect is illustrated in figure 7.5. For this reason we will not use the BPM detectors to normalise rates.

Detector	Efficiency (%)
GVA1	99.98 ± 0.01
BPM1	91.7 ± 0.1
TOF0	99.0 ± 0.1
BPM2	32.6 ± 0.2
TOF1	$100.0^{+0.0}_{-0.1}$

Table 7.1: Numerical detector efficiencies taken as an average over each of the results in figure 7.4.

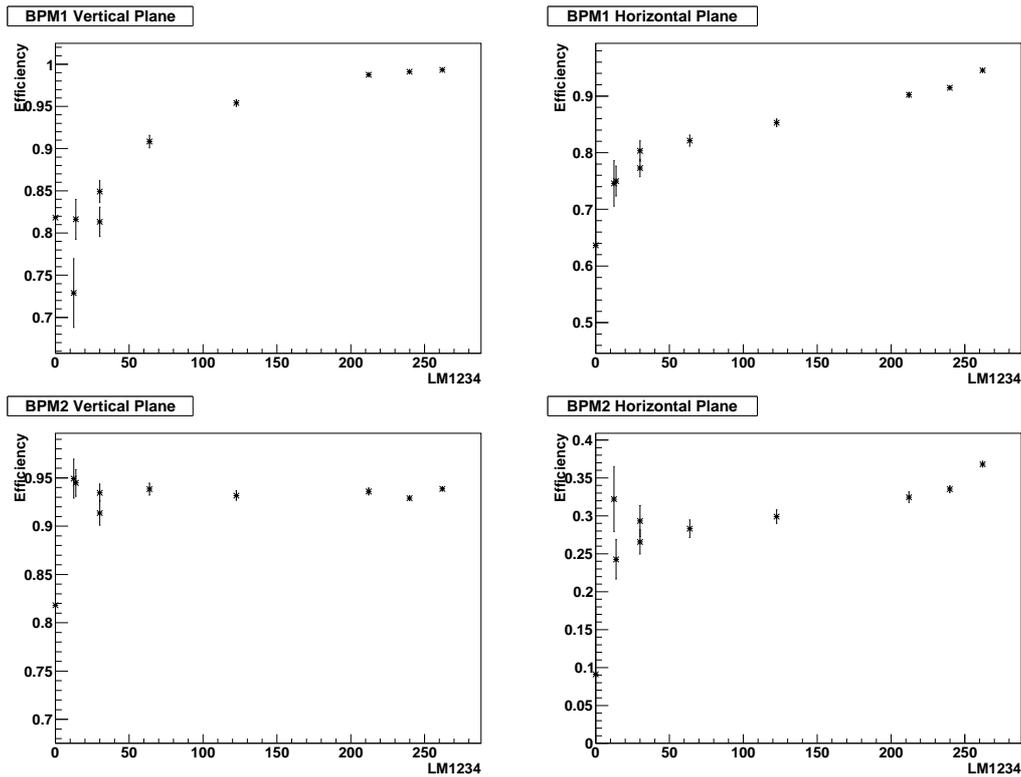


Figure 7.5: The efficiencies of individual BPM planes for BPM1 (top) and BPM2 (bottom). The vertical plane (bottom left) of BPM2 operates with a greater efficiency than the horizontal plane (bottom right).

7.5 Particle Rate Per Detector

Assuming constant conditions and within the saturation limits of the detectors, the particle rate should scale with luminosity counts at each detector along the beamline. Particle rate per luminosity count can be calculated for each spill. The LM counts themselves are calculated from an average over the total number of spills, in figure 7.6, for a set of runs² corresponding to EM matrix element (6,200). Each point in the graphs refers to one run.

Detector	Gradient	Intercept
GVA1	6.12±0.04	61.3±6.6
BPM1	0.30±0.01	2.6±1.4
TOF0	0.36±0.01	3.0±1.5
BPM2	0.05±0.01	0.54±0.57
TOF1	0.10±0.01	1.5±0.9

Table 7.2: Detector rate per luminosity monitor count (C1234).

The linear relationships in figure 7.6 allow the calculation of the gradients in table 7.2 which describe the constant of proportionality between detector and luminosity monitor counts.

7.6 Particle Identification and Momentum Reconstruction

The purpose of the MICE Beamline is to provide a muon beam of variable momentum and emittance as input to the MICE cooling channel. Step I provides a first opportunity to analyse the properties of the beam using the time of flight detectors TOF0 and TOF1.

We consider first the case of the negative polarity (6,200) beam, which was generated experimentally during run 2564. Figure 7.7 shows the time taken to travel between TOF0 and TOF1 for events in coincidence with a particle trigger, defined to be at TOF1. We recall equation 3.25,

$$p = \frac{m}{\sqrt{\left(\frac{t_2}{t_1}\right)^2 - 1}}, \quad (3.25)$$

²MICE run numbers: 2399, 2432, 2546, 2564, 2601, 2876, 2877, 2878, 2879, 2880, 2882, 2888 and 2889.

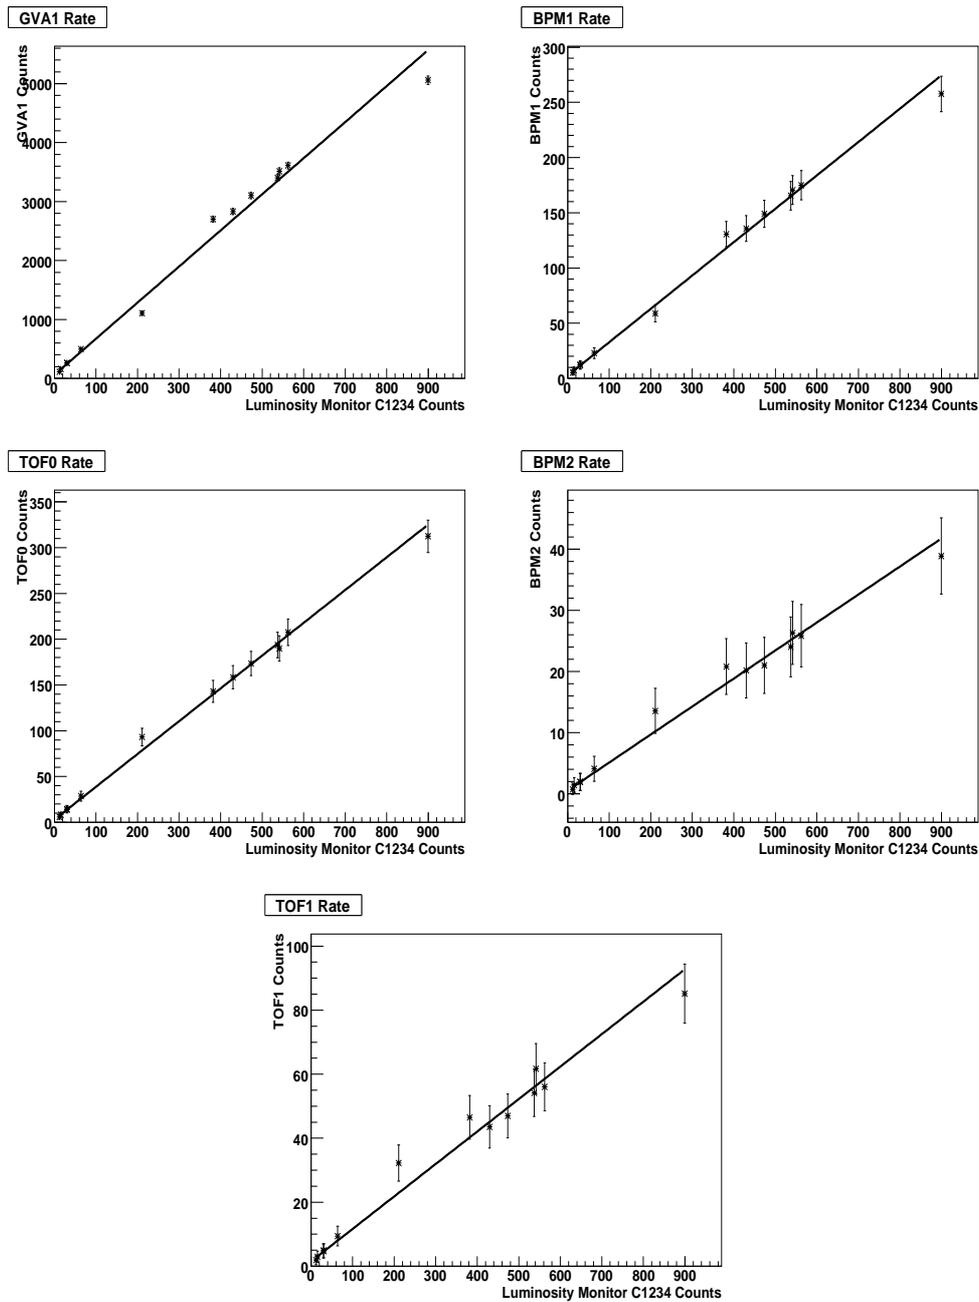


Figure 7.6: Particle rate for each detector. Detector rates are calculated by taking the average over all spills. Each data point corresponds to one run.

where p is the average momentum between the TOFs, m denotes particle mass, t_1 is the time taken by a particle travelling at the speed of light to traverse the distance between the TOFs and t_2 is the time taken by a particle travelling at a lesser speed v over the same distance. At TOF1 we expect to encounter only electrons, muons and pions. Due to their small mass, electrons will travel at approximately the speed of light. Muons and pions will

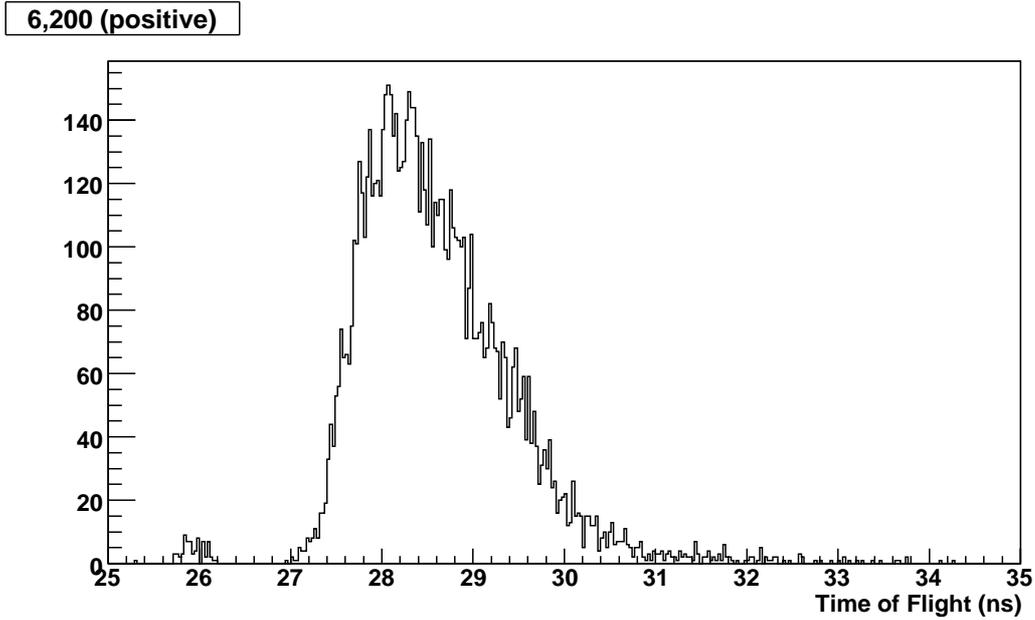


Figure 7.7: The reconstructed time of flight for particles travelling between TOF0 and TOF1.

travel at a slower rate according to their mass:

$$t_2 = t_1 \sqrt{1 + \frac{m^2}{p^2}}. \quad (7.1)$$

Figure 7.7 contains a well separated early peak around 25.93 ns, which clearly belongs to the electrons. Their small mass ($0.5 \text{ MeV}/c^2$) makes the peak well separated from muons (with a mass of $105.6 \text{ MeV}/c^2$) and pions (with a mass of $139.6 \text{ MeV}/c^2$). Relying solely upon the data it is less obvious where the muon peak ends and the pion peak begins.

Figure 7.8 provides predicted momenta distributions from simulation for pions and muons, for those events creating a hit in both detectors. Clearly muons are the dominant particle species, with pion contamination barely visible. Our simulation predicts a pion contamination of 0.5% in TOF0 and 0.3% in TOF1. Experimentally, momentum distributions are calculated from the time of flight between TOF0 and TOF1. A measurement of momentum using the time of flight system therefore yields a momentum at some point between the two detectors, therefore a pion contamination of 0.4% is assumed. We consider this level of contamination negligible.

The average momentum of simulated muons at TOF0 is 253.4 MeV/c with a width of

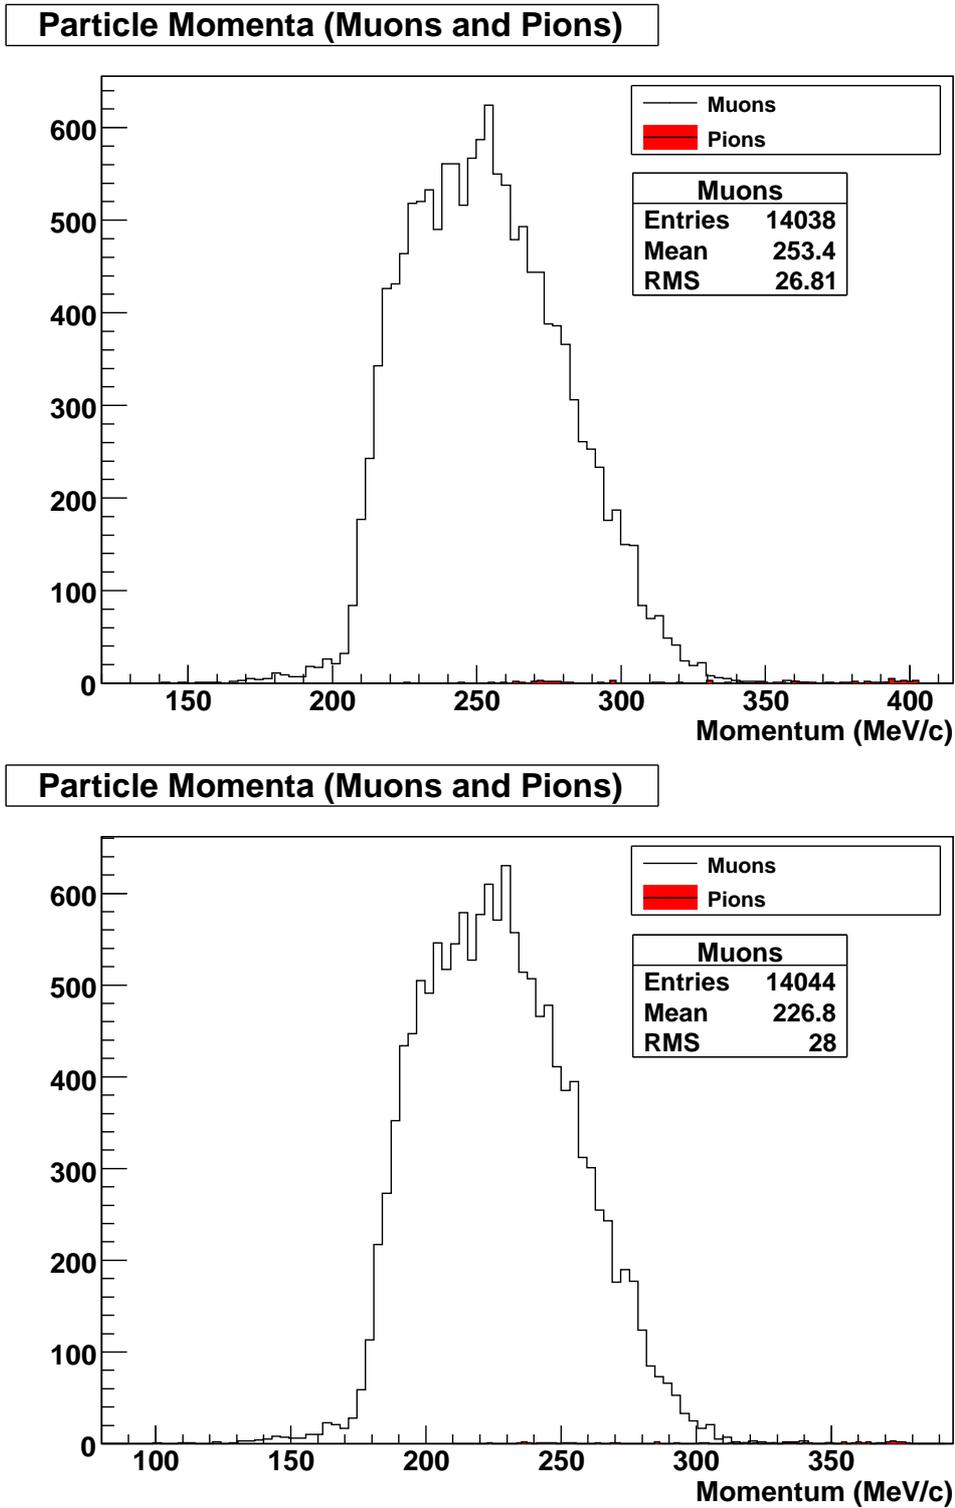


Figure 7.8: Momentum distributions for muons and pions at TOF0 (top) and TOF1 (bottom). The number of events, mean and RMS figures for the muon distribution are shown. The pion contamination, coloured in red, is very small and spread over a large range of momenta. As such it is very difficult to observe from the plots.

26.8 MeV/c, while the average at TOF1 is 226.8 MeV/c with a width of 28.0 MeV/c. The mean momentum between TOF0 and TOF1 will therefore be 240.1 MeV/c.

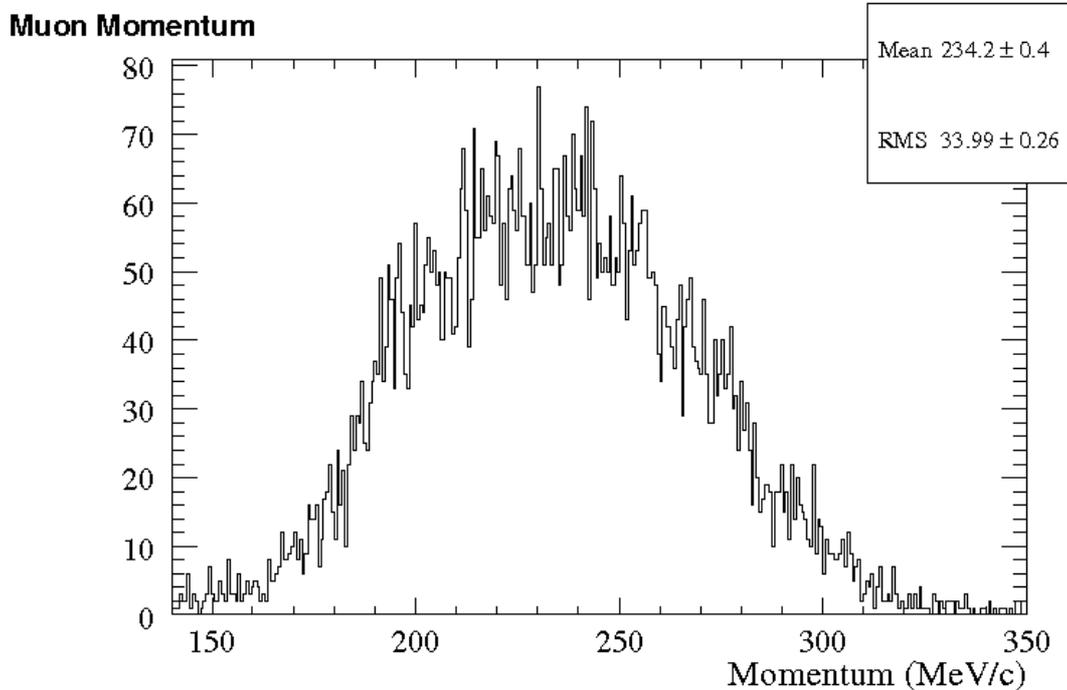


Figure 7.9: Momentum distribution between TOF0 and TOF1 for particles identified as muons in a (6,200) positive polarity beam.

When considering data we must isolate muons through a time of flight cut. In the above example a time of flight cut above 27 ns is applied. This cut varies depending on the momentum of the beam, with visual inspection of the time of flight distribution required to quantify the appropriate cut for each element of the EM matrix.

To calculate the muon momenta, we therefore convert the time of flight values directly to momenta using equation 3.25 for the large peak in its entirety. The distribution of muon momenta, taking t_2 for each particle directly from the data in figure 7.7 is given in figure 7.9. The graph suggests that a 200 MeV/c beam possesses a mean muon momentum of 234.2 MeV/c between TOF0 and TOF1, with a RMS width of 34.0 MeV/c.

This is 5.8 MeV/c lower than the simulated momentum. However, the difference can be explained by the fact that the particle trajectories are not straight lines in contrast to the assumption that TOF signals all emanate from the centre of the plane. Additional spread, due

to the slightly difference in the trajectories of particles hitting the edge of the TOFs compared with those hitting the centre of the detector should also be taken into account. Furthermore, the simulated beam is narrower in momentum due to the TOF time resolution, which should be convoluted with the TOF signal.

7.7 Predicted Performance of Step VI

In chapter 5 we have discussed the importance of reducing statistical errors in order to make a sufficiently precise measurement of ionisation cooling at Step III and beyond. Applying the same particle identification techniques described in section 7.6, we may calculate the number of good muons per spill for each element of the EM matrix. A good muon, for the purposes of this discussion, is a particle which reaches TOF1 and is identifiable as a muon from its time of flight between TOF0 and TOF1. In the final Step VI configuration, a good muon will also be required to reach the particle identification detectors downstream of the cooling channel, including TOF2.

Emittance (mm rad)	Momentum (MeV/c)		
	140	200	240
Negative Polarity			
3	7.47±0.14	11.5±0.1	9.18±0.13
6	4.59±0.08	8.15±0.12	6.87±0.13
10	6.34±0.20	7.21±0.08	5.85±0.09
Positive Polarity			
3	29.1±0.2	31.2±0.2	54.5±0.2
6	32.9±0.2	61.2±0.7	23.1±0.2
10	28.0±0.2	41.0±0.2	37.0±0.3

Table 7.3: Number of muons per 1000 C1234 counts, for runs considered in each element of the EM matrix. The positive polarity (6,200) magnetic optics were the most efficient in transmitting muons to TOF1.

Table 7.3 lists the number of muons per 1000 C1234 counts for each matrix element. Comparing this to the statistical requirements presented in section 5.3.2.5 for positive polarity 200 MeV/c beams, we arrive at an estimated duration for data taking for Step VI. The results in table 7.4 are calculated by assuming a beamloss of 2.0 Vms and dividing the total

Emittance (mm rad)	Time required (days)
3	71
6	4
10	99

Table 7.4: An estimate of the experimental running time required, at a beamloss of 2 Vms, to meet the statistical requirements of the MICE Experiment for positive polarity 200 MeV/c beams, rounded to the highest day. Though no formal calculation has been made of the error of these estimates, an error of 20% is suggested, to account for an unknown muon transmission and variations in muons per spill.

number of particles required, listed in table 5.1, by the number of muons per spill to arrive at the number of spills required. We assume a delay of 3.2 s between the start of successive spills and that the muons reaching TOF1 in Step I are all successfully transmitted through the Step VI cooling channel. As noted previously, the statistical error, and by extension the required number of muons, varies as a function of input emittance. For a 6 mm rad beam an order of magnitude less muons are required than in the 3 and 10 mm rad cases (table 5.1), which together with the increased rate listed in table 7.3, provides for a reduction in the time required for experimental running.

Should these runs be repeated with higher beamloss, the particle rate can be expected to increase, as shown in section 7.5. Increasing beamloss from 2.0 to 4.0 Vms would halve the necessary running time. This demonstrates that the goal of MICE Step I has been met: to construct a beamline that can deliver muons to carry out the full MICE programme.

7.8 Comparison with Simulation

The principal detectors of Step I are the Time of Flight (TOF) detectors which allow us to perform particle identification and will, as MICE evolves, allow for selection of muon events through the cooling channel. The ability to predict particle transmission and the beam profile at each TOF plane is crucial for the MICE collaboration to demonstrate its understanding of the beamline it has constructed. Such work is of additional value in validating the body of simulations which have informed the beamline design and will continue to prove essential to

the design of a neutrino factory.

7.8.1 Transmission Along Beamline

We wish to demonstrate the viability of the LM as a tool for normalisation both for simulation and experimental data taking. To this end, simulations of the MICE beamline and luminosity monitor, described in section 7.3, were generated for a given set of positive polarity optics (6,200).

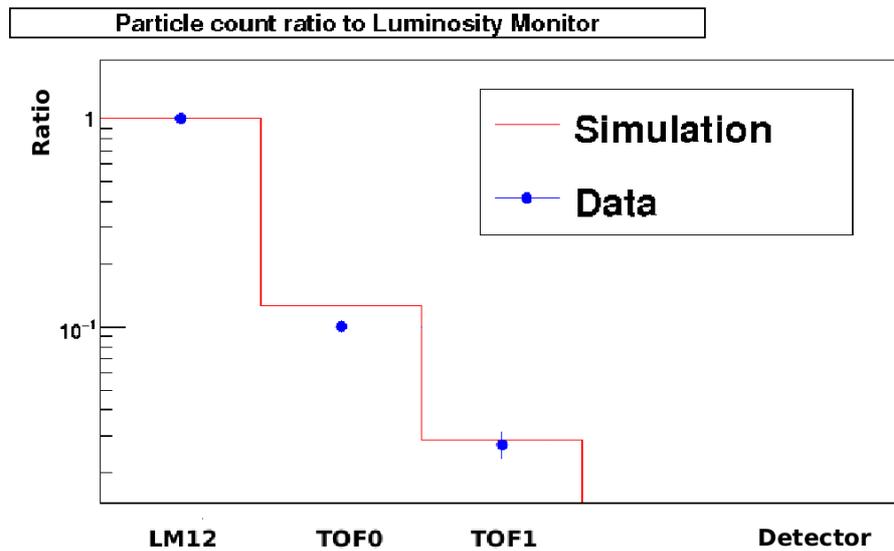


Figure 7.10: Rate through the MICE Beamline normalised to C12 LM counts, for experiment and simulation, for the positive polarity (6,200) beam.

The QGSP_BIC LM simulations from the previous chapter were combined with a simulation of EM matrix element (6,200). Comparable data was taken over 297 spills in run 2564. Figure 7.10 was constructed by taking the scaler rate in each TOF detector, normalised to luminosity counts (C12) for both simulations and experiment, folding in the efficiencies already described.

Figure 7.10 shows reasonable agreement between simulation and experiment, and similar charts may be constructed for any run, reinforcing our confidence in the LM as a tool for normalisation.

7.8.2 Simulated and Reconstructed TOF Profiles

During 2010, data was taken for each element of the EM matrix under a Step I geometry for both positive and negative polarities. Simulations using this same geometry have been described in section 7.3. Each plane consists of a number of slabs which are read out from two PMTs at either end. The slab-by-slab coincidences of PMTs are recorded in data, affording reconstruction of experimental profiles in each plane by applications within G4MICE. Time of flight cuts have been applied to select for muons which pass through both detectors in the manner described in section 7.6.

Simulated profiles have been generated using Grid computing resources. The simulations are configured with the precise magnetic currents recorded in the configuration database (chapter 4) for each experimental run corresponding to an EM matrix element. An equivalent slab-by-slab readout is inferred from the calculated slab position of each hit in each TOF plane. Events which produce muons in both planes of both TOFs are recorded for comparison with the time of flight tracks between both TOFs reconstructed from data. We note that the real detector presently suffers from a broken PMT affecting the rightmost slab of the x plane of TOF0. As such, simulated events which pass through this slab location are not recorded.

Both experimental and simulated profiles have been normalised according to their total number of hits and compared, with examples for the negative and positive polarity (6,200) EM matrix element shown in figures 7.11-7.14. No alignment correction has been applied to these figures however it is significant that the means of simulated and experimental distributions do not perfectly coincide, suggesting misalignment, which we will now quantify.

For the purposes of repeatability, we note that at present G4Beamline employs a right-handed co-ordinate system, where y is up, whilst G4MICE employs a left-handed co-ordinate system, due to a cabling error. This study has therefore translated the G4Beamline results into a left handed co-ordinate system.

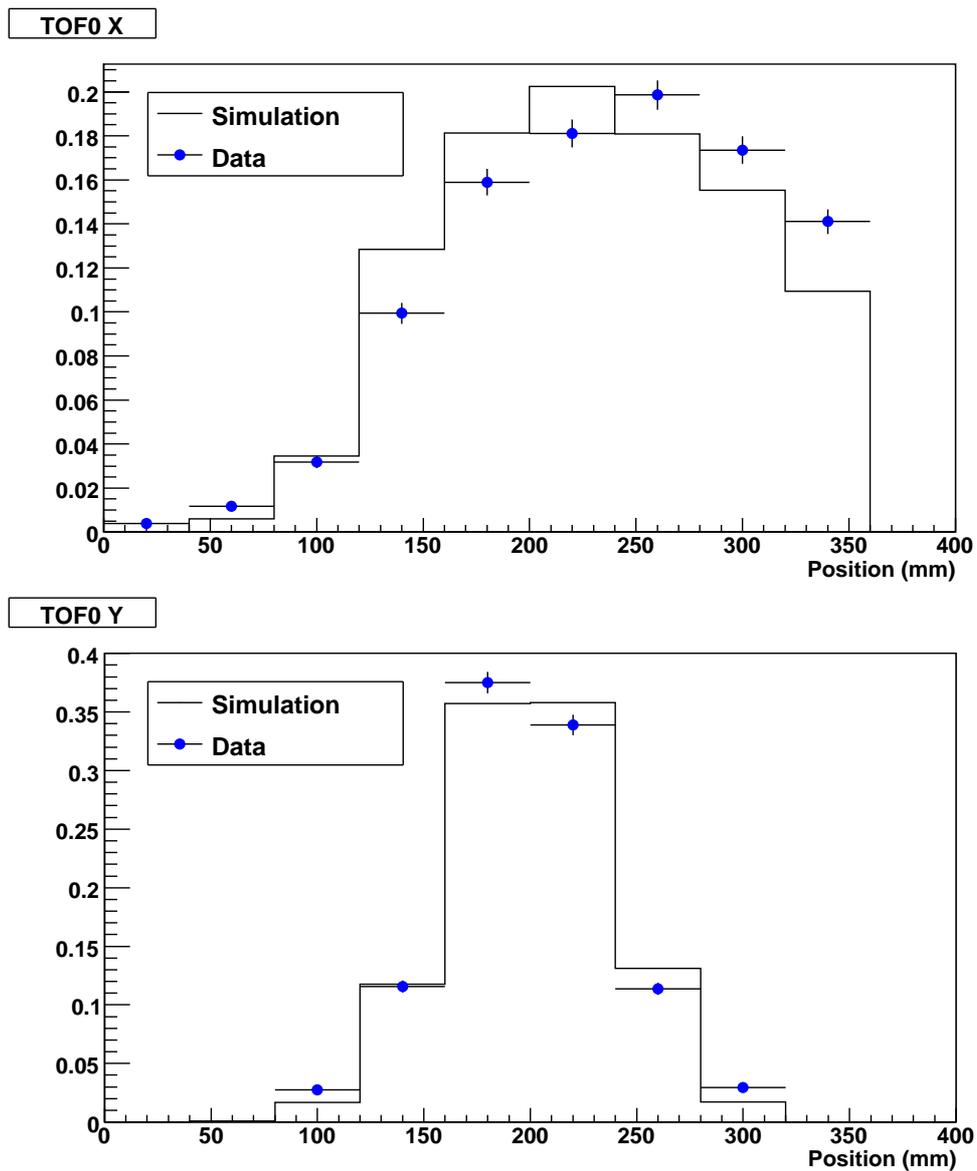


Figure 7.11: Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF0 for the negative polarity (6,200) matrix element.

7.8.2.1 Effective Misalignment

It has been noted that the real TOF detectors are misaligned relative to their design position. Rayner [177] has quantified the mechanical misalignment through comparison of observed detector centre positions with intended positions. However, there remains the possibility that distortions in the magnetic axis caused by the relative misalignment of the quadrupoles could produce a skew in the beam in a similar manner to a mechanical misalignment. Comparison of data and simulation provides an opportunity to quantify these ‘effective misalignment’

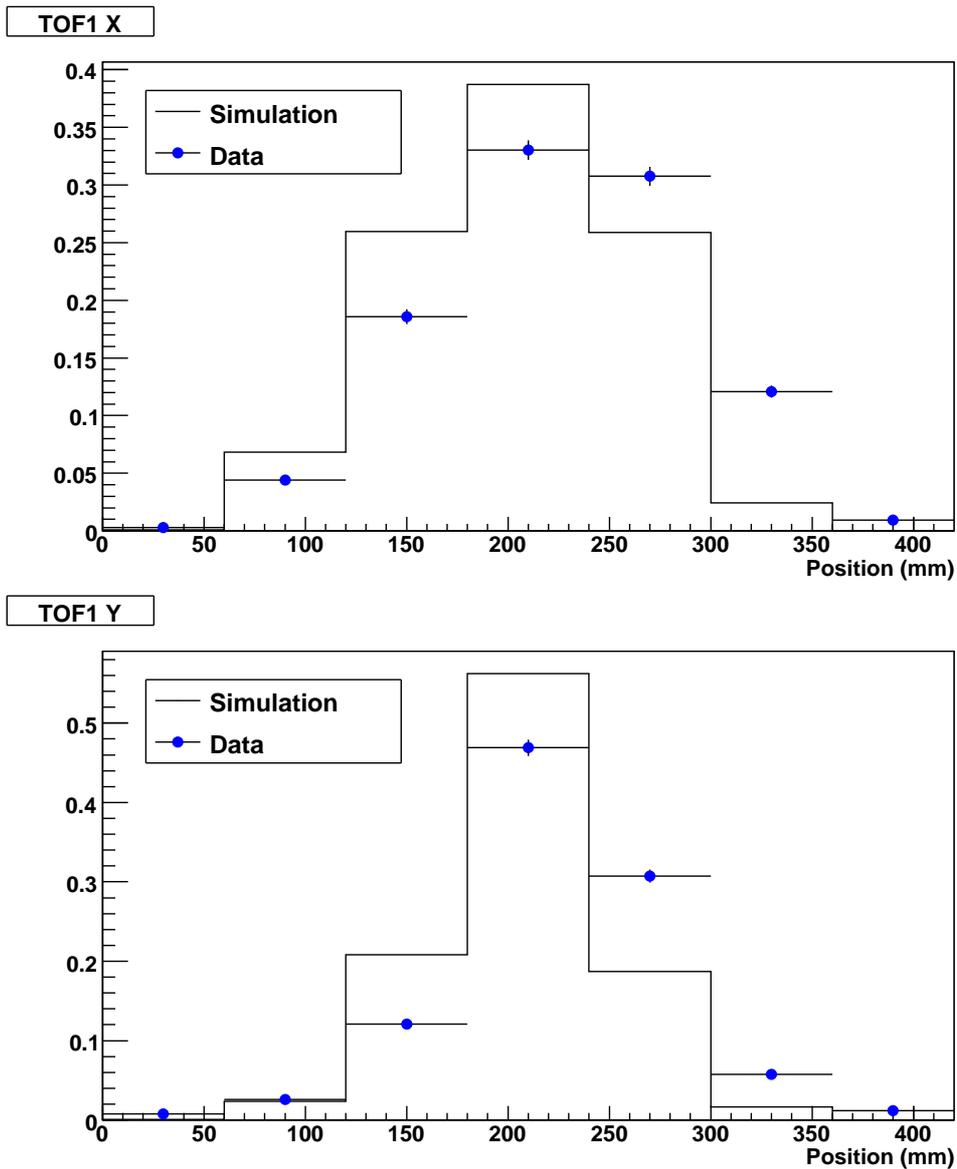


Figure 7.12: Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF1 for the negative polarity (6,200) matrix element.

corrections which include both mechanical effects and potential magnetic effects and should be applied to the results of future simulations.

Tables A.1-A.4 in Appendix A show the differences between experimental and simulated means, taken directly from profiles such as those shown in figures 7.11-7.14. These are relatively stable for each emittance and detector plane combination, with the exception of the TOF0 x plane.

We are able to recommend a set of effective misalignments by taking an average, weighted

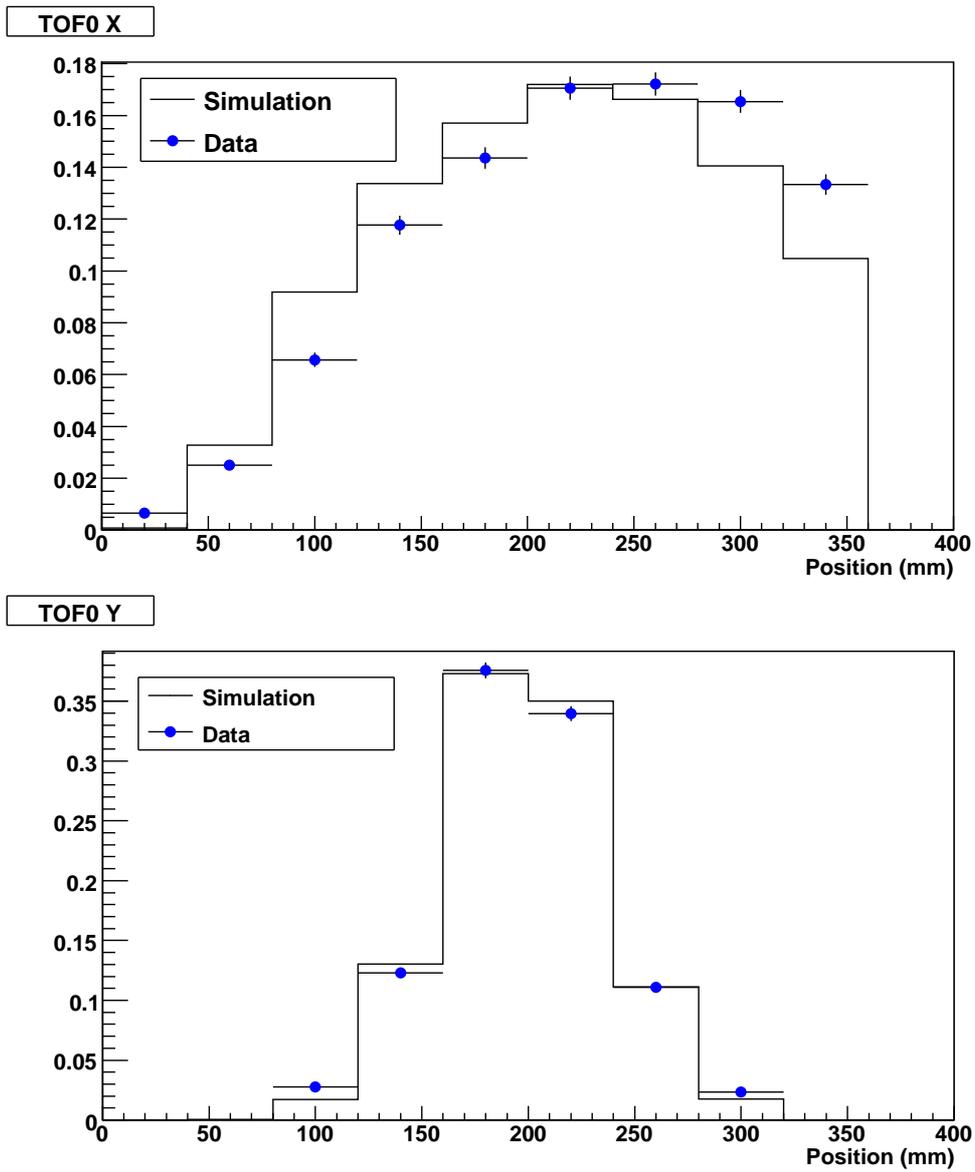


Figure 7.13: Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF0 for the positive polarity (6,200) matrix element.

according to errors for each plane and in the case of the TOF0 x plane, for each (emittance, plane) combination. These are listed in tables 7.5 and 7.6 and compared to the mechanical misalignments proposed in [177].

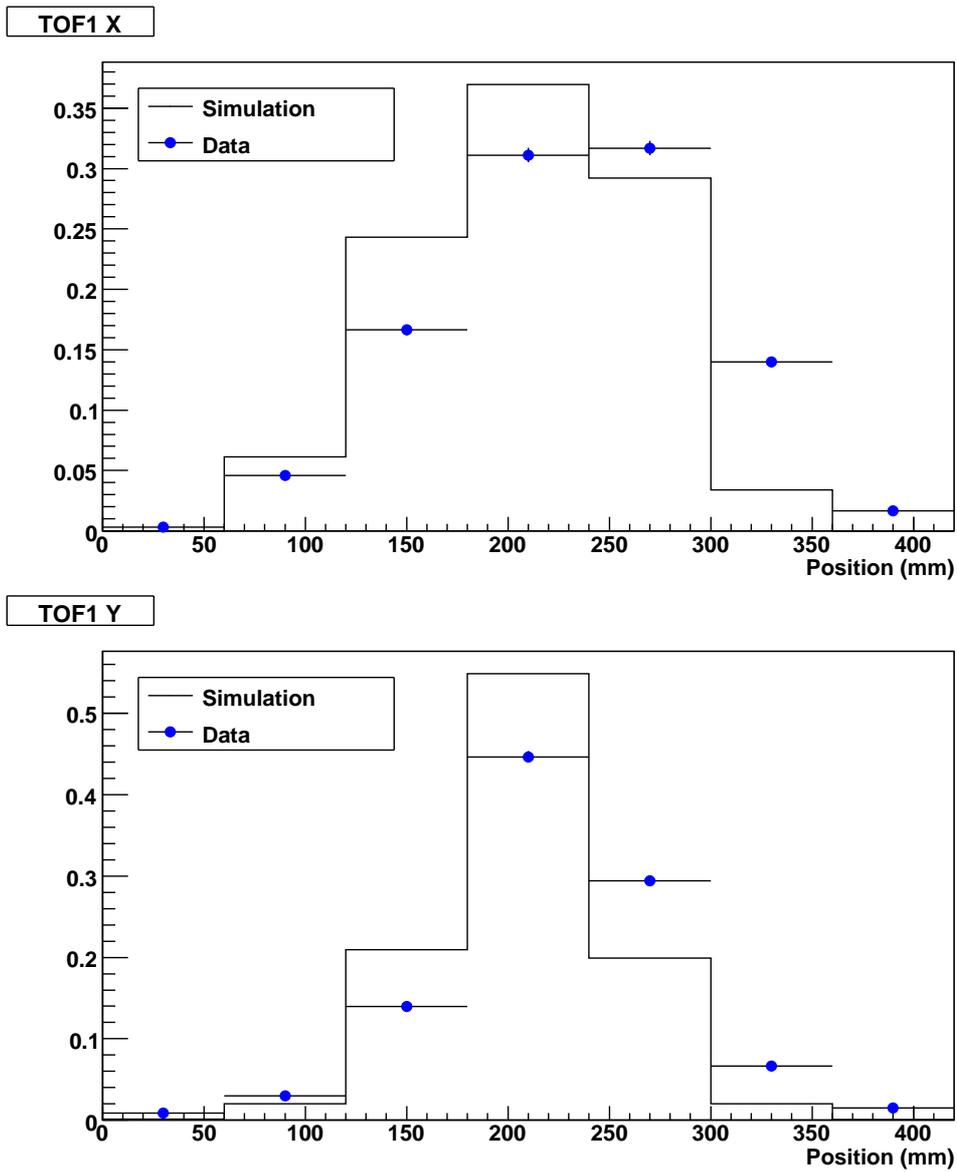


Figure 7.14: Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF1 for the positive polarity (6,200) matrix element.

Detector Plane	Effective Misalignment (mm)	Mechanical Misalignment (mm)
TOF0 X ($\epsilon = 3$ mm rad)	-3.09 ± 1.76	-1.80
TOF0 X ($\epsilon = 6$ mm rad)	-15.3 ± 2.1	-1.80
TOF0 X ($\epsilon = 10$ mm rad)	-13.9 ± 3.2	-1.80
TOF0 Y	1.74 ± 0.38	2.60
TOF1 X	-25.4 ± 0.6	-30.2
TOF1 Y	-16.8 ± 0.4	-17.6

Table 7.5: Effective misalignments which should be applied to MICE negative polarity MICE beamline simulations.

Detector Plane	Effective Misalignment (mm)	Mechanical Misalignment (mm)
TOF0 X ($\epsilon = 3$ mm rad)	7.96 ± 0.89	-1.80
TOF0 X ($\epsilon = 6$ mm rad)	-16.8 ± 1.3	-1.80
TOF0 X ($\epsilon = 10$ mm rad)	-11.7 ± 1.4	-1.80
TOF0 Y	1.47 ± 0.17	2.60
TOF1 X	-27.1 ± 0.3	-30.2
TOF1 Y	-16.0 ± 0.2	-17.6

Table 7.6: Effective misalignments which should be applied to MICE positive polarity MICE beamline simulations.

For TOF0 y , TOF1 x and TOF1 y planes, the effective misalignment is close to the mechanical misalignment. In the case of the x plane of TOF0, the difference between simulated and experimental means varies as a function of both emittance and momentum, causing fluctuations in the effective misalignment. This suggests the presence of non-mechanical magnetic effects close to TOF0, not included in the simulation. Asymmetries in the magnetic axis of the second quadrupole triplet, prior to TOF0 could produce a skew which would shift the experimental mean. This non-mechanical effect, if present, would be included in our effective alignment corrections and could cause them to fluctuate as a function of momentum.

These effective misalignments have been applied to the same simulations, with the results listed in tables A.5-A.8 in Appendix A, greatly reducing the overall discrepancy between simulation and experiment, recommending their use in future analyses. Figures 7.15-7.18 show the same (6,200) negative and positive polarity TOF profiles as before but with the effective alignment corrections we have calculated applied, demonstrating noticeable improvement.

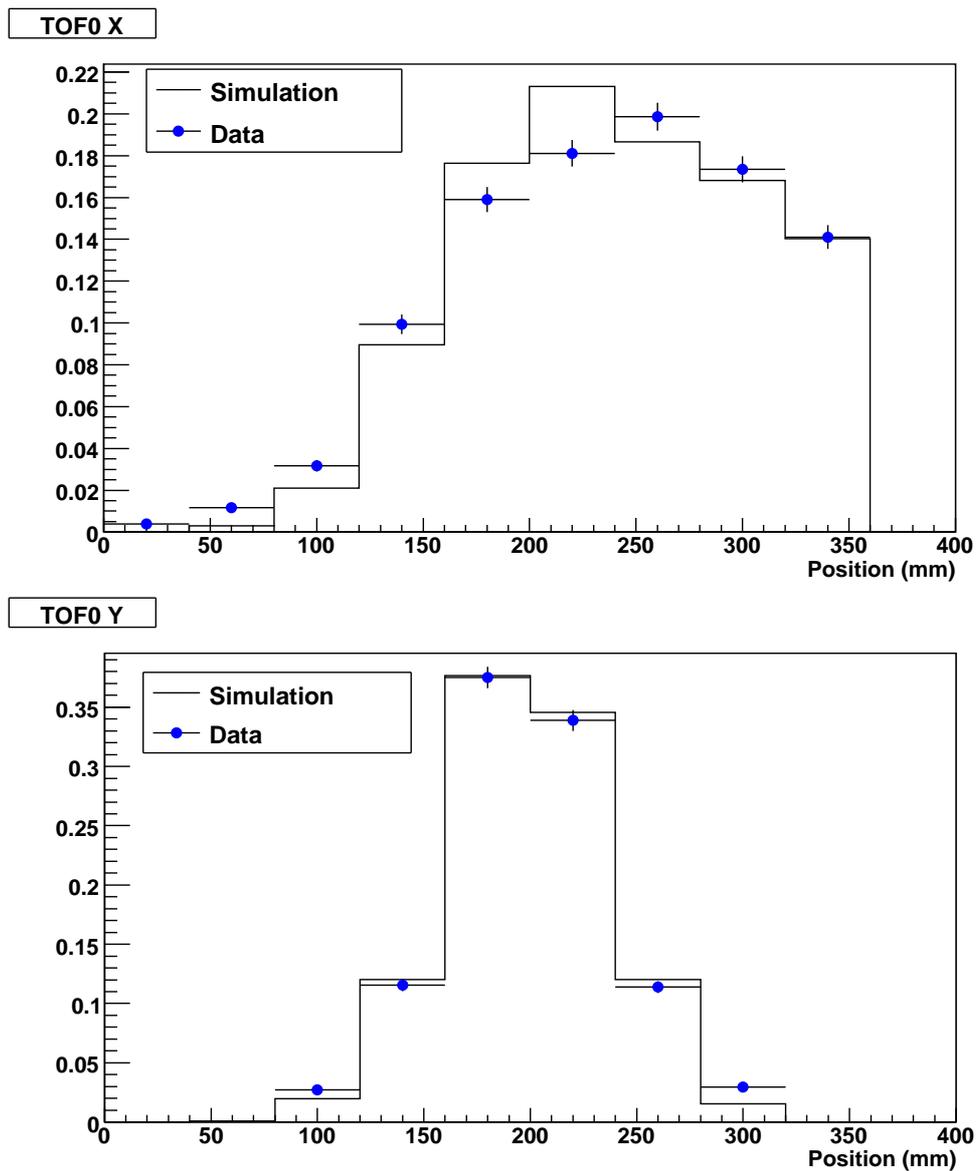


Figure 7.15: Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF0 for the negative polarity (6,200) matrix element with effective misalignments applied. The mechanical misalignment is from reference [177].

7.8.2.2 Profile Widths

In any simulated environment assumptions must be made about physical processes, in our case chiefly decay, energy loss and multiple scattering, which govern particle motion. Multiple scattering angles in particular are frequently sampled from Gaussian distributions. In reality we know that multiple scattering angles are not simply normally distributed and that non-Gaussian tails act to increase the effect of multiple scattering upon a sufficiently large sample. This can cause simulations to underestimate the effects of multiple scattering.

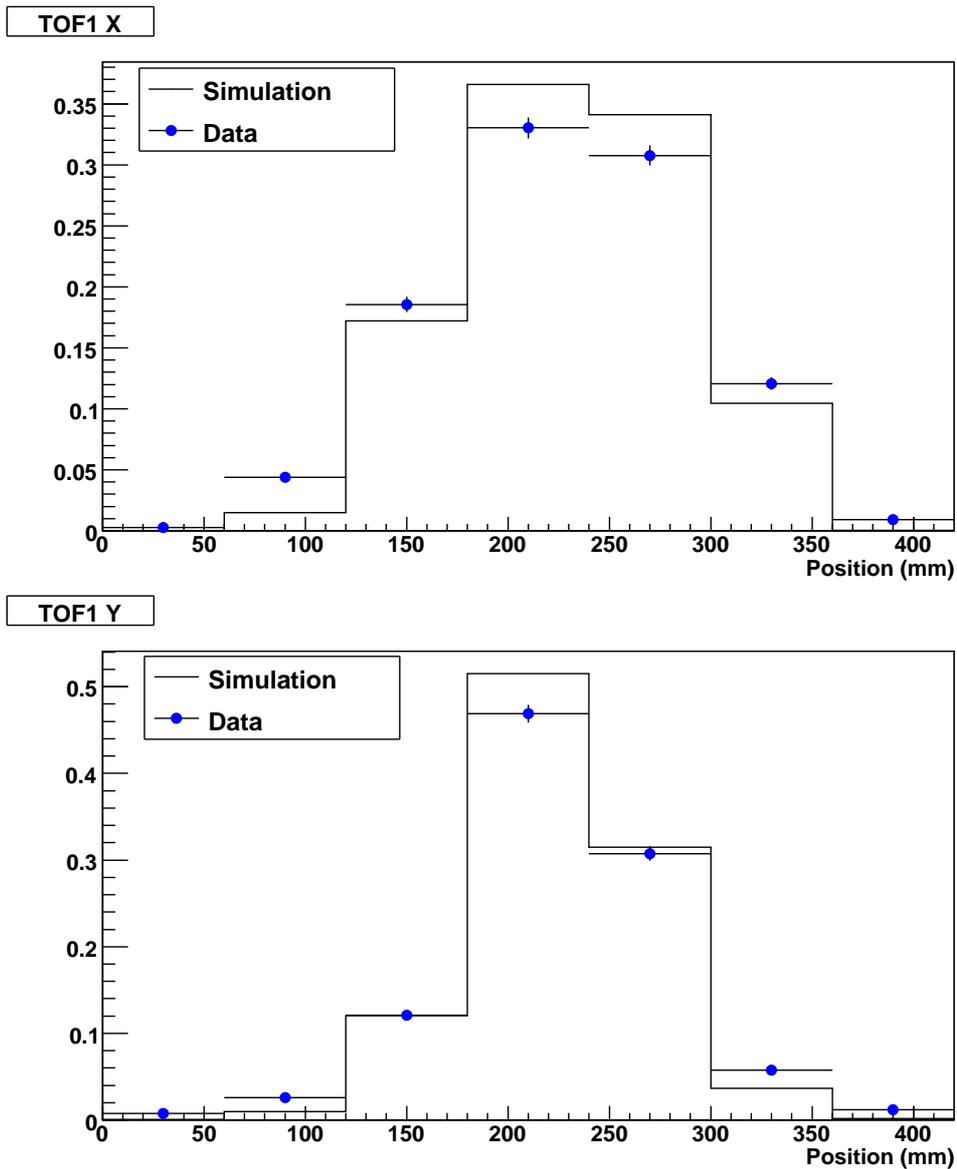


Figure 7.16: Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF1 for the negative polarity (6,200) matrix element with effective misalignments applied. The mechanical misalignment is from reference [177].

In addition, as previously noted, misalignment of the magnets relative to one another may distort the magnetic axis, producing asymmetries in the beam profile. Other possible sources of error include the misalignment of detectors relative to magnets, accuracy of magnetic field maps and fringe field approximations.

This study provides an opportunity to assess the accuracy of the simulation when compared directly with data. A Gaussian fit was applied to both simulated and experimental profiles, yielding standard deviations for each. These are recorded in tables B.1-B.4 in Ap-

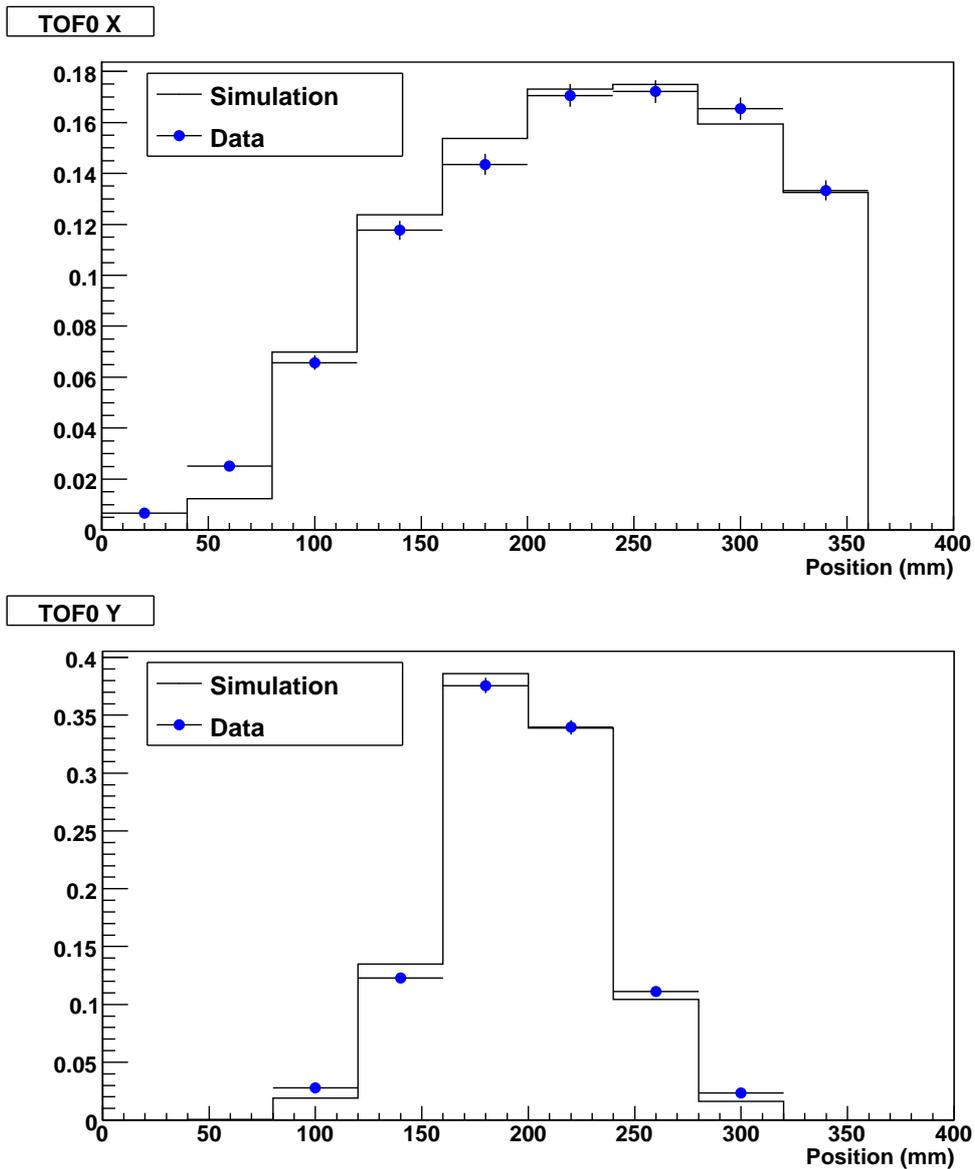


Figure 7.17: Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF0 for the positive polarity (6,200) matrix element with effective misalignments applied. The mechanical misalignment is from reference [177].

pendix B which show the ratio of experimental and simulated standard deviations and means for both positive polarity and negative polarity beams.

As expected, the width of the experimental profiles, described by the standard deviations, is greater than for the simulated profiles in the majority of cases. Since we require that a muon create a hit in both TOF0 and TOF1, relative misalignments between the two TOFs can cause a deficit of slab hits at the extremes of the detectors, as is particularly evident in data collected for the rightmost slab of the x plane of TOF0. This can give rise to inconsistencies

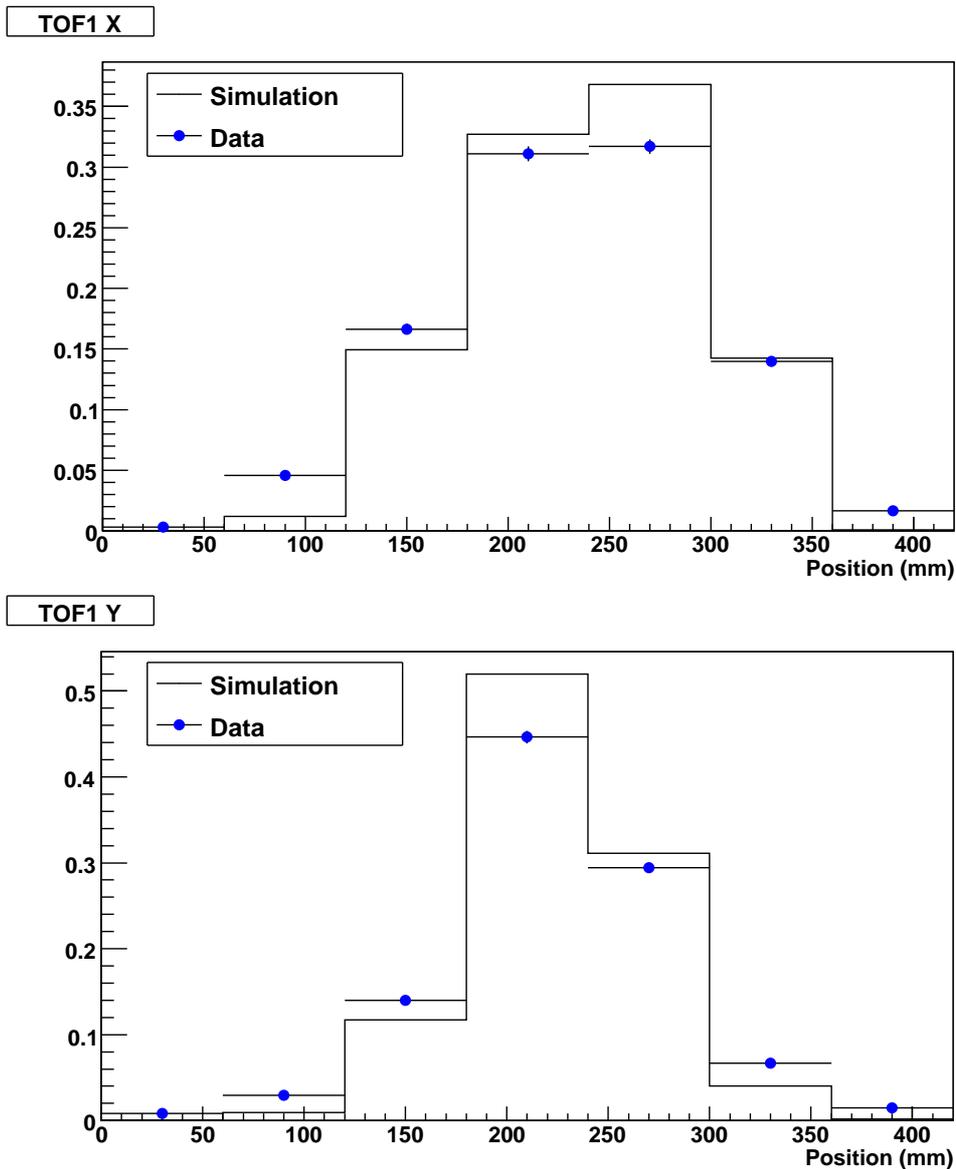


Figure 7.18: Slab by slab comparison between simulation and data for the x (top) and y (bottom) planes of TOF1 for the positive polarity (6,200) matrix element with effective misalignments applied. The mechanical misalignment is from reference [177].

when comparing simulation and data, therefore we apply the effective alignment corrections calculated in the previous section.

The influence of multiple scattering should recede as the momentum of the beam increases, evident in the general trend for σ to decrease as momentum increases, although individual exceptions exist, notably at the TOF0 x plane. In general we find that the measured width of the experimental beam profile is greater than that of the simulation, justifying our assumption that the simulation will underestimate multiple scattering. We have noted

that relative misalignments of the magnets may distort the magnetic axis, creating skew in the beam profiles. In addition potential inaccuracies in the simulated model of the magnetic system, including field maps and models of fringe fields could create noticeable differences between data and simulation. Any deficit in the material along the simulated beamline when compared with the true beamline can also contribute to an underestimate of multiple scattering. It should be noted that when making an emittance measurement in MICE experimentally, we may construct ‘tailor-made’ beams by selecting events in phase space, providing an opportunity to minimise discrepancies, when the first measurements of emittance are carried out.

Comparison of profile widths for experiment and simulation allows quantification of the accuracy of the MICE beamline simulations. The ratio of simulation to data appears relatively stable for a given detector plane within each row of the EM matrix, corresponding to a fixed emittance. From this we infer a percentage ratio between Monte Carlo and data (tables 7.7 and 7.8). We take the average ratio, weighted by the error in the ratio, for each combination of detector plane and emittance setting. This comparison of simulation to data provides a valuable opportunity to tune future simulations and will be used to perform the measurement of emittance at MICE.

Emittance (mm rad)	Profile Width Ratio $\frac{\text{Monte Carlo}}{\text{Data}}$ (%)			
	TOF0 X	TOF0 Y	TOF1 X	TOF1 Y
3	85.7 ± 1.1	79.8 ± 1.1	84.3 ± 0.7	98.4 ± 1.0
6	82.2 ± 1.4	97.2 ± 1.3	82.3 ± 0.8	89.5 ± 1.2
10	93.1 ± 2.6	99.5 ± 1.3	78.4 ± 0.8	89.4 ± 1.1

Table 7.7: Weighted ratio of the width of the Monte Carlo distribution with respect to the data at each TOF plane, for a negative polarity beam.

Emittance (mm rad)	Profile Width Ratio $\frac{\text{Monte Carlo}}{\text{Data}}$ (%)			
	TOF0 X	TOF0 Y	TOF1 X	TOF1 Y
3	78.0 ± 0.5	83.9 ± 0.4	83.7 ± 0.4	96.7 ± 0.5
6	75.6 ± 0.7	94.2 ± 0.6	76.6 ± 0.4	81.4 ± 0.5
10	61.6 ± 0.6	95.9 ± 0.5	78.2 ± 0.4	86.4 ± 0.5

Table 7.8: Weighted ratio of the width of the Monte Carlo distribution with respect to the data at each TOF plane, for a positive polarity beam.

7.9 Conclusions

In addition to the demonstration of ionisation cooling as a viable technology for the realisation of a neutrino factory cooling channel, a key goal of MICE is also to test the validity of its simulations. These simulations will be crucial to the success of future facilities. This chapter has performed the first analysis of the MICE beamline using the luminosity monitor, and provided a crucial comparison of simulated and experimental beam profiles at the TOF planes. Comparison of the width and mean values of beam profiles has enabled numerical quantification of the simulation accuracy and has identified effective misalignments which should be applied to the simulation output. In addition, by measuring the number of muons produced per spill and combining this with an earlier statistical analysis in chapter 5, a predicted running time has been calculated for the Step VI experimental runs. These results will be of significant value to MICE as it moves towards a world-first demonstration of ionisation cooling and realisation of a cooling channel cell.

Chapter 8

Conclusions

The historic discovery of oscillations in the neutrino sector raises questions central to our understanding of the universe. A new generation of experiments will seek to measure the neutrino mixing angles with unprecedented precision¹ and clarify the neutrino mass hierarchy. Work remains to quantify the CP-violating phase (δ) present in neutrino oscillations, which could provide for a fuller explanation of the matter-antimatter asymmetry observed in the universe.

The best experimental venue for future studies is a neutrino factory, described in chapter 2. A significant research and development effort is underway to realise the design of the main components of this future discovery facility. The International Scoping Study was completed in 2007 and provided for a baseline design of the detectors and acceleration system, into which a number of experiments have contributed, for example HARP (target material studies), MERIT (liquid mercury target demonstration) and notably MICE (muon cooling). The International Design Study has produced an Interim Design Report [70] benchmarking physics performance and providing a baseline description, ahead of further work in the coming year towards costing and formally scheduling such a facility.

MICE provides a proof of principle of the cooling channel for a neutrino factory, which exists to reduce the transverse size of the muon beam prior to injection into the acceleration

¹Preliminary results from T2K suggest that θ_{13} is non-zero $(0.03(0.04) < \sin^2 2\theta_{13} < 0.28(0.34)$ for normal (inverted) mass hierarchy) at a confidence level of 2.5σ [58]. More statistics will be needed to either confirm or deny this measurement.

system. It will be the venue for the first ever demonstration of ionisation cooling, a novel cooling technique effective over the short muon lifetime. The goal of the MICE experiment, explained in chapter 3 is to measure fractional change in transverse emittance, a measure of the size and divergence of the beam, of 10% to an accuracy of 1%.

This measurement will be made using two scintillating fibre trackers positioned before and after the MICE cooling channel. The effect of misalignment upon emittance measurement was considered in chapter 5 through simulations studying the impact of relative misalignment between the two detectors for a range of beam conditions. In order to contribute to a total error on fractional emittance change of 0.1% it was found that the spectrometer alignment tolerances were 1 mm in X and Y translation and 3 mrad in rotation. This level of precision should be achievable through careful engineering and surveying.

It is also necessary to study the statistical errors inherent in measuring a quantity which is itself derived from a matrix of covariances. An empirical approach was described, also in chapter 5, using large numbers of simulations exploiting the computing power of the Grid. It was found that due to correlations in the measurements of the two trackers upon a similar set of muons, the number of muons required for a contribution of 0.1% to the total error on fractional change in emittance is of order 10^5 , an order of magnitude less than in the uncorrelated case. This could potentially have a significant impact on the necessary running time of the experiment.

The final configuration of MICE relevant to our discussion thus far will not be immediately realised. MICE will evolve in an incremental manner towards the final Step VI design. Therefore, the experimental configuration will be changing frequently as a function of time. The MICE Configuration Database, described in chapter 4, allows experimenters to determine the geometrical description of the experiment valid at a given time. Surveys of the hall may uncover alignment errors which arise either due to mechanical misalignment during installation or as an effect of normal running (such as magnetic fields for example). The design of the Configuration Database supports addition of corrections to the stored configurations of the experiment to the conditions of the experiment, whilst retaining the previous

information, in order to reproduce analyses. This is achieved through implementation of a bi-temporal paradigm in which transaction time and validity time are considered. The database also supports calibrations and cabling configurations with electronic channel mapping useful for offline analysis. In addition, it records the beamline settings present for each run whilst also allowing the control room operator to reinforce the settings from a previous run, or from a tagged preset configuration. Alarm handler limits are also recorded, in order to assist in tuning of appropriate alarm tolerances.

MICE has currently concluded Step I of its program. During beamline commissioning and throughout Step I, a Luminosity Monitor, designed and built at the University of Glasgow, was used to monitor and analyse particle rates. This new detector grants the experiment an independent means of measuring particle production yields from the MICE target. These measurements can be compared to the beamloss monitors of the host ISIS synchrotron. Additionally the luminosity monitor allows us to measure particle rates in a region which is at a similar angle to the target as the MICE beamline, which can be used to normalise data taken with different conditions to the same luminosity, independent of the beamloss monitors. Data from the luminosity monitor allows us to study particle production in the target and verify existing simulations (chapter 6). Results obtained demonstrated that QGSP_BIC, QGSP_BERT and QGSC_BERT hadronic physics models supported by GEANT4 describe the particle yield data reasonably well, whilst LHEP and QGSC_CHIPS are not favoured.

A first analysis of Step I experimental data, using the luminosity monitor, was presented in chapter 7, demonstrating that the rate in each detector varied in proportion to beamloss. The role played by the time of flight detectors (TOF) in particle identification is of principal importance in Step I for characterising the beamline. The analysis demonstrated that the rate at each TOF predicted by simulation, normalised to luminosity monitor counts, was close to that observed experimentally, providing a useful check of the simulation code and geometries. Further analysis of the simulated and experimental beam profiles at each TOF detector yielded recommended offsets to account for mechanical misalignments of the TOFs and magnetic asymmetries in addition to quantifying the success of the simulation in predicting

the beam profile width due to effects including but not limited to multiple scattering. These analyses were performed on samples of muons, by applying time of flight cuts appropriate to each beam setting.

Combining the statistical analysis of chapters 5 and the data analysis in 7, one may predict the length of time required to run the experiment in order to make a measurement of fractional change in emittance to a statistical error of 0.1% for a beam of momentum 200 MeV/c at the central absorber of Step VI and input emittance settings of 3, 6 and 10 mm rad. The results in table 7.4 demonstrate that this can be achieved in a practical time frame for each setting considered, further reinforcing the viability of the MICE program.

MICE is presently in Step I of its program. In Step V the effect of muon ionisation cooling will be measured for the first time and systematically studied, providing an essential piece of research and development necessary to realise the cooling channel of a future neutrino factory. This exciting future facility, in addition to clarifying the neutrino mass hierarchy and measuring mixing angles and mass squared differences with greatly improved accuracy, will provide unprecedented reach and precision into the parameter space which governs neutrino oscillations and leptonic CP violation.

Appendix A

Misalignment Tables for the Time of Flight Counters

Listed in this appendix are the simulated and experimental means of the time of the beam profiles at the time of flight detectors, before (tables A.1-A.4) and after (tables A.5-A.8) applying the alignment corrections in tables 7.6 and 7.5.

APPENDIX A. MISALIGNMENT TABLES FOR THE TIME OF FLIGHT COUNTERS

	TOF0 \bar{x} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	207.4±0.8	214.0±2.3	6.7±2.4
(3,200)	210.0±1.0	218.8±1.3	8.8±1.6
(3,240)	211.0±0.7	218.8±0.9	7.8±1.2
(6,140)	231.2±1.7	220.8±3.2	-10.4±3.6
(6,200)	250.6±1.9	233.9±1.1	-16.8±2.2
(6,240)	259.8±1.4	241.5±0.9	-18.3±1.7
(10,140)	231.4±1.6	231.6±2.5	0.2±3.0
(10,200)	243.2±1.5	224.0±1.7	-19.2±2.3
(10,240)	252.6±2.3	241.6±1.1	-11.1±2.5
	TOF0 \bar{y} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	195.4±0.5	200.6±1.3	5.2±1.3
(3,200)	194.8±0.5	198.2±0.6	3.4±0.8
(3,240)	196.8±0.3	198.8±0.4	2.0±0.5
(6,140)	197.3±0.3	198.9±0.9	1.5±0.9
(6,200)	198.1±0.4	198.3±0.3	0.1±0.6
(6,240)	197.5±0.3	198.9±0.3	1.5±0.4
(10,140)	197.8±0.3	197.4±0.8	-0.4±0.8
(10,200)	197.4±0.2	199.0±0.3	1.6±0.3
(10,240)	197.9±0.3	199.1±0.3	1.1±0.4

Table A.1: Mean values for positive polarity beam position profiles at TOF0, with no alignment corrections applied.

APPENDIX A. MISALIGNMENT TABLES FOR THE TIME OF FLIGHT COUNTERS

	TOF1 \bar{x} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	245.7±0.4	218.3±1.2	-27.4±1.3
(3,200)	245.0±0.4	218.6±0.6	-26.4±0.7
(3,240)	244.7±0.3	211.6±0.5	-33.1±0.6
(6,140)	251.2±0.7	227.3±1.9	-23.9±2.1
(6,200)	232.8±0.8	213.0±1.1	-19.8±1.3
(6,240)	231.6±0.5	208.4±0.4	-23.1±0.6
(10,140)	251.2±0.8	210.5±1.3	-40.7±1.5
(10,200)	250.8±0.3	224.8±0.6	-25.9±0.7
(10,240)	246.5±0.6	221.0±0.6	-25.5±0.8
	TOF1 \bar{y} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	225.8±0.4	209.9±1.3	-15.9±1.3
(3,200)	225.9±0.4	210.0±0.6	-15.9±0.7
(3,240)	226.1±0.3	209.7±0.4	-16.4±0.5
(6,140)	225.9±0.5	209.7±1.3	-16.1±1.4
(6,200)	225.9±0.6	209.3±0.4	-16.6±0.7
(6,240)	225.3±0.4	209.8±0.3	-15.5±0.5
(10,140)	223.1±0.7	213.1±1.6	-9.9±1.7
(10,200)	226.0±0.3	209.6±0.5	-16.4±0.6
(10,240)	226.3±0.4	210.2±0.3	-16.1±0.5

Table A.2: Mean values for positive polarity beam position profiles at TOF1, with no alignment corrections applied.

APPENDIX A. MISALIGNMENT TABLES FOR THE TIME OF FLIGHT COUNTERS

	TOF0 \bar{x} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	208.8±2.4	208.4±4.0	-0.5±4.7
(3,200)	221.6±1.2	220.9±2.3	-0.6±2.6
(3,240)	231.2±1.9	224.4±2.0	-6.9±2.8
(6,140)	227.2±3.5	225.9±4.8	-1.3±5.9
(6,200)	255.8±2.1	239.5±2.0	-16.3±2.9
(6,240)	263.7±3.2	244.5±1.9	-19.2±3.7
(10,140)	236.8±6.1	234.5±4.5	-2.3±7.6
(10,200)	252.6±3.0	232.4±3.3	-20.3±4.4
(10,240)	260.2±4.1	250.3±3.9	-9.8±5.7
	TOF0 \bar{y} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	201.3±1.7	202.3±2.1	1.0±2.7
(3,200)	201.1±0.8	201.2±1.3	0.2±1.5
(3,240)	201.8±1.4	199.2±1.0	-2.5±1.7
(6,140)	198.7±0.7	201.2±1.4	2.5±1.6
(6,200)	199.3±0.6	200.9±0.7	1.6±1.0
(6,240)	199.9±0.7	202.5±0.6	2.5±0.9
(10,140)	198.8±1.4	203.2±1.3	4.4±2.0
(10,200)	199.1±0.4	201.3±0.7	2.2±0.8
(10,240)	199.1±0.6	200.7±0.6	1.6±0.8

Table A.3: Mean values for negative polarity beam position profiles at TOF0, with no alignment corrections applied.

APPENDIX A. MISALIGNMENT TABLES FOR THE TIME OF FLIGHT COUNTERS

	TOF1 \bar{x} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	246.1±1.3	208.3±2.1	-37.7±2.5
(3,200)	238.3±0.6	214.4±1.2	-23.9±1.3
(3,240)	233.4±0.9	208.7±0.9	-24.7±1.3
(6,140)	248.3±1.7	218.2±3.0	-30.1±3.4
(6,200)	227.2±1.0	205.6±1.0	-21.5±1.4
(6,240)	226.8±1.2	202.0±0.9	-24.8±1.5
(10,140)	243.7±3.7	209.7±2.2	-34.0±4.4
(10,200)	248.5±0.8	221.9±1.4	-26.6±1.6
(10,240)	241.9±1.2	215.9±1.1	-26.0±1.7
	TOF1 \bar{y} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	223.4±1.3	211.8±2.3	-11.6±2.7
(3,200)	226.0±0.5	210.7±1.2	-15.3±1.4
(3,240)	225.4±0.9	210.2±1.0	-15.2±1.3
(6,140)	225.6±1.1	212.8±2.1	-12.8±2.4
(6,200)	227.4±0.8	208.1±0.8	-19.4±1.1
(6,240)	227.7±1.0	208.6±0.7	-19.1±1.2
(10,140)	223.6±3.3	211.2±2.8	-12.4±4.3
(10,200)	225.1±0.7	210.1±1.1	-15.1±1.3
(10,240)	227.2±0.7	209.7±0.6	-17.5±0.9

Table A.4: Mean values for negative polarity beam position profiles at TOF1, with no alignment corrections applied.

APPENDIX A. MISALIGNMENT TABLES FOR THE TIME OF FLIGHT COUNTERS

	TOF0 \bar{x} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	207.4±0.8	202.7±2.2	-4.6±2.3
(3,200)	210.0±1.0	207.9±1.3	-2.2±1.6
(3,240)	211.0±0.7	207.8±0.9	-3.2±1.2
(6,140)	231.2±1.7	237.7±3.1	6.5±3.5
(6,200)	250.6±1.9	247.8±1.4	-2.8±2.3
(6,240)	259.8±1.4	253.5±1.1	-6.3±1.8
(10,140)	231.4±1.6	243.3±4.4	12.0±4.7
(10,200)	243.2±1.5	235.3±1.5	-7.9±2.1
(10,240)	252.6±2.3	251.2±1.1	-1.4±2.5
	TOF0 \bar{y} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	195.4±0.5	199.6±1.3	4.2±1.3
(3,200)	194.8±0.5	199.3±0.6	4.5±0.8
(3,240)	196.8±0.3	199.5±0.4	2.8±0.5
(6,140)	197.3±0.3	196.8±0.9	-0.5±1.0
(6,200)	198.1±0.4	196.7±0.3	-1.5±0.6
(6,240)	197.5±0.3	197.5±0.3	0.0±0.4
(10,140)	197.8±0.3	195.5±0.8	-2.2±0.8
(10,200)	197.4±0.2	197.5±0.3	0.1±0.4
(10,240)	197.9±0.3	197.7±0.3	-0.3±0.4

Table A.5: Mean values for positive polarity beam position profiles at TOF0, with corrections applied for effective misalignments.

APPENDIX A. MISALIGNMENT TABLES FOR THE TIME OF FLIGHT COUNTERS

	TOF1 \bar{x} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	245.7±0.4	239.1±1.2	-6.6±1.3
(3,200)	245.0±0.4	239.1±1.2	-5.9±1.3
(3,240)	244.7±0.3	237.5±0.5	-7.1±0.6
(6,140)	251.2±0.7	246.7±1.7	-4.5±1.8
(6,200)	232.8±0.8	233.7±0.6	0.9±1.0
(6,240)	231.6±0.5	230.7±0.5	-0.8±0.7
(10,140)	251.2±0.8	248.9±1.4	-2.3±1.7
(10,200)	250.8±0.3	243.4±0.6	-7.3±0.7
(10,240)	246.5±0.6	238.8±0.5	-7.7±0.8
	TOF1 \bar{y} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	225.8±0.4	226.7±1.3	0.9±1.3
(3,200)	225.9±0.4	226.4±0.6	0.4±0.7
(3,240)	226.1±0.3	225.9±0.4	-0.2±0.5
(6,140)	225.9±0.5	225.6±1.3	-0.3±1.4
(6,200)	225.9±0.6	225.8±0.4	-0.1±0.7
(6,240)	225.3±0.4	225.5±0.3	0.2±0.5
(10,140)	223.1±0.7	232.4±1.6	9.3±1.7
(10,200)	226.0±0.3	225.5±0.5	-0.5±0.6
(10,240)	226.3±0.4	225.7±0.3	-0.6±0.5

Table A.6: Mean values for positive polarity beam position profiles at TOF1, with corrections applied for effective misalignments.

APPENDIX A. MISALIGNMENT TABLES FOR THE TIME OF FLIGHT COUNTERS

	TOF0 \bar{x} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	208.8±2.4	212.6±4.1	3.7±4.8
(3,200)	221.6±1.2	222.9±2.4	1.4±2.7
(3,240)	231.2±1.9	228.6±2.0	-2.6±2.8
(6,140)	227.2±3.5	243.2±5.2	16.0±6.2
(6,200)	255.8±2.1	251.4±2.2	-4.4±3.1
(6,240)	263.7±3.2	258.3±2.1	-5.4±3.8
(10,140)	236.8±6.1	253.2±9.7	16.4±11.4
(10,200)	252.6±3.0	240.9±4.7	-11.8±5.6
(10,240)	260.2±4.1	263.3±3.7	3.2±5.6
	TOF0 \bar{y} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	201.3±1.7	201.3±2.1	-0.1±2.7
(3,200)	201.1±0.8	199.4±1.3	-1.7±1.5
(3,240)	201.8±1.4	197.6±1.0	-4.2±1.7
(6,140)	198.7±0.7	199.7±1.4	1.0±1.6
(6,200)	199.3±0.6	199.1±0.7	-0.2±1.0
(6,240)	199.9±0.7	200.5±0.6	0.6±0.9
(10,140)	198.8±1.4	201.1±1.4	2.3±2.0
(10,200)	199.1±0.4	199.6±0.7	0.5±0.8
(10,240)	199.1±0.6	198.7±0.6	-0.3±0.9

Table A.7: Mean values for negative polarity beam position profiles at TOF0, with corrections applied for effective misalignments.

	TOF1 \bar{x} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	246.1±1.3	231.7±2.0	-14.4±2.4
(3,200)	238.3±0.6	236.6±1.2	-1.7±1.3
(3,240)	233.4±0.9	228.1±0.9	-5.3±1.3
(6,140)	248.3±1.7	235.4±2.5	-13.0±3.0
(6,200)	227.2±1.0	226.6±1.1	-0.6±1.5
(6,240)	226.8±1.2	234.4±1.0	7.5±1.5
(10,140)	243.7±3.7	244.4±2.5	0.7±4.5
(10,200)	248.5±0.8	241.4±1.3	-7.2±1.6
(10,240)	241.9±1.2	233.3±1.0	-8.6±1.6
	TOF1 \bar{y} (mm)		
Matrix Element	Data	Sim	Difference
(3,140)	223.4±1.3	228.3±2.2	4.9±2.6
(3,200)	226.0±0.5	227.6±1.2	1.5±1.3
(3,240)	225.4±0.9	227.2±1.0	1.8±1.3
(6,140)	225.6±1.1	230.0±2.1	4.4±2.4
(6,200)	227.4±0.8	225.4±0.9	-2.1±1.2
(6,240)	227.7±1.0	225.2±0.7	-2.5±1.2
(10,140)	223.6±3.3	232.8±3.0	9.1±4.5
(10,200)	225.1±0.7	226.7±1.1	1.6±1.3
(10,240)	227.2±0.7	225.4±0.6	-1.8±1.0

Table A.8: Mean values for negative polarity beam position profiles at TOF1, with corrections applied for effective misalignments.

Appendix B

Beam Profile Widths at the Time of Flight Counters

This appendix contains the ratios of simulated and experimental beam profile widths (tables B.1-B.4) relevant to the discussion in chapter 7 but included here to aid the flow of discussion.

	TOF0 σ_x (mm)		
Matrix Element	Data	Sim	Ratio
(3,140)	91.3±1.9	93.6±4.3	1.02±0.05
(3,200)	93.5±0.9	71.9±1.1	0.769±0.014
(3,240)	86.8±1.3	87.9±1.7	1.01±0.02
(6,140)	119±3	85.4±2.7	0.719±0.030
(6,200)	84.5±1.5	72.6±1.4	0.860±0.022
(6,240)	90.1±2.1	75.9±1.2	0.842±0.024
(10,140)	108±5	129±10	1.19±0.11
(10,200)	125±4	123±5	0.987±0.048
(10,240)	118±3	105±3	0.882±0.033
	TOF0 σ_y (mm)		
Matrix Element	Data	Sim	Ratio
(3,140)	72.1±2.0	60.6±1.2	0.841±0.028
(3,200)	72.6±0.9	58.7±0.8	0.809±0.015
(3,240)	73.6±1.6	56.4±0.5	0.766±0.018
(6,140)	42.8±0.6	40.1±1.3	0.936±0.034
(6,200)	40.7±0.6	39.1±0.6	0.960±0.019
(6,240)	34.3±0.6	34.6±0.4	1.00±0.02
(10,140)	43.8±1.3	42.9±1.2	0.980±0.040
(10,200)	34.0±0.3	32.9±0.6	0.966±0.019
(10,240)	37.0±0.5	38.2±0.5	1.03±0.02

Table B.1: Values, in mm, of simulated and experimental standard deviations for negative polarity beam profiles at TOF0. The ratios of simulated and experimental values have been listed.

	TOF1 σ_x (mm)		
Matrix Element	Data	Sim	Ratio
(3,140)	68.7±1.0	57.3±1.2	0.834±0.021
(3,200)	63.0±0.4	52.5±0.6	0.834±0.011
(3,240)	60.4±0.6	51.5±0.4	0.854±0.011
(6,140)	88.8±1.2	68.5±1.7	0.772±0.022
(6,200)	64.5±0.6	51.3±0.5	0.795±0.011
(6,240)	62.2±0.8	56.0±0.7	0.900±0.015
(10,140)	99.7±3.2	72.8±1.4	0.730±0.027
(10,200)	74.6±0.6	57.3±0.7	0.769±0.011
(10,240)	74.6±0.9	60.6±0.5	0.813±0.012
	TOF1 σ_y (mm)		
Matrix Element	Data	Sim	Ratio
(3,140)	69.0±1.1	63.9±1.3	0.927±0.025
(3,200)	60.6±0.5	60.5±0.9	0.998±0.016
(3,240)	58.5±0.8	58.2±0.6	0.995±0.017
(6,140)	65.6±1.0	59.5±1.6	0.906±0.028
(6,200)	50.6±0.7	44.9±0.7	0.887±0.019
(6,240)	51.2±0.9	46.0±0.6	0.899±0.019
(10,140)	94.9±2.8	84.3±2.0	0.889±0.034
(10,200)	60.3±0.6	55.3±0.8	0.917±0.016
(10,240)	44.6±0.7	38.9±0.5	0.870±0.017

Table B.2: Values, in mm, of simulated and experimental standard deviations for negative polarity beam profiles at TOF1. The ratios of simulated and experimental values have been listed.

	TOF0 σ_x (mm)		
Matrix Element	Data	Sim	Ratio
(3,140)	105±1	83.3±1.5	0.791±0.015
(3,200)	106±1	84.2±0.8	0.793±0.010
(3,240)	110±1	84.7±0.6	0.772±0.007
(6,140)	131±2	78.5±1.4	0.600±0.013
(6,200)	96.7±1.3	90.6±1.0	0.937±0.017
(6,240)	103±1	80.2±0.6	0.780±0.009
(10,140)	123±2	122±4	0.991±0.039
(10,200)	144±2	92.2±0.9	0.640±0.009
(10,240)	127±2	72.9±0.5	0.572±0.009
	TOF0 σ_y (mm)		
Matrix Element	Data	Sim	Ratio
(3,140)	71.6±0.5	60.5±0.7	0.845±0.012
(3,200)	66.0±0.5	55.1±0.3	0.835±0.008
(3,240)	64.2±0.3	53.9±0.2	0.840±0.006
(6,140)	43.0±0.3	41.1±0.7	0.956±0.018
(6,200)	40.4±0.4	38.2±0.2	0.945±0.011
(6,240)	36.9±0.2	34.6±0.2	0.937±0.008
(10,140)	45.3±0.3	42.6±0.5	0.939±0.013
(10,200)	34.9±0.1	33.3±0.3	0.956±0.008
(10,240)	38.7±0.3	37.6±0.2	0.972±0.009

Table B.3: Values, in mm, of simulated and experimental standard deviations for positive polarity beam profiles at TOF0. The ratios of simulation and data have been listed.

	TOF1 σ_x (mm)		
Matrix Element	Data	Sim	Ratio
(3,140)	70.1±0.3	57.9±0.7	0.826±0.011
(3,200)	65.7±0.3	61.4±0.5	0.934±0.009
(3,240)	63.7±0.2	51.9±0.2	0.816±0.004
(6,140)	90.4±0.5	71.0±1.2	0.785±0.014
(6,200)	67.6±0.5	51.4±0.3	0.761±0.007
(6,240)	64.4±0.3	48.4±0.2	0.751±0.005
(10,140)	96.0±0.7	70.0±0.7	0.729±0.009
(10,200)	74.4±0.3	57.9±0.3	0.777±0.005
(10,240)	74.3±0.4	60.7±0.3	0.817±0.006
	TOF1 σ_y (mm)		
Matrix Element	Data	Sim	Ratio
(3,140)	68.9±0.3	63.6±0.9	0.923±0.014
(3,200)	61.1±0.4	59.2±0.4	0.968±0.009
(3,240)	57.1±0.2	55.8±0.3	0.978±0.007
(6,140)	67.4±0.4	59.0±1.0	0.875±0.015
(6,200)	55.1±0.5	44.9±0.3	0.815±0.010
(6,240)	55.9±0.4	44.7±0.3	0.801±0.007
(10,140)	93.3±0.6	80.0±1.0	0.858±0.012
(10,200)	60.7±0.3	54.9±0.3	0.903±0.007
(10,240)	49.3±0.4	39.6±0.3	0.804±0.009

Table B.4: Values, in mm, of simulated and experimental standard deviations for positive polarity beam profiles at TOF1. The ratios of simulation and data have been listed.

Bibliography

- [1] F. A. Scott. Energy Spectrum of the Beta-Rays of Radium E. *Phys. Rev.*, 48(5):391–395, Sep 1935.
- [2] W. Pauli. Dear radioactive ladies and gentlemen. *Phys.Today*, 31N9:27, 1978. Translation with originally provided by Pauli for a conference in Tübingen in 1930.
- [3] H. Bethe and R. Peierls. The ‘neutrino’. *Nature*, 133:532, 1934.
- [4] Frederick Reines and Clyde L. Cowan. The neutrino. *Nature*, 178:446–449, 1956.
- [5] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105(4):1413–1415, Feb 1957.
- [6] BR Martin, G Shaw. *Particle Physics*. John Wiley & Sons, Second edition, 2006.
- [7] M. Goldhaber, L. Grodzins, and A.W. Sunyar. Helicity of Neutrinos. *Phys.Rev.*, 109:1015–1017, 1958.
- [8] R.P. Feynman and Murray Gell-Mann. Theory of Fermi interaction. *Phys.Rev.*, 109:193–198, 1958.
- [9] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Sov.Phys.JETP*, 26:984–988, 1968.
- [10] G. Danby, J.M. Gaillard, Konstantin A. Goulianos, L.M. Lederman, Nari B. Mistry, et al. Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys.Rev.Lett.*, 9:36–44, 1962.

- [11] Martin L. Perl, G.S. Abrams, A. Boyarski, Martin Breidenbach, D. Briggs, et al. Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation. *Phys.Rev.Lett.*, 35:1489–1492, 1975.
- [12] K. Kodama et al. Observation of tau neutrino interactions. *Phys.Lett.*, B504:218–224, 2001. http://www.fnal.gov/pub/presspass/press_releases/donut.html.
- [13] ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration, SLD Collaboration, LEP Electroweak Working Group, SLD Electroweak Group, SLD Heavy Flavour Group. Precision electroweak measurements on the Z resonance. *Phys.Rept.*, 427:257–454, 2006.
- [14] R Davis. A Review of the Homestake Solar Neutrino Experiment. *Progress in Particle and Nuclear Physics*, 32, 1994.
- [15] The Super-Kamiokande Collaboration. Solar neutrino results in Super-Kamiokande-III. *Phys.Rev.*, D83:052010, 2011.
- [16] The SNO Collaboration. Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory. *Phys.Rev.Lett.*, 87:071301, 2001.
- [17] E Kearns. The Emerging Picture of Neutrino Oscillation. *Presentation at NuFact09: 11th International Workshop on Neutrino Factories, Superbeams and Beta Beams, Illinois Institute of Technology*, 2009. http://vmsstreamer1.fnal.gov/VMS_Site.03/Lectures/NuFact/presentations/090720Kearns.pdf.
- [18] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Sov.Phys.JETP*, 26:984–988, 1968.
- [19] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog.Theor.Phys.*, 28:870–880, 1962.

- [20] A. Cervera, A. Donini, M.B. Gavela, J.J. Gomez Cadenas, P. Hernandez, et al. Golden measurements at a neutrino factory. *Nucl.Phys.*, B579:17–55, 2000.
- [21] B Kayser. Neutrino Oscillation Studies with Atmospheric Neutrinos. In FJP Soler, CD Froggatt and F Muheim, editor, *Neutrinos in Particle Physics, Astrophysics and Cosmology*. CRC Press, 2006.
- [22] S.P. Mikheev and A.Yu. Smirnov. Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos. *Sov.J.Nucl.Phys.*, 42:913–917, 1985.
- [23] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys.Rev.*, D17:2369–2374, 1978.
- [24] E.W. Otten and C. Weinheimer. Neutrino mass limit from tritium beta decay. *Rept.Prog.Phys.*, 71:086201, 2008.
- [25] Christian Weinheimer. Direct neutrino mass search. *International School of Physics Enrico Fermi: Course 152: Neutrino Physics, Varenna, Lake Como, Italy, 23 Jul - 2 Aug 2002*, 2002. arXiv:hep-ex/0210050.
- [26] Claude Amsler et al. Review of Particle Physics. *Phys.Lett.*, B667:1–1340, 2008.
- [27] Ettore Majorana. Teoria simmetrica dell'elettrone e del positrone (Translation: Theory of the Symmetry of Electrons and Positrons). *Nuovo Cim.*, 14:171–184, 1937.
- [28] K Zuber. Neutrinoless Double Beta Decay. In FJP Soler, CD Froggatt and F Muheim, editor, *Published in Neutrinos in Particle Physics, Astrophysics and Cosmology*. CRC Press, 2009.
- [29] The CHOOZ Collaboration. Search for neutrino oscillations on a long base-line at the CHOOZ nuclear power station. *Eur. Phys. J.*, C27:331–374, 2003.
- [30] Double Chooz Collaboration. Double Chooz Official Site. <http://doublechooz.in2p3.fr/Scientific/scientific.php>.
- [31] S. Abe et al. Precision Measurement of Neutrino Oscillation Parameters with KamLAND. *Phys.Rev.Lett.*, 100:221803, 2008.

- [32] T. Araki, S. Enomoto, K. Furuno, Y. Gando, K. Ichimura, et al. Experimental investigation of geologically produced antineutrinos with KamLAND. *Nature*, 436:499–503, 2005.
- [33] John N. Bahcall, Aldo M. Serenelli, and Sarbani Basu. New solar opacities, abundances, helioseismology, and neutrino fluxes. *Astrophys.J.*, 621:L85–L88, 2005.
- [34] J. N. Abdurashitov, T. J. Bowles, M. L. Cherry, B. T. Cleveland, R. Davis, S. R. Elliott, V. N. Gavrin, S. V. Girin, V. V. Gorbachev, T. V. Ibragimova, A. V. Kalikhov, N. G. Khairnasov, T. V. Knodel, K. Lande, I. N. Mirmov, J. S. Nico, A. A. Shikhin, W. A. Teasdale, E. P. Veretenkin, V. M. Vermul, D. L. Wark, P. S. Wildenhain, J. F. Wilkerson, V. E. Yants, and G. T. Zatsepin. Measurement of the Solar Neutrino Capture Rate by SAGE and Implications for Neutrino Oscillations in Vacuum. *Phys. Rev. Lett.*, 83(23):4686–4689, Dec 1999.
- [35] Till Kirsten. Retrospect of GALLEX/GNO. *J.Phys.Conf.Ser.*, 120:052013, 2008.
- [36] Y Koshio. PhD Thesis. <http://www-sk.icrr.u-tokyo.ac.jp/sk/ykphd/chap3-3.html>.
- [37] Super-Kamiokande Collaboration. Solar neutrino measurements in Super-Kamiokande-I. *Phys.Rev.*, D73:112001, 2006.
- [38] The SNO Collaboration. Determination of the ν_e and Total ^8B Solar Neutrino Fluxes with the Sudbury Neutrino Observatory Phase I Data Set. *Phys. Rev. C* 75 045502, 2007.
- [39] John N. Bahcall, M.H. Pinsonneault, and Sarbani Basu. Solar models: Current epoch and time dependences, neutrinos, and helioseismological properties. *Astrophys.J.*, 555:990–1012, 2001.
- [40] Edited by FJP Soler, CD Froggatt, F Muheim. *Neutrinos in Particle Physics, Astrophysics and Cosmology*. CRC Press, 2006.

- [41] The MINOS Collaboration. Search for ν_μ to ν_e oscillations in the MINOS experiment. *Int.J.Mod.Phys.*, A26:179–189, 2011.
- [42] WWM Allison et al. The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton year exposure of Soudan2. *Physics Letters B* 449 137, 1999.
- [43] M. Ambrosio et al. Status report of the MACRO experiment for the year 2001. *hep-ex/0206027*, 2002.
- [44] The Super-Kamiokande Collaboration. Three flavor neutrino oscillation analysis of atmospheric neutrinos in Super-Kamiokande. *Phys. Rev. D*, 74(3):032002, Aug 2006.
- [45] The K2K Collaboration. Measurement of Neutrino Oscillation by the K2K Experiment. *Phys.Rev.*, D74:072003, 2006.
- [46] A. Piepke. Final results from the Palo Verde neutrino oscillation experiment. *Prog.Part.Nucl.Phys.*, 48:113–121, 2002.
- [47] H.A. Tanaka. The Search for $\nu_\mu \rightarrow \nu_e$ oscillations at MiniBooNE. *FERMILAB-CONF-07-345-E*, page 004, 2007.
- [48] J.J. Napolitano, S.J. Freedman, G.T. Garvey, M.C. Green, K.T. Lesko, et al. Construction and Performance of a Large Area Liquid Scintillator Cosmic Ray Anticoincidence Detector. *Nucl.Instrum.Meth.*, A274:152–164, 1989.
- [49] Thomas Schwetz. The LSND puzzle in the light of MiniBooNE results. *CERN-PH-TH-2008-105*, 2008. arXiv:hep-ph/0805.2234.
- [50] Sergei N. Gninenko. A resolution of puzzles from the LSND, KARMEN, and MiniBooNE experiments. *Phys.Rev.*, D83:015015, 2011.
- [51] A.A. Aguilar-Arevalo et al. A Search for Electron Antineutrino Appearance at the $\Delta m^{*2} \sim 1\text{-eV}^{*2}$ Scale. *Phys.Rev.Lett.*, 103:111801, 2009.
- [52] The MiniBooNE Collaboration. Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations. *Phys.Rev.Lett.*, 105:181801, 2010.

- [53] Alessandro Menegolli. Status of the ICARUS T600 detector at the LNGS. *J.Phys.Conf.Ser.*, 203:012107, 2010.
- [54] Heather Ray. OscSNS: Precision Neutrino Measurements at the Spallation Neutron Source. *J.Phys.Conf.Ser.*, 136:022029, 2008.
- [55] J. Marteau. The OPERA experiment: a direct search of the $\nu_\mu \longrightarrow \nu_\tau$ oscillations. *11th ICATPP Conference on Astroparticle, Particle, Space Physics, Detectors and Medical Physics Applications*, 2009.
- [56] OPERA Collaboration. Observation of a first ν_τ candidate in the OPERA experiment in the CNGS beam. *Phys.Lett.*, B691:138–145, 2010.
- [57] Danuta Kielczewska. The T2K experiment at J-PARC. *Acta Phys.Polon.*, B41:1565–1578, 2010.
- [58] K Abe et al. Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam. *arXiv:hep-ex/1106.2822*, 2011.
- [59] M Messier. The NOvA Experiment at Fermilab. *Presentation given at the 34th International Conference on High Energy Physics (ICHEP08)*, 2008.
- [60] The EUROnu Collaboration. EUROnu Design Study. <http://www.euronu.org/>.
- [61] Jean-Eric Campagne. The SPL-Frejus physics potential. *Nucl.Phys.Proc.Suppl.*, 155:185–186, 2006.
- [62] Jelena Maricic. The Long Baseline Neutrino Oscillation Experiment at DUSEL. *J.Phys.Conf.Ser.*, 203:012109, 2010.
- [63] K. Nakamura. HYPER-KAMIOKANDE: A next generation water Cherenkov detector for a nucleon decay experiment. *Front.Phys.*, 35:359–363, 2000.
- [64] F. Dufour. T2KK, a Mton water Cherenkov detector in the line of the T2K beam. *Nucl.Phys.Proc.Suppl.*, 188:201–203, 2009.

- [65] S. Geer. Neutrino beams from muon storage rings: Characteristics and physics potential. *Phys. Rev. D*, 57(11):6989–6997, Jun 1998.
- [66] P Zucchelli. A Novel Concept for a $\bar{\nu}_e/\nu_e$ Neutrino Factory: The Beta-Beam. *Phys Lett B*, 532:166–172, 2002.
- [67] Thomas Schwetz, Mariam Tortola, and J.W.F. Valle. Global neutrino data and recent reactor fluxes: status of three-flavour oscillation parameters. *New J.Phys.*, 13:063004, 2011.
- [68] D.G. Koshkarev. Proposal for a Decay Ring to Produce Intense Secondary Particle Beams at the SPS, 1974. CERN Internal Report CERN/ISR-DI/74-62 (1974).
- [69] D. Cline and D. Neuffer. A Muon Storage Ring for Neutrino Oscillations Experiments. *AIP Conf.Proc.*, 68:846–847, 1980.
- [70] The International Design Study for a Neutrino Factory (IDS-NF) Collaboration. Interim Design Report. <https://www.ids-nf.org/wiki/FrontPage/Documentation/IDR>.
- [71] S Ozaki, RB Palmer, MS Zisman, JC Gallardo. Feasibility Study II of a Muon Based Neutrino Source. *Technical Report BNL-52623, Brookhaven National Laboratory, Upton, NY*, 2001.
- [72] P Huber, M Lindner, M Rolinec, W Winter. Optimisation of a Neutrino Factory Oscillation Experiment. *Phys. Rev.*, D74:073003, 2006.
- [73] EURISOL Collaboration. EURISOL Design Study. <http://www.eurisol.org/site02/index.php>.
- [74] M Benedikt. Beta-Beam Baseline Design Study Within EURISOL. http://beta-beam.web.cern.ch/beta-beam/Presentations/ISS_CERN_MBenedikt.pdf.
- [75] Fermilab. Long Baseline Neutrino Experiment. <http://lbne.fnal.gov/>.
- [76] C Densham, M Dracos, A Longhin, M Zito. Baseline scenario for the EUROnu superbeam studies. *EUROnu WP2 Note 09-05*, 2009.

- [77] FJP Soler, C Le Couedic, F Breton. Pion Production for Neutrino Factory - Challenges. *Proceedings of the 12th International Workshop on Neutrino Factories, Super Beams and Beta Beams, Mumbai, India 20-25 October 2010*, 2010.
- [78] J. Pasternak. Status of the Proton Driver. Presentation to EUROnu on 18 January 2011, <http://alturl.com/hqpuu>.
- [79] Project X Collaboration. Reference Design Report. <http://projectx-docdb.fnal.gov/cgi-bin/ShowDocument?docid=776>.
- [80] K T McDonald et al. The MERIT High-Power Target Experiment at the CERN PS. *Proceedings of IPAC 10, Kyoto, Japan*, 2010.
- [81] J.R.J. Bennett, C.N. Booth, R.A. Brownsword, C.J. Densham, T.R. Edgecock, et al. Solid target studies in the UK. *Nucl.Phys.Proc.Suppl.*, 155:291–292, 2006.
- [82] O Caretta, C Densham, P Loveridge. The Development of Fluidised Powder Target Technology for a Neutrino Factory or Muon Collider. http://www.phy.princeton.edu/mumu/target/Densham/densham_092210.pdf.
- [83] J.R.J. Bennett. A high power, radiation cooled, rotating toroidal target for neutrino production. *AIP Conf.Proc.*, 542:253–257, 2000.
- [84] Tristan Davenne, Otto Caretta, Peter Loveridge. A Static Packed Bed Target. <http://alturl.com/fydrf>.
- [85] JRJ Bennett et al. Solid Target Studies for T2K and for a Neutrino Factory. *Beams for European Neutrino Experiments (BENE) 06*, 2006. <http://alturl.com/udij3>.
- [86] Thomas Edgecock, John Back, Roger Bennett, Christopher Booth, Goran Skoro, et al. Overview of Solid Target Studies for a Neutrino Factory. *Proceedings of IPAC 2010*, 2010. <http://alturl.com/w8vm8>.
- [87] J Back. Target Presentation. *ECFA Neutrino Panel Meeting*, 2011. <http://alturl.com/oyksi>.

- [88] FJP Soler, C Le Couedic, F Breton. Optimum Proton Energy at Neutrino Factory using HARP Data. *Fifth Neutrino Factory International Design Study, Fermilab*, 9 April 2010, 2010.
- [89] D. Neuffer, M. Martini, G. Prior, C. Rogers, and C. Yoshikawa. Muon capture in the front end of the IDS neutrino factory. *WEPE068*, 2010.
- [90] Albright, Carl H. and others. The neutrino factory and beta beam experiments and development. *Neutrino Factory/Muon Collider Collaboration*, 2004. aka Feasibility Study II-a.
- [91] R Palmer. Optimization of Cooling vs Acceleration Acceptance. *International Scoping Study Meeting (KEK)*, 2006. <http://alturl.com/5e8hr>.
- [92] The ISS Accelerator Working Group. Accelerator Design Concept for Future Neutrino Facilities. *2009 JINST 4 P07001*, 2009.
- [93] R. Edgecock. Introduction to the non-scaling electron model FFAG EMMA. *Nucl.Phys.Proc.Suppl.*, 155:321–322, 2006.
- [94] Andrew Brian Laing. Optimisation of detectors for the golden channel at a neutrino factory, 2010. Ph.D. Thesis (Advisor: FJP Soler).
- [95] International Design Study for the Neutrino Factory (IDS-NF). <https://www.ids-nf.org/wiki/FrontPage>.
- [96] S Geer. Muon Colliders and Neutrino Factories. *Annual Review of Nuclear and Particle Physics*, 59:347–365, 2009.
- [97] The Muon Collider Collaboration. Mu+ Mu- Collider: A Feasibility Study. *Snowmass 96 Proceedings Fermilab-Conf-96/366*, 1996. <http://www.cap.bnl.gov/mumu/pubs/snowmass96.html>.
- [98] Helmut Wiedemann. *Particle Accelerator Physics*. Springer, third edition, 1993.

- [99] ED Courant, HS Snyder. Theory of the alternating gradient synchrotron. *Annals Phys.*, 3:1–48, 1958. Also reprinted in 'The Development of Coliders'. Edited by Claudio Pellegrini and Andrew M. Sessler. N.Y., AIP, 1995. pp. 23-70.
- [100] Elena Wildner. An Introduction to Accelerator Physics, NuFact08 Summer School. <http://benasque.org/2008nufact/cgibin/talks/allprint.pl>.
- [101] Mario Conte, William W MacKay. *An Introduction to the Physics of Particle Accelerators*. World Scientific, 2008.
- [102] S Van der Meer. Stochastic Damping of Betatron Oscillations in the ISR. *CERN Internal Report CERN/ISR-PO/72-31*, 1972.
- [103] H Poth. Electron Cooling: Theory, Experiment, Application. *Physics Reports*, 196(3-4):134–135, 1970. doi:10.1016/0370-1573(90)90040-9.
- [104] The MICE Collaboration. The Muon Ionisation Cooling Experiment Official Web-page. <http://mice.iit.edu>.
- [105] Steve Virostek. MICE Magnets. *Muon Accelerator Program Review - Fermilab*, 2010. http://mice.iit.edu/mico/meetings/mico120/mico120_virostek_solenoidandrfcc.pdf.
- [106] T Roberts. MICE Beamline Magnet Drawings. <http://hep04.phys.iit.edu/cooldemo/micenotes/public/pdf/MICE0065/MICE0065.pdf>.
- [107] A Dobbs, M Apollonio, K Long, J Pasternak, D Adams. The MICE Muon Beam: Status and Progress. MICE Note 300: <http://mice.iit.edu/micenotes/public/pdf/MICE0300/MICE0300.pdf>.
- [108] G H Eaton. Muon Facilities and Experiments; Now and in the New Millenium. *CCLRC Technical Report RAL-TR-1998-019*, 1998.
- [109] P Hodgson, C Booth, PJ Smith, J Tarrant. MICE Target Hardware. MICE Note 312: <http://hep04.phys.iit.edu/cooldemo/micenotes/public/pdf/MICE0312/MICE0312.pdf>.

- [110] Company website. E. I. du Pont de Nemours and Company. http://www2.dupont.com/DuPont_Home/en_US/index.html.
- [111] W Lau, SQ Yang, MA Green, S Ishimoto, J Swanson. The Development of a 6061-Aluminum Window for the MICE Liquid Absorber. *Internal MICE Note 130*, 2005. <http://alturl.com/djhb4>.
- [112] C Rogers, L Coney, P Snopok. Wedge Absorber Simulations for the Muon Ionisation Cooling Experiment. MICE Collaboration Meeting 26, <http://cm26.ucr.edu>.
- [113] S Ishimoto, S Suzuki, M Yoshida, MA Green, Y Kuno, W Lau. Liquid Hydrogen Absorber for MICE. *Internal MICE Note 301*, 2010. <http://alturl.com/v5pgv>.
- [114] D Li, MA Green, SP Virostek, MS Zisman, W Lau, AE White, SQ Yang. Progress on the RF Coupling Coil Module Design for the MICE Channel. *Proceedings of the Particle Accelerator Conference 2005*, doi: 10.1109/PAC.2005.1591297, pages 2869–2871, 2005.
- [115] Applied Fusion Ltd. Applied Fusion Ltd Web Page. <http://www.applied-fusion-ltd.co.uk/home.htm>.
- [116] A Bross. Fermilab Detector Update. *MICE Collaboration Meeting 23*, 2009. <http://alturl.com/5mjmk>.
- [117] A Blondel et al. MICE Status Report October 2009. *MICE Collaboration Internal Note 267*, 2009. <http://alturl.com/2cz65>.
- [118] R. Bertoni et al. The design and commissioning of the MICE upstream time-of-flight system. *Nucl.Instrum.Meth.*, A615:14–26, arXiv:hep-ex/1001.4426, 2010.
- [119] R Bertoni, M Bonesini, A de Bari, G Cecchet, Y Karadzhov, S Terzo. The Construction and Laboratory Tests for the MICE TOF0/1 Detectors. *MICE Collaboration Internal Note 241*,

- <http://hep04.phys.iit.edu/cooldemo/micenotes/public/pdf/MICE0241/MICE0241.pdf>, 2009.
- [120] R Bertoni, M Bonesini, A de Bari, G Cecchet, Y Karadzhov, R Mazza. The Construction of the MICE TOF2 Detector. *MICE Collaboration Internal Note 286*, <http://hep04.phys.iit.edu/cooldemo/micenotes/public/pdf/MICE0286/MICE0286.pdf>, 2010.
- [121] DA Sanders. MICE Particle Identification Systems. *Proceedings of PAC09*, 2009.
- [122] Mark Rayner. First Results of MICE. *MICE Funding Agency Committee*, 2010. <http://alturl.com/9i7cv>.
- [123] The MICE Collaboration. MICE Status Report December 2008. *MICE Collaboration Internal Note 230*, 2008. <http://alturl.com/m22h2>.
- [124] The MICE Collaboration. MICE Technical Reference Document, 2005. http://mice.iit.edu/trd/MICE_Tech_ref.html.
- [125] Alessandro Manfredini et al. KL Manual. <http://mice.iit.edu/mico/manuals/KL-manual-15Dec09.pdf>.
- [126] The KLOE Collaboration. The KLOE detector : technical proposal. *LNF-93-002-IR*, 1993.
- [127] M. Ellis, P.R. Hobson, P. Kyberd, J.J. Nebrensky, A. Bross, et al. The Design, Construction and Performance of the MICE Scintillating Fibre Trackers. *physics.ins-det*, 2010. arXiv:1005.3491.
- [128] D0 experiment. <http://www-d0.fnal.gov/>.
- [129] H Sakamoto. Status of the MICE Tracker System. *Proceedings of IPAC10*, 2010. <http://alturl.com/zmiwb>.

- [130] S.P Virostek M.A Green F Trillaud M.S Zisman. Fabrication, testing and modeling of the mice superconducting spectrometer solenoids. *Proceedings of IPAC10*, 2010. <http://alturl.com/694wz>.
- [131] CT Rogers, D Adey, M Ellis, R Gamet, M Littlefield, H Nebrensky. Study of the Proton Absorber in the MICE Beamline. MICE Internal Note 294: <http://hep04.phys.iit.edu/cooldemo/micenotes/public/pdf/MICE0294/MICE0294.pdf>.
- [132] Malcolm Ellis. MICE Software Design and Physics Performance. *Astroparticle, Particle, and Space Physics, Detectors and Medical Physics Applications - Proceedings of the 10th Conference*, World Scientific Publishing., pages 763–770, 2008.
- [133] Muons Inc. G4Beamline Simulation Software. <http://g4beamline.muonsinc.com>.
- [134] Muons Inc. Muons Inc. Home Page. <http://www.muonsinc.com/>.
- [135] K. Yonehara et al. MANX, a 6-D Muon Cooling Experiment for RAL. *Proceedings of EPAC 2009*, 2009. <http://alturl.com/wodv9>.
- [136] Robert Abrams et al. A Proposal to Fermilab for an Accelerator Experiment to Study Muon Beam Cooling Techniques for Muon Colliders, Neutrino Factories and Stopping Muon Beams. <http://tinyurl.com/37rsaep>.
- [137] EPICS Collaboration. Experimental Physics and Industrial Control System. <http://www.aps.anl.gov/epics/index.php>.
- [138] W3C. SOAP Specification. <http://www.w3.org/TR/soap12-part1/>, 2007.
- [139] JJ Nebrensky. Notes from Data Flow Workshop. *MICE Note 255*, 2009. <http://alturl.com/ebooc>.
- [140] F Rolland. *The Essence of Databases*. Prentice Hall, 1997. Chapter 3.
- [141] F Rolland. *The Essence of Databases*. Prentice Hall, 1997. Chapter 4.
- [142] Open Source PostgreSQL Project. Postgresql. <http://www.postgresql.org/>.

- [143] The MySQL Open Source Project. MySQL Home page. <http://www.mysql.com/>.
- [144] Oracle Corporation. Oracle Home Page. <http://www.oracle.com/us/products/database/index.html>.
- [145] Richard T. Snodgrass. *Developing Time-Oriented Database Applications in SQL*. Morgan Kaufmann Publishers, Inc., 1999.
- [146] INCITS Technical Committee H2. The SQL 92 Standard. <http://www.contrib.andrew.cmu.edu/shadow/sql/sql1992.txt>, 1992.
- [147] Christian S. Jensen et al. Dengfeng Gao. Join operations in temporal databases. *International Journal on Very Large Databases*, 2005.
- [148] W3C. Extensible Markup Language (XML) 1.0 (Fifth Edition). <http://www.w3.org/TR/xml/>, 2008.
- [149] The Apache Software Foundation. Apache Tomcat. <http://tomcat.apache.org/>.
- [150] W3C. Http 1.1 specification. <http://www.w3.org/Protocols/>, 2010.
- [151] J Gostling, B Joy, G Steele, G Bracha. *The Java Language Specification, Third Edition*. Addison Wesley, 2005.
- [152] B Stroustrup. *The C++ Programming Language: 3rd Edition*. Addison Wesley, 1997.
- [153] David Forrest. MICE Configuration Database Front-End. <http://micewww.pp.rl.ac.uk:8080/configdb/generic/>.
- [154] Guido van Rossum and the Python Software Foundation. Python. <http://www.python.org/psf/>.
- [155] I D Reid. The G4MICE Geometry Interface to the Configuration Database. Presentation given at MICE Collaboration Meeting 29, <http://alturl.com/4gtsn>.
- [156] David Forrest. The Configuration Database System, 2010. MICE Internal Note 327, <http://alturl.com/hq3ck>.

- [157] David Forrest. Database Workshop 2010. <http://alturl.com/gciko>.
- [158] FJP Soler David Forrest. A new application of the grid: Muon ionization cooling for a neutrino factory. *Philisophical Transactions of the Royal Society A*, 14:4103–4113, 2010.
- [159] David Forrest and FJP Soler. A Study of Statistical Errors in MICE. *AIP Conference Proceedings*, 1222:490–493, 2010.
- [160] J Cobb. Statistical Errors on Emittance Measurements. Internal MICE Note 268 <http://hep04.phys.iit.edu/cooldemo/micenotes/public/pdf/MICE0268/MICE0268.pdf>.
- [161] David Forrest. The Grid & MICE. *MICE Internal Note 246*, 2009. <http://alturl.com/88xqt>.
- [162] Ulrik Egede, et al. Ganga User Interface for Job Definition and Management. *Proceedings 4th International Workshop on Frontier Science: New Frontiers in Subnuclear Physics, Milan, Biococca, Italy, 12-17 Sep 2005*, 369, 2005. <http://www.lnf.infn.it/sis/frascatiseries/Volume40/volume40.pdf>.
- [163] gLite Collaboration. gLite Workload Management System. <http://web.infn.it/gLiteWMS/>.
- [164] Job Description Language. <http://www.dcc.fc.up.pt/ines/aulas/0910/CG/jdl.pdf>.
- [165] Ganga Project Home Page. <http://ganga.web.cern.ch/>.
- [166] David Forrest and FJP Soler. Alignment Errors on Emittance Measurements for MICE. *Proceedings of Science*, NUFACT08:122, 2009. <http://alturl.com/6u4gx>.
- [167] Hamamatsu Photonics. Hamamatsu Photonics Official Site. <http://www.hamamatsu.com/>.
- [168] MA Clarke-Gayther, A Borden, GM Allen. Global Beam Loss Monitoring Using Long Ionisation Chambers at ISIS. *Proc. 4th European Particle Accelerator Conference (EPAC 94)*, 1994.

- [169] K Walaron. Neutrino Factory Targets and the MICE Beam. *PhD Thesis (University of Glasgow)*, 2007.
- [170] D Adams. Private Communication. ISIS.
- [171] FJP Soler. Status of MICE, NuFact09 Conference Proceedings. *AIP Conf.Proc.*, 1222:288–292, 2010.
- [172] CERN. Geant4. <http://geant4.web.cern.ch/geant4/>.
- [173] M Drummond. The MICE Luminosity Monitor. *BSc Project, University of Glasgow*, 2011.
- [174] The MICE Collaboration. Characterization of the Muon Beam Line for the Muon Ionisation Cooling Experiment. *In Preparation*, 2011.
- [175] C Rogers, D Adey, M Ellis, R Gamet, M Littlefield, H Nebrensky. Study of the Proton Absorber in the MICE Beamline. <http://alturl.com/87mgy>.
- [176] Marco Appolonio. The (Emittance,Momentum) Matrix for the MICE Beamline. http://mice.iit.edu/bl/MATRIX/index_mat.html.
- [177] M Rayner. Geometry Requirements for Analysis. *MICE Collaboration Meeting 28 (Sofia)*, 2010. <http://alturl.com/zyhaj>.