THE UNIVERSITY OF
WARWICK

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Please complete the following:

Full name: …………………………………………………………………………………………….

University ID number: ………………………………………………………………………………

2. THESIS DEPOSIT

2.1 I understand that under my registration at the University, I am required to deposit my thesis with the University in BOTH hard copy and in digital format. The digital version should normally be saved as a single pdf file.

2.2 The hard copy will be housed in the University Library. The digital version will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.3 below) this will be made openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service.
[At present, theses submitted for a Master's degree by Research (MA, MSc, LLM, MS or MMedSci) are not being deposited in WRAP and not being made available via EthOS. This may change in future.]

2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the hard copy thesis for a limited period. It is also possible to apply separately for an embargo on the digital version. (Further information is available in the *Guide to Examinations for Higher Degrees by Research*.)

2.4 *If you are depositing a thesis for a Master's degree by Research, please complete section (a) below. For all other research degrees, please complete both sections (a) and (b) below:*

(a) Hard Copy

I hereby deposit a hard copy of my thesis in the University Library to be made publicly available to readers (please delete as appropriate) EITHER immediately OR after an embargo period of ……….................... months/years as agreed by the Chair of the Board of Graduate Studies.

I agree that my thesis may be photocopied. YES / NO (*Please delete as appropriate*)

(b) Digital Copy

I hereby deposit a digital copy of my thesis to be held in WRAP and made available via EThOS.

Please choose one of the following options:

EITHER My thesis can be made publicly available online. YES / NO (*Please delete as appropriate*)

OR My thesis can be made publicly available only after…..[date] (Please give date)
YES / NO (*Please delete as appropriate*)

OR My full thesis cannot be made publicly available online but I am submitting a separately identified additional, abridged version that can be made available online.
YES / NO (*Please delete as appropriate*)

OR My thesis cannot be made publicly available online. YES / NO (*Please delete as appropriate*)

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following:

Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. DECLARATIONS

(a) I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.

- The digital version of the thesis I am supplying is the same version as the final, hard-bound copy submitted in completion of my degree, once any minor corrections have been completed.

- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.

- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

(b) IF I HAVE AGREED (in Section 2 above) TO MAKE MY THESIS PUBLICLY AVAILABLE DIGITALLY, I ALSO DECLARE THAT:

- I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.

- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.

5. LEGAL INFRINGEMENTS

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and return it to the Graduate School Office when you submit your thesis.

Student's signature: ....…… Date: ....

Beam instrumentation and investigations into muon cooling at MICE

by

David Adey

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

February 2012

THE UNIVERSITY OF
WARWICK

Contents

List of Tables

List of Figures

Acknowledgments

I would like to thank my supervisors at Warwick, Steve Boyd and John Back, for their comments and advice whilst completing this thesis. I am indebted also to my colleagues from the MICE collaboration, especially Alan Bross, Malcolm Ellis, Ken Long, Chris Rogers and Paul Rubinov, for their help in conducting this research. Finally, I am grateful, as ever, to my family for their support.

Declarations

The material from Chapter One is taken entirely from available publications and other sources for the relevant topics. Chapter Two consists of material taken from publications and internal documents relating to MICE, and from discussions within the collaboration. Chapter Three consists of material from internal documents and discussions and development of my own algorithms, software and analysis. Specifically, Sections 3.1 to 3.4 are taken from sources, and Sections 3.5 to 3.10 is my own work. The premise behind Chapter Four was inspired by discussions with MICE colleagues, but all subsequent work is my own. Chapter Five is entirely my own work.

Abstract

The Muon Ionisation Cooling Experiment (MICE) aims to make a proof of principle measurement of ionisation cooling of muon beams for use in a future neutrino factory or muon collider. To complete this measurement, two precision scintillating fibre spectrometers have been constructed, for which data acquisition systems and calibration procedures have been developed. In addition, reconstruction algorithms have been written and the performance of the detectors tested with cosmic rays.

To explore the behaviour of the MICE cooling channel under various conditions associated with the beam emittance and RF phase, a method of particle weighting using approximate Voronoi diagrams was used and shown to be a viable method of statistical weighting. Investigations have also been made into approximating the MICE cooling channel with transfer matrices, using a neural network approach to fit the data. This showed that whilst the effect of the magnetic field on the particle trajectory can be accounted for in a simple transfer matrix, the inclusion of stochastic effects is essential to an accurate approximation of emittance transferral through the channel.

Chapter 1

Introduction

1.1 The Status of Neutrino Physics

The discovery of the continuous energy distribution of the electron in β decays led to the postulation of the neutrino by Wolfgang Pauli in 1930, a suggestion which, given the near impossibility of verification perceived at the time, Pauli referred to as "*a very terrible thing*". Over the following years, the theory of weak interactions expanded to account for the conservation of spin and parity, and with the discovery of the muon in 1937 and the observation of its decay, Pontecorvo proposed the universality of electron and muon behaviour, first articulating the concept of the particle family. With the postulation and confirmation of laws of symmetry came their exceptions, in particular that of parity violation in β decay as found by C. S. Wu [1]. With the arrival of nuclear power there was now a high flux source of neutrinos which could be exploited in an inverse β decay process to make a measurement of the neutrino's existence, and in 1956 the neutrino was finally observed directly by Reines and Cowan [2].

Observations of parity violation in the lepton sector led to the proposition of the left-handedness of neutrinos, confirmed by electron capture experiments in 1958 [3]. Lepton number conservation introduced by Konopinski and Mahmoud [4] in 1953 was well supported by the initial reactor experiments and attempts to measure $\bar{\nu}_e + {}^{37}Cl \rightarrow {}^{37}Ar + e^-$, which failed. Additionally, the failure to measure muon decay via an electron and photon implied the conservation not only of lepton number, but of the individual lepton generation number.

The development of the standard model by Glashow [5], Weinberg [6] and Salam [7] and its confirmation by experiments at CERN and Fermilab, including the discovery of the $W^{\pm}$ and Z particles [8, 9], began a rapid expansion of the list of fundamental particles, finally including the τ lepton and the top and bottom quarks. Whilst confirmed in its predictions in many cases, the standard model requires that neutrinos be massless. The number of interacting neutrino generations is confirmed at three from completed experiments [10], although this accounts only for interacting neutrinos. The possibility of the existence of right-handed, sterile neutrinos, which would interact only by gravity, is still an issue.

Whilst oscillations had been observed in strange quark interactions, a comprehensive model of neutrino mixing was not proposed until 1967 with concurrent models of two flavour mixing and preliminary suggestions of oscillation by Maki, Nakagawa, Sakata [11] and Pontecorvo. In the former's model, the flavour eigenstates of the ν_e and ν_μ are mixings of mass eigenstates. The first hints of oscillations were observed as the solar neutrino anomaly by the Homestake experiment, where the measured flux of electron neutrinos from nuclear interactions in the Sun was too low. Confirmed most recently by the SNO experiment, the deficiency is attributed to the ν_e oscillating into ν_μ and ν_τ whilst still in the Sun.

With the model of neutrino oscillations dependent on the massive eigenstates and the angles of their mixing, progress has been made into the measurement of these properties. At the time of writing, two of the three mixing angles are well defined, as are the squared differences between the mass states. Limits exist on θ_{13}, although the orientation of the mass hierarchy, as well as the value of the CP phase δ are unknown. Furthermore, the nature of neutrinos has yet to be determined in whether they are Dirac or Majorana particles. Experiments looking for neutrino-less double β decay hope to answer this question.

1.1.1 Flavour oscillation

Neutrinos are described as a mixing between mass and weak eigenstates, where the weak state corresponds to the lepton flavours e, μ and τ. This gives rise to parameters which describe how the states mix, and to the prediction of the states oscillating between one another.

Defining the weak, or flavour, eigenstate ν_α of the neutrino in terms of a linear composition of mass eigenstates ν_k, where these mass eigenstates are the eigenstates of the Hamiltonian, the neutrino state may be written as

$$|\nu_\alpha\rangle = \sum_k U^*_{\alpha k}|\nu_k\rangle \tag{1.1}$$

where U, the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, is defined in Equation 1.2 [12], in which $\theta_{\alpha\beta}$ is the mixing angle between those eigenstates, and s or

c are the sine or cosine of that angle respectively. Possible CP violating effects can occur through the inclusion of $e^{i\delta}$.

$$U = \begin{pmatrix} c\theta_{12}s\theta_{13} & s\theta_{12}c\theta_{13} & s\theta_{13}e^{-i\delta} \\ -s\theta_{12}c\theta_{13} - c\theta_{12}s\theta_{23}s\theta_{23}e^{i\delta} & c\theta_{12}c\theta_{13} - s\theta_{12}s\theta_{23}s\theta_{23}e^{i\delta} & s\theta_{23}c\theta_{13} \\ s\theta_{12}c\theta_{13} - c\theta_{12}s\theta_{23}s\theta_{23}e^{i\delta} & -c\theta_{12}c\theta_{13} - s\theta_{12}s\theta_{23}s\theta_{23}e^{i\delta} & c\theta_{23}c\theta_{13} \end{pmatrix} \tag{1.2}$$

As eigenstates of the Hamiltonian, the mass states should evolve over a time t as

$$|\nu_\alpha(t)\rangle = e^{-iE_k t}|\nu_k\rangle \tag{1.3}$$

where E_k is the energy of the mass eigenstate. With the mass states themselves composed of flavour states, the evolution of a single flavour state over time can be defined as

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau}\left(\sum_k U^*_{\alpha k}e^{-iE_k t}U_{\beta k}\right)|\nu_\beta\rangle \tag{1.4}$$

where the state becomes a superposition of other flavour states. This ultimately leads to the probability of one flavour state oscillating into another, which is given in the general case in Equation 1.5,

$$P_{\nu_\alpha\to\nu_\beta} = \sum_{k,j} U^*_{\alpha k}U_{\beta k}U_{\alpha j}U^*_{\beta j}exp\left(-i\frac{\Delta m^2_{kj}L}{2E}\right) \tag{1.5}$$

where L is the distance travelled by the neutrino, and is more commonly used in the two-flavour case given by:

$$P_{\nu_\alpha\to\nu_\beta} = sin^2 2\theta sin^2\left(\frac{\Delta m^2 L}{4E}\right) \tag{1.6}$$

A measurement of oscillation and the mixing parameters can then be made by measuring the flux of a neutrino beams at two different values of L, as is the case with accelerator and reactor experiments, or by having a well defined model of the source neutrino flux, as in a solar neutrino experiment. The value of L/E is key in designing an experiment to measure the maximum oscillation between states. The energy and baseline must be selected to correspond to this maximum, which creates issues for accelerator design and selecting detector locations.

1.2 Current Experiments

With the observation of neutrino oscillations a wide range of experiments have been undertaken to investigate the phenomena and make measurements of the mixing parameters. The experiments rely on neutrinos from accelerators, nuclear reactors and natural sources, utilising various detector technologies. The experiments described here are long baseline experiments, where the baseline L in the L/E component of the oscillation probability has been optimised, often at distances of several hundred kilometres.

MINOS

The Main Injector Neutrino Oscillation Search (MINOS) [13] is an accelerator based long baseline experiment from Fermilab to Soudan. Using the main injector ring of the Tevatron, 120 GeV protons produce a neutrino beam with an energy range of 1 to 30 GeV. The near and far detector setup creates baselines of L = 1 km and 735 km, with both being similar steel-scintillator calorimeters. This has led to a ν_e appearance measurement from the ν_μ beam, the results of which are shown in Figure 1.1 [14].

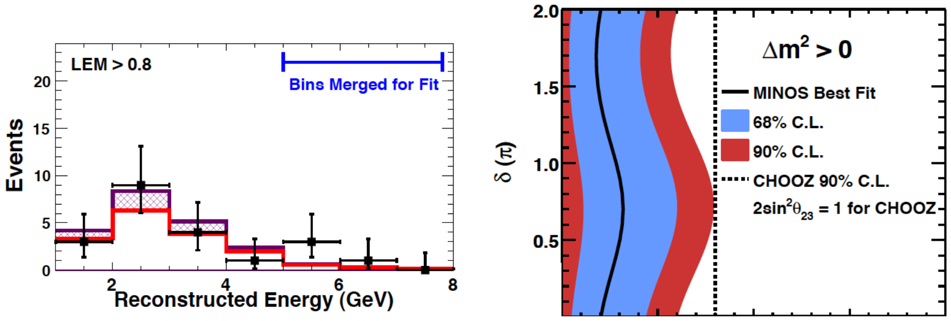

(a) Signal (b) θ_{13} Best fit

Figure 1.1: Oscillation results from MINOS. The excess events as a function of energy in Figure (a) show the oscillation of neutrinos, which when fitted to find the mixing angles generate the range of possible values defined in Figure (b).

T2K

The Tokai to Kamiokande (T2K) experiment [15] is a long baseline neutrino experiment with a goal of measuring, or setting a limit on, θ_{13} and improving the accuracy of the measurements of Δm^2_{23} and θ_{23}. The θ_{13} measurement is achieved with a ν_e

appearance in a ν_μ beam, whilst the other measurements are made with a ν_μ disappearance. The neutrino beam is produced at the J-PARC accelerator by colliding high energy protons into a solid target, producing pions, and then focusing the pions and allowing them to decay weakly into a muon and a muon neutrino. T2K uses an off-axis beam to improve the energy spread of the neutrinos at Super-Kamiokande.

T2K consists of two detectors in a near-far arrangement. The far detector is the Super-Kamiokande water Cerenkov detector, offset from the beam axis by 2.5° for a better beam energy distribution. The detector contains a volume of 50000 tons of water instrumented with 13000 photomultiplier tubes. The near detector, ND280, contains both on- and off-axis components, and is tasked with measuring the flux, content and interaction rates of the un-oscillated beam.

At the time of writing, preliminary results from T2K gave a non-zero measurement of θ_{13} at a confidence level of 90%, with a mean value of $\sin^2 2\theta_{13}$ of 0.11. This is shown in Figure 1.2 [16].

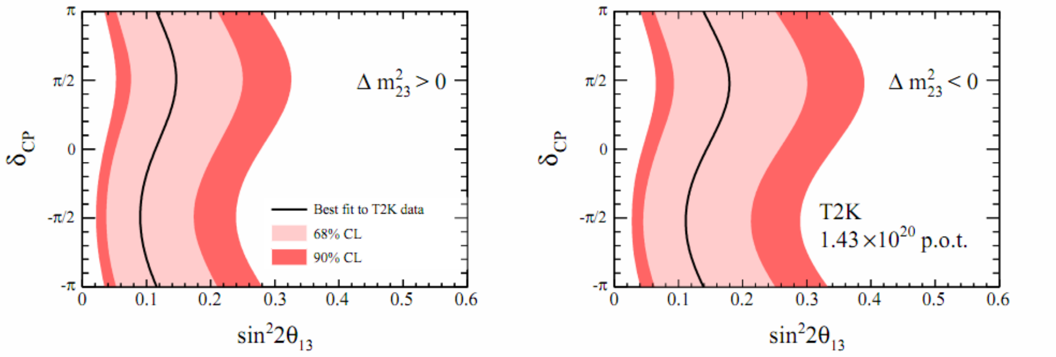

Figure 1.2: T2K preliminary results. By fitting the excess events in the ν_e appearance channel to an oscillation probability, such as that given in Equation 1.6, the values of the mixing angle θ_{13} and the phase δ can be found for standard and inverted mass hierarchies.

If θ_{13} is large, it is conceivable the parameters of the PMNS matrix could be measured to the desired accuracy with current, or similar, experiments. The value of $e^{i\delta}$ may still remain beyond their reach, requiring a significantly more sensitive experiment.

OPERA

The Oscillation Project with emulsion tracking apparatus (OPERA) [17] uses a neutrino beam produced at CERN and directed towards the Gran Sasso laboratory.

OPERA is a detector which consists of thin layers of emulsion interspaced with scintillator and lead; the fine granularity along the z axis enables the reconstruction of the τ decay vertex, resulting in the first direct observation of the ν_τ from oscillation [18].

Reactor Experiments

Reactor based neutrino experiments use the high flux from the nuclear interactions inside the reactor core. Anti-electron neutrinos from beta decay in the reactor are detected in the reciprocal process in detectors at varying baselines. There are currently three reactor based experiments preparing to start experimental running: Double Chooz [19], Daya bay [21] and Reno [22]. Reactors have the benefit of making an unambiguous measurement of θ_{13} without the interference of matter effects or CP violation.

Double Chooz has produced preliminary results supporting a non-zero measurement of θ_{13}, albeit with a lower best-fit value of $\sin^2 2\theta_{13} = 0.086$ [20].

1.2.1 Global best fits

Combining the data from completed experiments provides a best fit to the oscillation parameters from global data. This consists of measurements for most parameters, given in Table 1.1, and a limit on θ_{13}, which is shown in Figure 1.3.

Parameter	Value
$\sin^2\theta_{12}$	$0.306^{+0.018}_{-0.015}$
$\sin^2\theta_{23}$	$0.420^{+0.08}_{-0.03}$
$\sin^2\theta_{13}$	$0.021^{+0.007}_{-0.008}$
Δm^2_{12}	$7.58^{+0.22}_{-0.26} \times 10^{-5}$ eV
Δm^2_{23}	$2.35^{+0.12}_{-0.09} \times 10^{-3}$ eV

Table 1.1: Global best fit for neutrino mixing parameters [23].

The suggestion of a non-zero θ_{13} by the T2K experiment is a recent measurement, prior to which it was felt that existing experiments could be incapable of making a measurement of θ_{13} to any precision, requiring the construction of a far more sensitive setup. Additionally, the value of δ is strongly linked to θ_{13} and will almost certainly require a more sensitive experiment to measure. The key components of any next generation experiment would be a high flux, improved flavour selection and sensitive control of the L/E oscillation parameter.

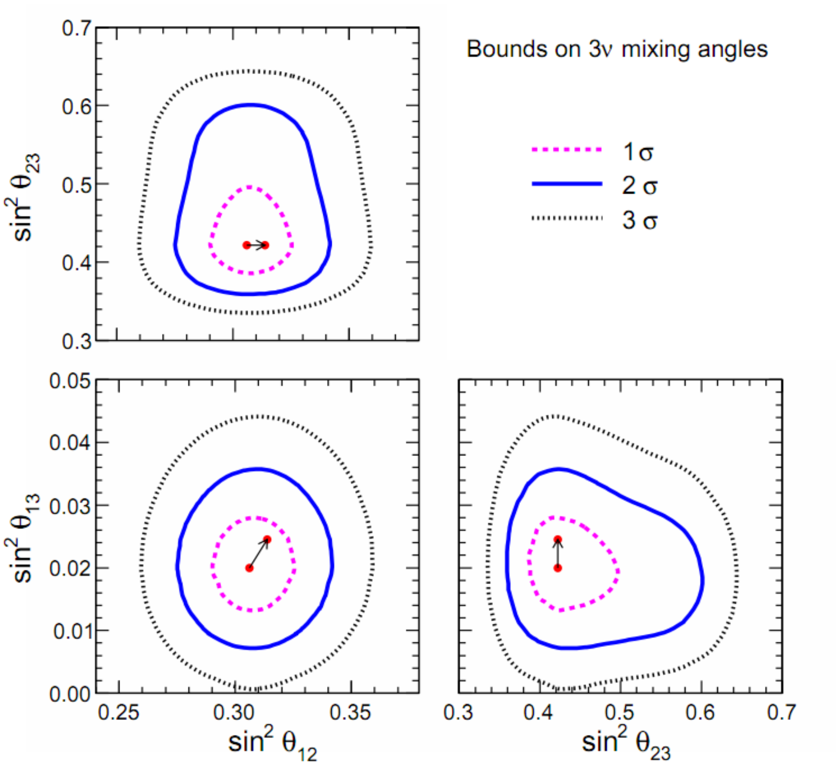

Figure 1.3: Global best fit mixing angles with confidence levels. The shifted central value corresponds to differing reactor flux models.

1.3 A Neutrino Factory

It has already been noted that limitations arise in current oscillation experiments due to their sensitivities to small parameters. If θ_{13} remains very small, the excess events required to make a measurement of it could be beyond the statistical limit of the current generation of experiments. A higher flux beam which could greatly reduce the statistical limit would therefore be of use in this region.

The three flavour oscillation scenario provides a number of possible channels in which to observe the neutrinos, although this is then limited by the ability of the accelerator to select the charge of the muon, and the detectors in their ability to identify particle types. A multi-faceted beamline could provide separate neutrino beams to different detector types, allowing for the observation of some or all of the interaction channels.

Given the limitations of current experiments to measurements of the mass hierarchy and mixing angles, and in particular to the CP phases, it becomes desirable to construct an experiment which can probe these values to ever greater accuracy. Candidate experiments consist of superbeams, β beams and a neutrino factory. A superbeam would comprise a power upgrade to an existing facility such as J-PARC, and an expansion of the Kamiokande detector into a 1 Mton water Cerenkov detector. A β beam experiment involves the acceleration of isotopes which produce a β decay, generating a boosted $\bar{\nu}_e$ beam.

The sensitivities of superbeams, β beams and a neutrino factory to measurements of θ_{13} and δ are given in Figures 1.4 and 1.5, where a neutrino factory is clearly more sensitive in every respect to the values of interest.

Not only would the high flux achievable by a neutrino factory, simulated at 10^{21} muon decays per year, provide statistics capable of measuring a limit on θ_{13} to $\sin^2 2\theta_{13} = 10^{-5}$, but the long baselines required given the high muon energy would force the CP asymmetries to be dominated by matter effects, allowing for an absolute measurement of the mass hierarchy. The sensitivity to the phase δ would also cover the viable regions at a neutrino factory, although with a lower than expected performance due to degeneracies. The sensitivity of a neutrino factory to θ_{13} and the mass hierarchy are given in Figures 1.6 and 1.7, respectively.

1.3.1 Physics concept

The key component in such an experiment is the production of the neutrinos themselves. Existing high flux neutrino experiments such as T2K have used a super-beam operating mode, which involves a high power proton beam hitting a high Z target,

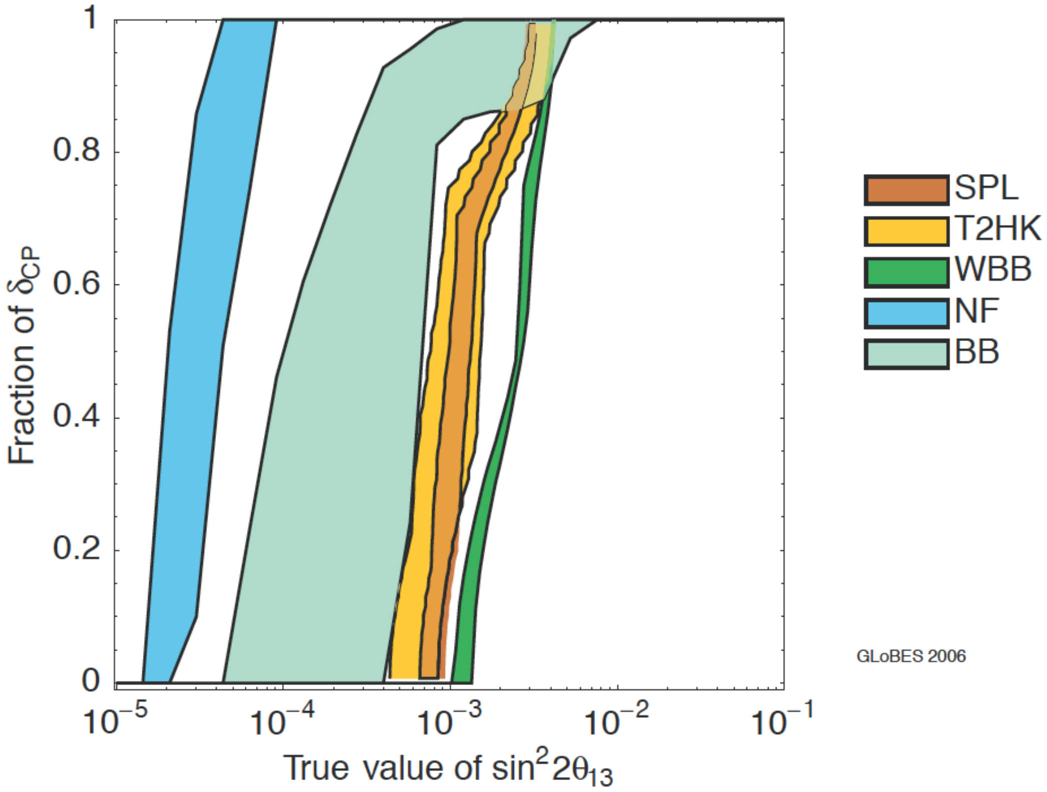

Figure 1.4: Comparison of the discovery potential of θ_{13} of future neutrino oscillation experiments for a neutrino factory (NF), β beam (BB) and superbeam setups. The right edge of each coloured band represents a conservative setup (and therefore minimum sensitivity) and the left edge represents an optimised setup (maximum sensitivity) [24].

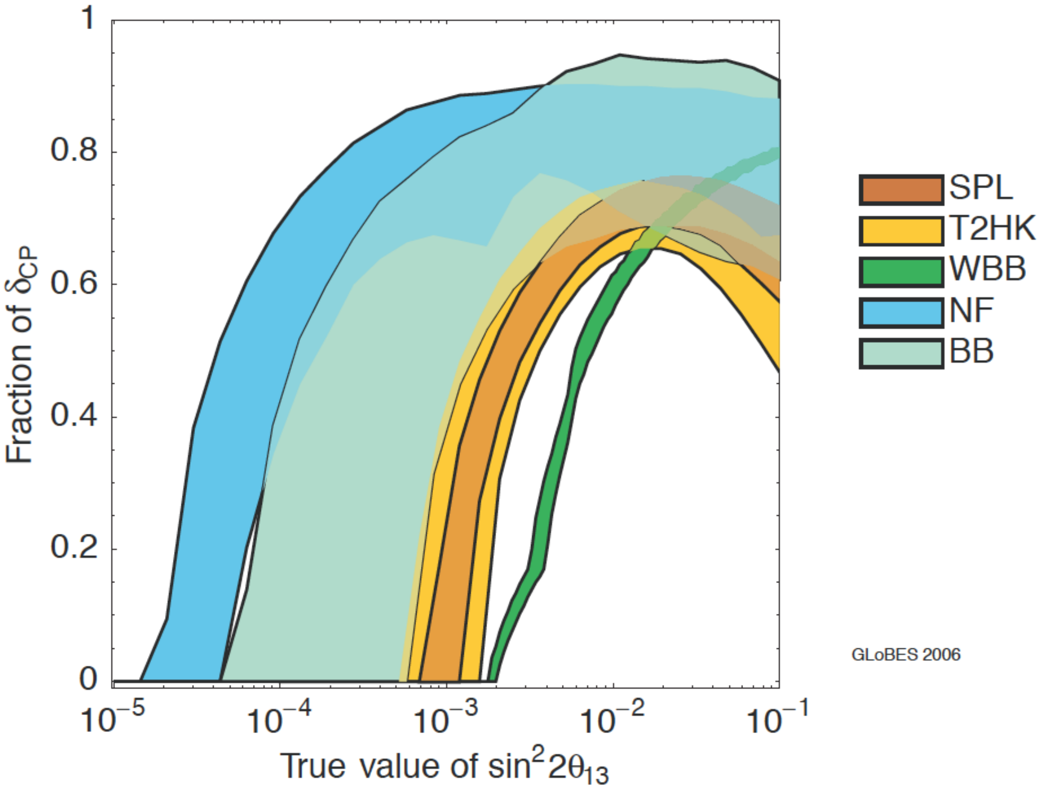

Figure 1.5: Comparison of the discovery potential of δ of future neutrino oscillation experiments for a neutrino factory (NF), β beam (BB) and superbeam setups. The right edge of each coloured band represents a conservative setup (and therefore minimum sensitivity) and the left edge represents an optimised setup (maximum sensitivity) [24].

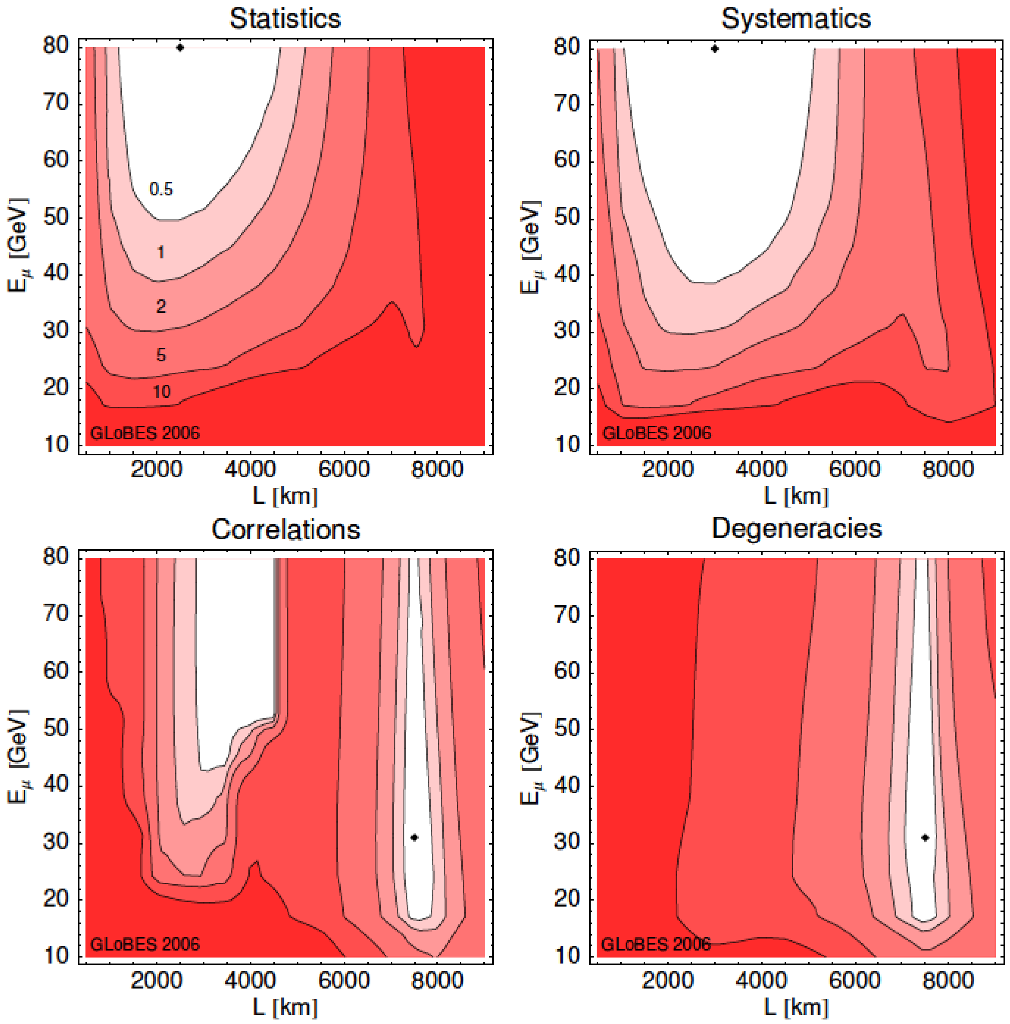

Figure 1.6: Neutrino factory 5σ sensitivity to θ_{13} due to varied effects, where the white region is the maximum sensitivity and the contours show factor increases given on the figures [24].

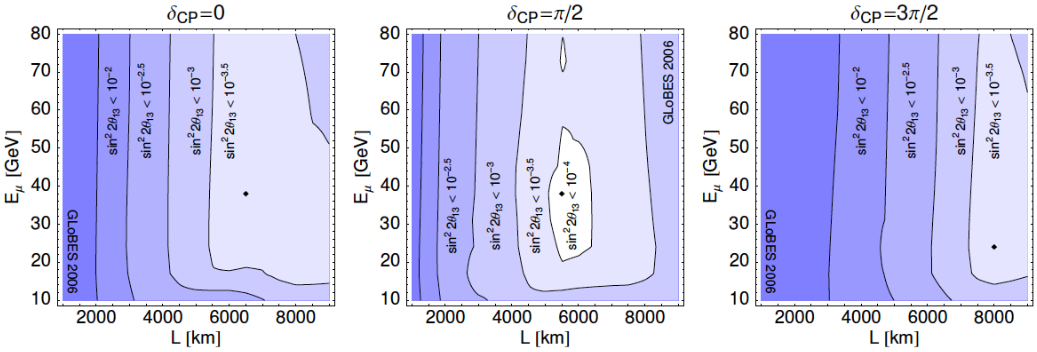

Figure 1.7: Neutrino factory 3σ sensitivity to the mass hierarchy, given for three values of δ and plotted against the energy and length of the baseline. The contours correspond to the minimum value of θ_{13} at which the hierarchy could be ascertained [24].

with the resulting pion beam focussed by a horn; the pions are then allowed to decay into muons and muon neutrinos, and the muons decay into electrons, electron neutrinos and muon anti-neutrinos in a decay pipe oriented towards the detectors.

A neutrino factory would begin with this concept, but instead of allowing the muons to decay they would be re-accelerated to a controlled energy and allowed to decay in re-circulating decay rings. This will produce a high flux mono-energetic neutrino beam with a selection on either neutrino or anti-neutrino, provided by a beamline which can select the charge of the muons. The particle production chain would then be:

$$\begin{aligned} p + p &\rightarrow \pi^{+}\pi^{-}\pi^{0} \\ \pi^{+} &\rightarrow \mu^{+}\nu_{\mu} \\ \mu^{+} &\rightarrow e^{+}\nu_{e}\bar{\nu}_{\mu} \end{aligned}$$

This means the neutrino beam produced by a neutrino factory will be half ν_e and half ν_μ if the π^+ is selected, and the reverse if the π^- is selected. The high flux beam could then be directed towards near and far detectors; the change in baseline allowing for measurement of the L/E dependance in neutrino oscillation probabilities. Currently, two-detector experiments such as T2K employ a near detector of the order of hundreds of metres from the decay point, and a far detector of the order of hundreds of kilometres from the decay point.

Construction of a neutrino factory places demands on the performance of accelerator components which are not currently met: high power proton sources, targets capable of withstanding multi-MW beams, large acceptance accelerating devices and cooling which can operate within the lifetime of a muon are all prerequisites of building such an accelerator complex, such as that shown in Figure 1.8.

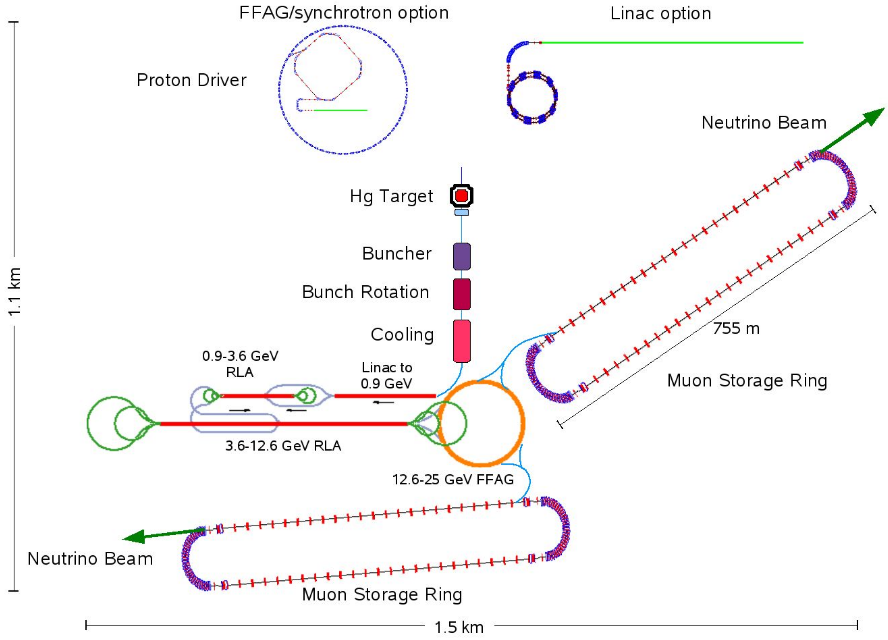

Figure 1.8: Layout of a neutrino factory with either a linear accelerator or FFAG proton driver. This figure shows the racetrack option for the muon decay rings [26].

These components comprise the front end of a neutrino factory, which is almost identical to that of a muon collider, where the muons would enter recirculating collider rings rather than decay rings. Muon colliders have the advantage over linear accelerators of multiple acceleration turns without the losses due to synchrotron radiation, which has limited the energy of lepton rings such as LEP. Requirements on the transverse and longitudinal parameters of the muon beam are stricter for a muon collider, but the required technological advancements remain the same. In fact, construction of a muon collider would produce a neutrino factory almost as a by-product.

1.3.2 Accelerator

Proton Driver

The initial acceleration of the proton beam prior to its interaction with a target will take the form of a proton driver, which could be comprised of a number of different components, depending on the final design. The requirements for a proton driver are given in Table 1.2. The options can be summarised as:

- H^- ion source with a linear accelerator (linac) and non-linear fixed field alternating gradient (FFAG) driver rings.
- H^- ion source and linac, with a pair of 50 Hz booster and 25 Hz driver recirculating synchrotrons (RCS).
- Full energy H^- linac with accumulator and bunch compression rings.
- A combination of linac, synchrotron and FFAG options.

Parameter	Value
Average Beam Power	4 MW
Reptition Rate	50 Hz
Proton energy	10 ± 5 GeV
Proton bunch length	2 ± 1 ns

Table 1.2: Requirements for a neutrino factory proton driver

The non-full power linear accelerator option is based on a proposed upgrade to the ISIS facility at RAL. This would involve a 324 MHz, 200 MeV linac with a 50 Hz rate producing a 30 mA beam current. The booster synchrotron would then have four straight sections each with seven quadrupole triplets followed by 90° arcs. Depending on the bunch formation the synchrotron will use 2.1 - 3.6 MHz or 1.3 - 2.2 MHz RF. The full energy linac option is based on the Superconducting Proton Linac (SPL) providing a 5 GeV acceleration over approximately 530 m.

Target

The high power proton beam necessary to produce the required flux for a neutrino factory poses challenges for the construction of the target. Three options are under consideration for a neutrino factory target: a mercury jet (baseline), a solid target and a powdered jet. The choice of material is especially of concern for the solid target, both in terms of the produced pion yield and the target's ability to withstand the high thermal and kinetic shock of the proton beam.

In order to minimise the integrated energy deposition in a solid target, regimes involving a rotating target wheel, in which a new section of the target is presented to the beam at each pulse, allowing other sections to cool, are under investigation. The stress placed on the target when subjected to a 4 MW beam varies with the structure of the beam, including the micro-bunch structure and physical size.

To minimise the accumulated shock, liquid targets could be used with a high Z material such as mercury. The instantaneous energy deposition would then be distributed into a pool of mercury before it is re-circulated into the jet. The use of mercury generates its own safety concerns especially given the high pressure and heat involved, as well as the accumulated radiation. The MERIT experiment [27] looked into the use of mercury jets.

To avoid the safety and technical issues involved with high-pressure mercury, a powdered jet could be used in its place, consisting of a similarly high Z material. Achieving consistent flows with a powder of a dense material poses a technical challenge and is being investigated.

Front End

After production of pions at the target they must be collected and allowed to decay into muons. The muon beam will then have to undergo manipulation before re-acceleration and storage in the decay rings. All of these functions are grouped under the title of the neutrino factory front end, an example of which is shown in Figure 1.9.

The pions produced in the target will have a large transverse momentum spread produced by the high Z target. In order to maximise the capture rates, the particles must be produced in a high solenoidal field. The field strength at the interaction point should be approximately 20 T with a bore radius of 7.5 cm, tapering to 1.75 T. This would be followed by a solenoidal decay channel or drift space.

The muon beam should then undergo bunching. This can be achieved either by allowing the beam to drift, generating a momentum-time correlation which can be exaggerated into bunches by an RF field (decelerating the early, high momentum particles and accelerating the late, low momentum particles) or by bunching the beam in the drift space by varying the RF frequency with distance.

Once the phase rotation and bunching has been completed, the beam must match the downstream acceptance of the muon accelerator. The high transverse phase space of the beam means very few particles will transmit through the accel-

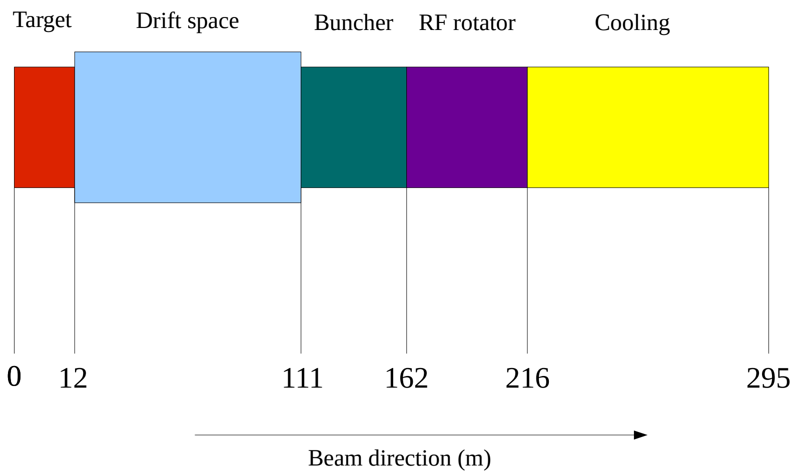

Figure 1.9: Neutrino factory front end, with target, pion capture and beam manipulation.

eration chain without some reduction in the form of cooling. The effect of cooling in the neutrino factory front end can be seen in Figure 1.10, with an increasing transmission of muons.

The cooling section will take the form of an extended channel approximately 300 m in length. Ionisation cooling is the method chosen for muons, and is discussed in more detail in Section 1.4.5.

Acceleration and Storage

Once the muon beam has been captured and is manipulated to meet the acceptance and phasing requirements of the downstream accelerator, the particles can be brought up to the desired energy and stored in decay rings. The acceleration will comprise a linear pre-accelerator followed by successive chains of re-circulating linear accelerators (RLA). The current design energies would be 0.9 GeV in the pre-linac, 3.6 GeV in the first re-circulating linac and 12.6 GeV in the second. Additional acceleration up to 50 GeV may then be performed with FFAGs. A schematic is given in Figure 1.11.

One issue with the RLAs lies in the arc sections. As the particles gain energy with each pass through the linac, the dipoles in the arcs can no longer successfully

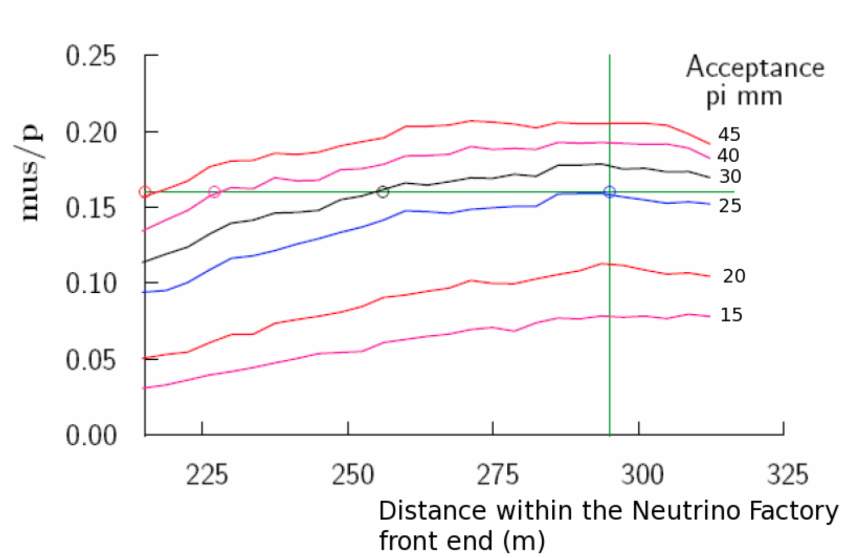

Figure 1.10: Muons per incident proton as a function of the cooling channel length for various downstream acceptances. The x-axis shows the distance from the target (m).

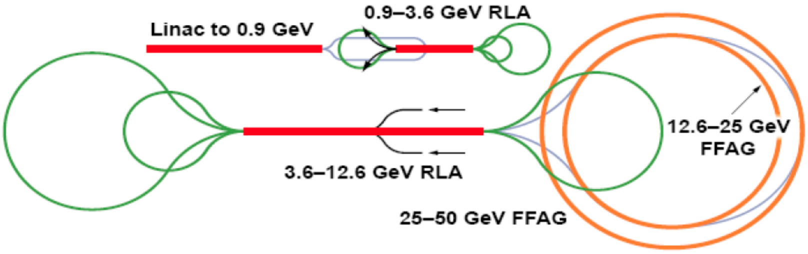

Figure 1.11: Layout of the downstream muon acceleration scheme.

bend the beam, generating the need for multiple arcs designed for different particle energies. To mitigate this need, FFAGs could be used in the arc sections as well.

The race track shaped decay ring has two long, straight sections with 180° curves. The ratio of the straight lengths to the curved sections should be maximised to ensure the majority of the muons decay whilst travelling along the desired baseline. With baselines considered at 7500 km and 3500 km, the straight sections would need to be tilted at large angles. The μ^+ and μ^- could either be injected into separate rings, or counter-rotated in the same ring. The time difference between the decays would then be $\mathcal{O}(100\text{ ns})$.

1.3.3 Detectors

The $\nu_e \leftrightarrow \nu_\mu$ channel in a Neutrino Factory benefits from a highly accurate measurement of energy provided by a calorimeter design, either magnetised iron or totally active scintillator. Current neutrino baseline experiments such as T2K rely on the large mass of a water Cerenkov detector, and this could be extended to larger volumes and photo-multiplier tube (PMT) surface areas to be used as the far detector in a Neutrino Factory. Liquid argon time projection chambers provide scope for high resolution tracking, and emulsion cloud chambers similar to those employed by OPERA give the opportunity for identifying the τ neutrino decay vertices. The choice of detector depends on the desired channel, as well as cost and engineering constraints.

Water Cerenkov

Water Cerenkov detectors for a neutrino factory would see a scale increase on existing large detectors such as Super-Kamiokande. The Hyper-K detector would consist of 5 optically separate cylindrical water volumes, each 50 m high with a diameter of 43 m. A separate design would involve cubic volumes of $60\times60\times60\text{ m}^3$ surrounded by 57000 inward facing 20” photomultipliers, with an outer region 2.5 m thick instrumented with 15000 8” PMTs to act as a background veto.

Magnetised Iron

A large magnetised iron detector would be similar in design to the current MINOS experiment on a larger scale. The Magnetised Iron Neutrino Detector (MIND) consists of 4 cm iron plates interspaced with a 1 cm active layer, covering a volume of $14 \times 14 \times 20\text{ m}^3$ for a total mass of 60 kton. Solid or liquid scintillator layers, with a position resolution of 1 cm, placed within a strong magnetic field will give momentum and charge information. The field itself could be either a toroid or dipole.

The high stopping power of the iron, together with good momentum resolution obtained from calorimetry of the muon and hadron showers (where the muon itself has been identified by its penetration), make this detector a viable choice. However, it may lack the sensitivity to lower energy muons to make an accurate measurement of CP violation. High momentum muons are a primary source of error given charge misidentification from their curvature in the magnetic field.

Totally active scintillator

A totally active calorimeter would involve a $15 \times 15 \times 100\ \mathrm{m}^3$ volume of scintillating bars (approximately 22.5 kton), placed within a 0.5 T magnetic field. Simulations of the detector suggest a position resolution of 4.5 mm with a momentum resolution of 10% uniformly across the momentum range. The low density allows for efficient reconstruction of momentum down to approximately 600 MeV/c.

Liquid Argon

Liquid argon (LAr) provides the possibility to construct a high resolution drift chamber with a large active volume, given the long drift length of charges produced in ionising tracks. Timing is also obtained by the prompt arrival of the scintillation light. A Neutrino Factory far detector would require a LAr mass of 100 kton, to which end additional R&D is required, looking at both the feasibility of the argon production and charge detection devices.

Emulsion cloud chamber

Sensitivity to the $\nu_e \to \nu_\tau$ channel is missing in large resolution detectors due to their inability to identify the τ vertex before it decays. Fine layered emulsion detectors such as OPERA provide this feature, although suffer from a slow readout time. Adopting the OPERA geometry, an Emulsion Cloud Chamber would have $44\,\mu$m films sandwiched by $205\,\mu$m plastic sheets, at a total of 4 kton. A magnetised option is also under consideration.

1.4 Neutrino Factory R&D

The technological challenges in building a neutrino factory are diverse, and a number of experiments are underway to solve these problems.

1.4.1 Front End Test Stand

Based at the Rutherford Appleton Laboratory, the Front End Test Stand (FETS) [28] is a proton driver experiment to demonstrate a low energy chopper, which will create a longitudinal bunch structure to match any downstream RF period. It consists of a H^- ion source, a low energy beam transport, acceleration from a radio frequency quadrupole and a medium energy beam transport and chopper. The chopper uses a fast-low combination for a fast rise time and long flat-top field. This allows for particles outside the desired bunch structure to be kicked out of the beam. With the addition of the medium energy transport, formed of a channel of 11 quadrupoles, FETS simulations suggest a chopping efficiency of approximately 99%.

1.4.2 Target Studies

The MERIT experiment looked at the technological viability of a liquid mercury target. RAL also has a program looking at solid targets by pulsing high voltage and current through thin wires of various elements, investigating the thermal and mechanical stresses the materials would undergo [29]. Simulations are also being used to investigate the energy deposition on the liquid mercury target and the surrounding target area, including the solenoidal focusing coils [30].

1.4.3 MuCool

One of the key challenges for an ionisation cooling channel is the requirement to operate high gradient RF fields in a strong magnetic field [31]. The field not only limits the achievable voltage, but can cause electron discharges to repeatedly hit the same point on a cavity, causing damage. The MuCool program centred at Fermilab [32] aims to investigate the issues surrounding operating RF in a magnetic field.

1.4.4 FFAGs

Fixed Field Alternating Gradient (FFAG) accelerators are proposed at various stages in a neutrino factory, both in the proton driver and the muon acceleration. FFAGs have the advantage of a very large acceptance and the ability to circulate and accelerate particles of differing momenta. FFAGs have the additional advantage of not needing to scale the magnetic field over time, as in a synchrotron, saving on both power supply cost and the duty factor involved in ramping the fields.

The EMMA experiment [33][34] at Daresbury Laboratory, UK, aims to be the first non-scaling FFAG. Its 16 m circumference contains 42 cells, producing an electron beam of 10-20 MeV, for which beam commissioning is currently underway [35].

1.4.5 Cooling

The cost of producing accelerator components with an acceptance (the phase space aperture through which particles can pass) large enough to accommodate the muon beam is a considerable issue for a Neutrino Factory. To reduce the necessary acceptance, some form of beam cooling is required. Methods exist for reducing the phase space occupied by a beam, but these operate over a much longer period than the lifetime of a muon. These consist of:

- **Adiabatic damping**
  In electron rings, as the beam is bent around a circular accelerator the electrons lose energy by synchrotron radiation. This acts as a form of cooling and reduces the phase space density of the beam. This is unsuitable for a muon as synchrotron radiation is heavily dependant on the particle mass.

- **Electron cooling**
  By introducing a mono-chromatic electron beam for short sections in an ion accelerator, the ions will lose energy by Coulomb interactions until the two beams are in thermal equilibrium [36]. The electrons can then be removed resulting in a net reduction in momentum for the ions, which will be re-accelerated longitudinally.

- **Stochastic cooling**
  A fast kicker magnet is used to push high transverse momentum particles back into the beam bunch [37]. It is a slow, iterative process which can take of the order of minutes to be effective.

Ionisation cooling is proposed as a fast alternative to these methods capable of reducing the emittance of a muon beam. This is discussed in more detail in Section 1.5.

1.5 Accelerator Concepts

The motion of particles in an accelerator resemble that of a forced harmonic oscillator, with a sinusoidal transverse motion, effected by the magnetic fields, which has a periodic dependance on how these fields change along the accelerator. Different magnet types affect the particles in different ways, often with opposing effects in the vertical and horizontal planes, as is the case with dipole and quadrupole magnets. Dipole magnets bend the particle according to its charge and momentum in one plane, with no effect (to first order) in the other plane, and are used to control the direction of a particle around the curvature of an accelerator chain. Quadrupole magnets allow for the focusing and defocusing of a beam, with the effect being opposite in the planes in the same magnet. They are used to stop the beam from diverging and being lost. Higher pole magnetics, such as sextupole and octupole, can be used to compensate for undesired effects in the quadrupoles. Solenoids allow for focusing in both planes simultaneously but with a lesser strength than quadrupoles.

Longitudinally, particles undergo acceleration by feeling the force of an electric field. To repeat this force at a steady rate, Ising [38] and Wideroe [39] proposed to use a time varying field across a cavity, referred to now as radio frequency (RF) cavities. This allows for a particle to see the effect of many high gradient electric fields over a short distance, without feeling any opposing effects. It is best, then, to separate the concepts of particle accelerators into distinct areas: the transverse effects, the longitudinal effects, the behaviour of particle ensembles and the effect of material interactions.

Transverse Dynamics

The motion of a charged particle in an electro-magnetic field is defined via the Lorentz force as

$$\frac{d\vec{p}}{dt} = e(\vec{E} + \vec{v} \times \vec{B}) \tag{1.7}$$

which originates in Lagrange's equation where the Lagrangian is

$$L = -\gamma mc^2 - e\Phi + e\vec{v}.\vec{A} \tag{1.8}$$

The Hamiltonian can then be defined in Equation 1.9 where $\vec{P} = \partial L/\partial \vec{v} = \vec{p} + e\vec{A}$ is the canonical momentum derived from the mechanical momentum $\vec{p}$. This gives a generic Hamiltonian for a charged particle in an electro-magnetic field, represented in canonical coordinates.

$$H = \vec{P}.\vec{v} - L = c[m_2c^2 + (\vec{P} - e\vec{A})^2]^{1/2} + e\Phi \tag{1.9}$$

By transferring the Hamiltonian into the coordinate system for a reference particle (the Frenet-Serret coordinate system) the motion along the path of the particle, s, can become the independent variable instead of time [40]. The Hamiltonian is then related to the conjugate momentum in the coordinate, $-p_s$. The Hamiltonian is then given by

$$\tilde{H} = -\left(1+\frac{x}{p}\right)\left[\frac{(H-e\phi)^2}{c^2} - m^2c^2 - (p_x - eA_x)^2 - (p_y - eA_y)^2)\right]^{1/2} - eA_s \tag{1.10}$$

Defining Hamilton's equations as $x' = \frac{\partial \tilde{H}}{\partial p_x}$ and $p'_x = -\frac{\partial \tilde{H}}{\partial x}$, and extracting these terms from the Hamiltonian, the motion of the particle due to the transverse motion in magnetic fields can be given by Equations 1.11 and 1.12, which simplify to Equation 1.13, known as Hill's equation, where $k(s)$ is a periodic function attributed to the magnetic focusing.

$$x'' - \frac{\rho + x}{x^2} = \pm\frac{B_y}{B\rho}\frac{p_0}{p}\left(1+\frac{x}{p}\right)^2 \tag{1.11}$$

$$y'' = \pm\frac{B_x}{B\rho}\frac{p_0}{p}\left(1+\frac{y}{p}\right)^2 \tag{1.12}$$

$$x'' + k(s)x = 0 \tag{1.13}$$

The solution to Hill's equation is shown in Equations 1.14 and 1.15. It is a periodic trajectory with an amplitude defined by a constant emittance and the betatron function $\beta(s)$, a parameter of the magnetic lattice, and the phase of the particle, $\phi(s)$, with respect to the phase of a reference particle ϕ_0.

$$x = \sqrt{\beta(s)\epsilon}\cos(\phi(\mathrm{s}) + \phi_0) \tag{1.14}$$

$$x' = \sqrt{\frac{\epsilon}{\beta(s)}}\sin[\phi(\mathrm{s}) + \phi_0] + \left[\frac{\beta'(\mathrm{s})}{2}\right]\sqrt{\frac{\epsilon}{\beta(\mathrm{s})}}\cos[\phi(\mathrm{s}) + \phi_0] \tag{1.15}$$

From this, the position and momentum of a particle depends on the changing value of the β function, which is described as being a function of the magnetic arrangement, or lattice. Further parameters can be defined as

$$\alpha = -\frac{\partial\beta}{\partial s} \tag{1.16}$$

$$\gamma = \frac{1+\alpha^2}{\beta} \tag{1.17}$$

where α, β and γ are the Twiss functions. The variable β is a function of the magnetic lattice defining the envelope in which the beam propagates. The variables α and γ are functions of β which effect the orientation of the phase space ellipse shown in Figure 1.12.

Thus, depending on the value of the betatron function and phase, ϕ, the position and divergence of the particle can be found to follow an ellipse, as shown in Figure 1.12 [41]. This ellipse can be defined by the Courant-Snyder invariant as

$$\epsilon = \gamma x^2 + 2\alpha x x' + \beta x'^2 \tag{1.18}$$

and so the position and momentum of a particle can be traced through the magnetic lattice.

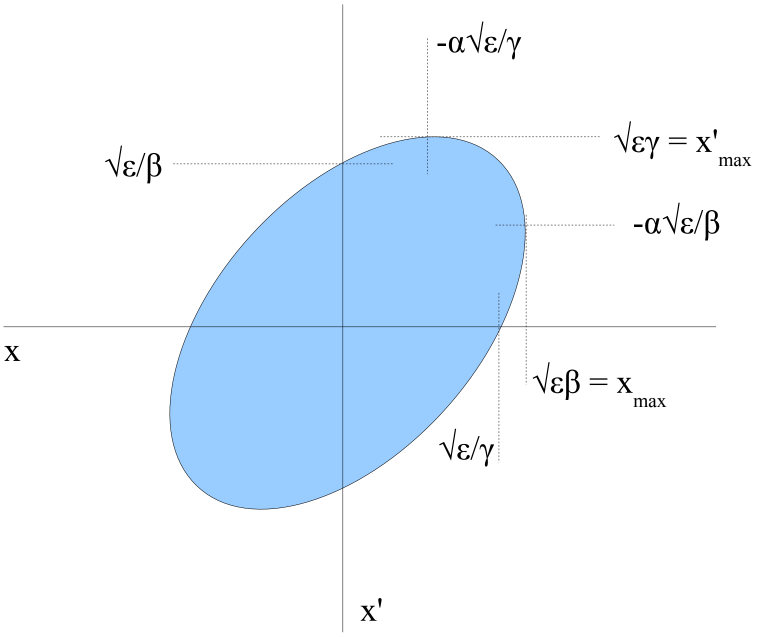

Figure 1.12: Phase space ellipse of a set of particles, where the area of the ellipse is equal to $\pi\epsilon$.

Transfer Matrices

As the transformation of a particle's position in phase space from one point to another in the periodic lattice is linear, it can be defined in terms of a matrix. Given a particle state (x, x'), the transformation is given by:

$$\begin{pmatrix} x \\ x' \end{pmatrix}_2 = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix}_2 \begin{pmatrix} x \\ x' \end{pmatrix}_1 \tag{1.19}$$

By defining μ as the phase advance per cell, the common transfer matrix in a periodic lattice, the Twiss matrix, is given by:

$$M = \begin{pmatrix} cos\mu + \alpha sin\mu & \beta sin\mu \\ -\gamma sin\mu & cos\mu - \alpha sin\mu \end{pmatrix} \tag{1.20}$$

This can be expanded to map the Twiss parameters between two points, as given by:

$$\begin{pmatrix} \beta \\ \alpha \\ \gamma \end{pmatrix}_2 = \begin{pmatrix} M_{11}^2 & -2M_{11}M_{12} & M_{12}^2 \\ -M_{11}M_{21} & M_{11}M_{22} + M_{12}M_{21} & -M_{12}M_{22} \\ M_{21}^2 & -2M_{21}M_{22} & M_{22}^2 \end{pmatrix} \begin{pmatrix} \beta \\ \alpha \\ \gamma \end{pmatrix}_1 \tag{1.21}$$

The use of such matrices is extremely useful in tracking the phase space state of a particle or beam between points in an accelerator.

Longitudinal Dynamics

The acceleration is performed by passing the charged particle across a voltage difference. In radio frequency (RF) systems, the voltage is varied as a sine wave such that the voltage seen by the particle is given by

$$V = V_0 \mathrm{sin}\phi_\mathrm{s} \tag{1.22}$$

Circular accelerators are set up such that the frequency of the RF wave is an integer multiple, or harmonic number, of the revolution rate of the particles around the ring, and the particles are grouped into bunches so that they all arrive on the accelerating edge of the RF wave.

A reference phase ϕ_s can be defined where a particle initially at this phase will always arrive at the RF cavity when the voltage is the same value. Around this phase the slope of the sine wave ensures that early particles will see less voltage, and be relatively decelerated, and late particles will see more voltage and be relatively accelerated, moving them all towards the reference phase. The region in which this holds is known as the RF bucket. Particles falling outside the bucket will be continuously decelerated until the trailing bunch absorbs them. The size of the bucket depends on the proximity of the reference phase to 90°, where the bucket size drops to zero. The RF bucket and resulting change in energy is shown in Figures 1.13 and 1.14.

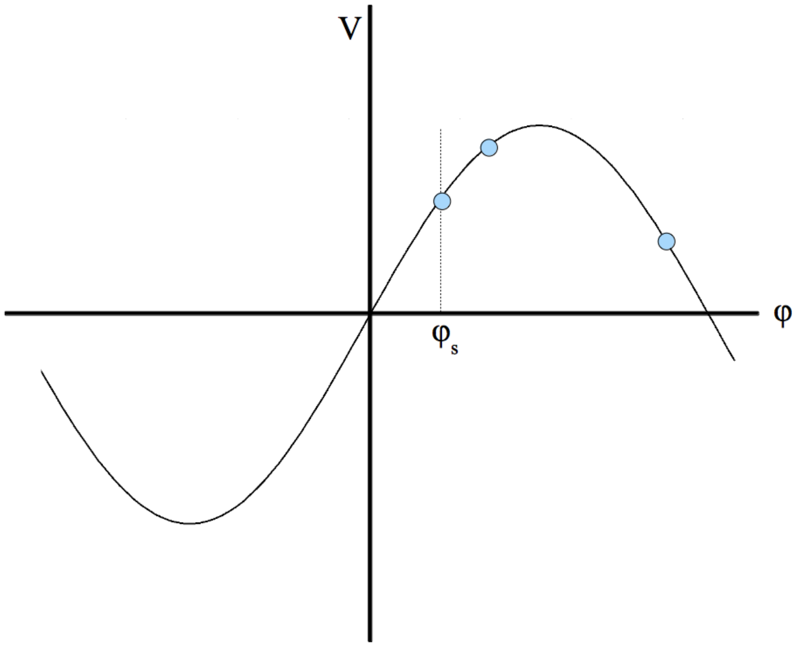

Figure 1.13: Voltage seen by particles inside and outside the RF bucket. ϕ is the phase of particle with respect to the RF wave, and ϕ_s a reference phase.

The reference phase in Figure 1.13 refers to the stationary central point in Figure 1.14, with particles separated from the reference phase but still within the RF bucket moving in an ellipse. The particle at a phase greater than $\pi - \phi_s$ is then pushed further from the reference phase and lost.

Emittance

From the elliptical nature of particle motion in a magnetic lattice, the amplitude of a single particle can be defined as the single particle emittance. This is equivalent to the Courant-Snyder invariant.

For many particles, the emittance of the beam can be defined in terms of the variances of the beam, such that it is equivalent to the determinate of the covariance matrix. This is given in Equation 1.23, where n is the dimension of the beam, which corresponds to the area of the ellipse shown previously.

$$\epsilon^n = det(\Sigma_n) \tag{1.23}$$

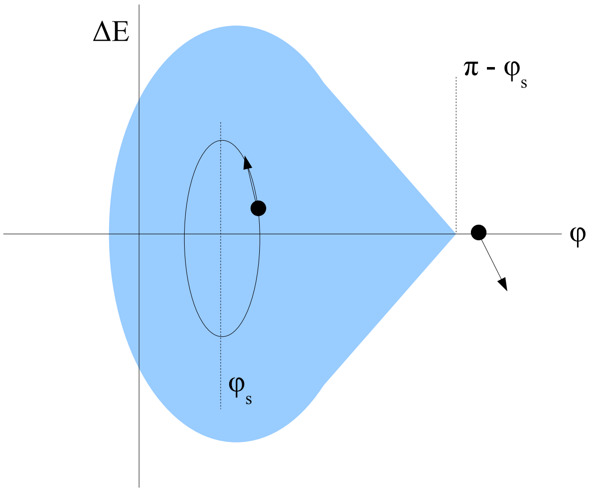

Figure 1.14: Change in energy for particles inside and outside the RF bucket. ϕ is the phase of particle with respect to the RF wave, and ϕ_s a reference phase.

Given the elliptical motion of the particle and the relationship between the physical coordinates and the Twiss parameters, the covariance matrix is also related and can be written as:

$$\Sigma = \begin{pmatrix} \sigma_x^2 & \sigma_{xx'} \\ \sigma_{xx'} & \sigma_{x'}^2 \end{pmatrix} = \epsilon \begin{pmatrix} \beta & -\alpha \\ -\alpha & \gamma \end{pmatrix} \tag{1.24}$$

This means that the emittance refers to volume of the beam in phase space and is closely related to the covariances of the beam particle ensemble and the Twiss parameters. These values can then be derived from the covariances of the beam.

Liouville's theorem

From the solution to Hill's equation, the particle trajectory depends on the constant emittance. The invariance of the emittance through the magnetic lattice derives from Liouville's theorem, which states that the area A within a closed contour in phase space is conserved, as defined in Equation 1.25, where p and q are the Hamiltonian canonical co-ordinates.

$$A = \int p \, dq \tag{1.25}$$

As the area of a phase space ellipse is related to the beam emittance via $A = \pi\epsilon$, Liouville's theorem states that emittance is conserved in a magnetic lattice, in the absence of material. This remains true despite any level of focusing or bending.

When material of radiative effects are included, Liouville's theorem does (or appears to) break down. These are of interest to deliberately reducing the beam emittance and are discussed later.

Given the definitions supplied in the transverse dynamics section of this chapter, it is clear that as the beam is accelerated, the divergence will decrease, leading to what may appear as transverse emittance reduction and the violation of this theorem. This is because the coordinates used so far are not canonical. In order to properly scale the emittance with acceleration it must be multiplied by a relativistic $\beta\gamma$ factor. This then defines the normalised emittance, as:

$$\epsilon_N = \beta\gamma\epsilon \tag{1.26}$$

It is this normalised emittance which is of interest to studies of emittance reduction. Scaling the emittance in this way is valid when using the trace space or angular co-ordinates $x' = dx/dz$. When momentum coordinates are used the normalised emittance is found in a different way, such as that given in the following section.

Solenoidal Dynamics

Designs of neutrino factory cooling channels rely on the use of solenoids to transport and focus the beam. Labelling a beam which meets the solution for a solenoidal lattice as 'matched', it becomes necessary to define how a matched beam in a solenoid should behave.

The covariance matrix Σ for such a beam is given by Equation 1.27 [43]. Here the Twiss functions and emittance remain as β, α, γ and ϵ, with L the average canonical angular momentum and κ the magnetic focusing of the solenoid.

$$\Sigma = \begin{pmatrix} \frac{\epsilon_N \beta m c}{P_z} & & & \\ -mc\epsilon_N\alpha & mcP_z\epsilon_N\gamma & & \\ 0 & mc\epsilon_N(\beta\kappa - L) & \frac{\epsilon_N \beta m c}{P_z} & \\ mc\epsilon_N(\beta\kappa - L) & 0 & -mc\epsilon_N\alpha & mcP_z\epsilon_N\gamma \end{pmatrix} \tag{1.27}$$

This matrix satisfies the condition that $\epsilon_N^4 = det(\Sigma)/m^4$, where m is the particle mass, from which the normalised emittance can be extracted. With the MICE lattice designed so that any beam matched in the first solenoid should be transported through the cooling channel, the covariance matrix of the beam in the first spectrometer solenoid should match this form.

Ionisation Cooling

To manipulate the transverse and longitudinal emittance of a muon beam a new method is required. Ionisation cooling uses an ionising medium to reduce the total momentum of a particle before reaccelerating longitudinally with RF cavities. By constructing absorber modules within strongly focusing solenoids, the β function can be manipulated to achieve the maximum level of cooling, as shown by Equation 1.28. Following each absorber with an RF cavity capable of restoring the lost energy is a requirement not so much for the reduction of divergence but to maintain the beam transmission. The RF cavities must themselves be placed within solenoidal fields to maintain the lattice and couple the optical functions between the focusing absorber modules. A diagram of the energy loss and reaccelerating phases is given in Figure 1.15.

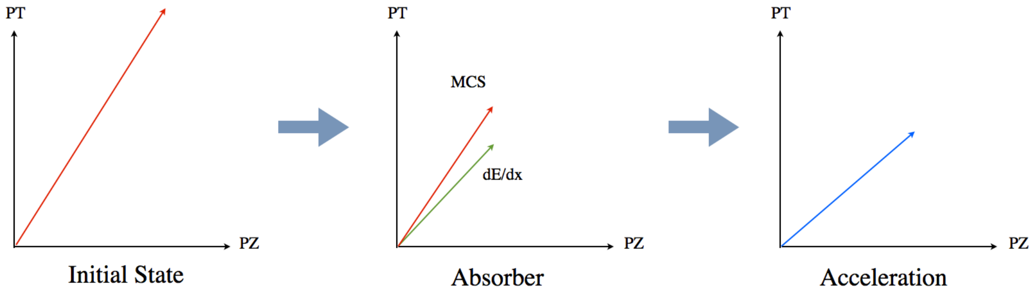

Figure 1.15: Changes to the transverse and longitudinal momenta of a particle at varying stages in an ionisation cooling cell.

The initial, high transverse momentum state undergoes energy losses in a medium which reduce momenta in all three dimensions, as well as some multiple scattering which increases the transverse momentum. In order to achieve effective cooling, the level of multiple scattering must be minimised, which imposes restrictions on all sections of the cooling channel apparatus. After each absorber section the particle is re-accelerated by RF cavities. The gradient and frequency of the cavities varies between setup designs, but 16 MV/m and 201 MHz are typical values.

The rate of change of the normalised emittance in a medium is given by Equation 1.28,

$$\frac{d\epsilon_N}{dX} = -\frac{\epsilon_N}{\beta^2 E}\langle\frac{dE}{dX}\rangle + \frac{\beta_t(0.014\mathrm{GeV})^2}{2\beta^3 \mathrm{E m_\mu X_0}} \tag{1.28}$$

where ϵ_N is the normalised emittance, β the relativistic beta, β_t the transverse β function and X_0 the radiation length of the absorber material. The point at which the cooling due to energy loss and the heating due to multiple scattering is equal occurs at an equilibrium emittance, given by:

$$\epsilon_{eq} = \frac{\beta_t(0.014GeV)^2}{2\beta X_0}\langle\frac{dE}{dX}\rangle^{-1} \tag{1.29}$$

From this, it can be seen that the equilibrium emittance is proportional to the β function at the interaction point. To achieve a minimum equilibrium emittance, the β function should be minimised. This creates a case for strongly focusing the beam in the absorber material.

Particle Interactions

Since the MICE cooling channel intentionally forces the beam to interact with material, the effects of charged particles in various materials must be considered. Due to the realistic engineering requirements of constructing the cooling channel, material other than that in the absorber will be presented to the beam. The processes of interest are the level of energy losses and multiple Coulomb scattering.

For ionising particles, the mean energy loss as a function of the distance in a material dE/dX is given by the Bethe-Bloch equation:

$$-\frac{dE}{dX} = Kz^2\frac{Z}{A}\frac{1}{\beta^2}\left[\frac{1}{2}\ln\frac{2m_ec^2\beta^2\gamma^2T_{max}}{I^2} - \beta^2 - \frac{\delta}{2}\right] \tag{1.30}$$

Charged particles interacting in a medium undergo many small angle deflections by means of Coulomb scattering. The effect of these combined processes (multiple Coulomb scattering) can be defined by the angle θ_0 which represents 98% of the resulting Gaussian distribution of angles. The value is valid up to this region, where it is a fit to Moliere scattering, and is given by:

$$\theta_0 = \frac{13.6MeV}{\beta cp}z\sqrt{\frac{x}{X_0}}\left(1 + 0.038\ln\frac{x}{X_0}\right) \tag{1.31}$$

1.6 Aims of this Thesis

With the need for ionisation cooling laid out, a range of research areas present themselves within the context of the Muon Ionisation Cooling Experiment (MICE), on which this thesis is based. A particular need of demonstrating ionisation cooling is the ability to accurately measure the emittance of a muon beam. This thesis will therefore investigate the use of a scintillating fibre tracker to measure the momentum and position of individual particles, and by explaining the calibration, reconstruction and analysis procedures, demonstrate the tracker's viability for the necessary measurements.

It is also the goal of this research to use the simulated performance of MICE to investigate the details of ionisation cooling. In particular, methods of manipulating the data from MICE to investigate subtle effects of ionisation cooling will be looked into, with attention given to methods of particle weighting. Finally, given the documented use of transfer matrices to describe the behaviour of a beam in an accelerator, a period of research was devoted to investigating whether such a method could be employed for ionisation cooling, using neural networks to find the parameters of the matrix. It is the aim of this thesis to show that this is a viable option for parameterising ionisation cooling.

Chapter 2

The Muon Ionisation Cooling Experiment

2.1 Overview

Given the requirements for a low emittance for muon beams in a Neutrino Factory, some method of reducing the phase space volume must be employed. Current technologies for emittance reduction, such as electron and stochastic cooling, do not operate at the required rate for dealing with muons given their short lifetime, and so a new method must be employed. The Muon Ionisation Cooling Experiment (MICE) aims to prove the viability of ionisation cooling, discussed in Chapter 1, by measuring a 10% reduction in the transverse emittance to within 1%, requiring an absolute measurement on the emittance to within 0.1%.

MICE is based at the Rutherford Appleton Laboratory, UK, where a proton synchrotron used for a neutron spallation source (ISIS) is also used for the production of muons. A parasitic target is dipped into the beam and pions are produced by the interaction of high energy protons with the target. The pions then decay in the MICE beamline into muons.

A collection of detectors then measures the time, position and momentum of each individual muon before and after the cooling channel with the single particle amplitudes contributing to an overall beam emittance. MICE is unique in its ability to measure emittance on a single particle basis with mm-MeV-ps resolutions.

2.2 Target and Beamline

2.2.1 ISIS

The ISIS proton synchrotron consists of a 35 keV H^- ion source which feeds into a radio frequency quadrupole (RFQ). The RFQ operates at 202.5 MHz accelerating the ions into bunches 5 ns long at an energy of 665 keV before a linear accelerator increases the energy of the beam to 70 MeV. After injection into the main ISIS accelerator ring, the beam is formed into two bunches which are then subjected to an energy gain of approximately 140 keV per revolution over several thousand orbits. The resulting 800 MeV proton beam is then extracted into two beamlines which transport the protons to the tungsten targets. This process is repeated at a rate of 50 Hz.

Whilst in the main accelerator ring, the proton bunches are around 100 ns long at maximum acceleration, with a repetition rate of approximately 3 MHz, with two bunches in each pulse crossing a specific location in the accelerator (such as the MICE target). This will form the microstructure of the MICE beam which is of concern to the tracker electronics discussed later.

2.2.2 Target

MICE runs parasitically in the ISIS beam by mechanically inserting a titanium target into the proton beam at peak acceleration, and transporting the resulting particles down a separate beamline and into the MICE experimental hall. A linear motor drives a hollow titanium tube into the ISIS beam pipe at an acceleration of 80 g, catching one pulse at peak acceleration but avoiding the injection phase of the subsequent pulse. Pions are then produced by the hadronic interactions of the protons with the high Z target; a selection of particles coming off the target are then captured by the beamline.

The effective cross section presented by the target to the beam is proportional to the depth into the beam to which the target is driven. Concerns of running in this mode consist of the total energy deposition in the target (although given the reasonably low repetition rate, this is not a large problem), and the flux of secondary particles. Whilst some secondaries are collected by the MICE beamline, the vast majority are absorbed elsewhere in the ISIS synchrotron vault, leading to possible activation of accelerator components. The rate of these secondaries seen inside the ISIS vault is therefore of concern. This effect can be measured with ionising beam loss monitors used in the normal running of ISIS. Some beam loss is expected during the regular injection-acceleration-extraction cycle and is taken into account.

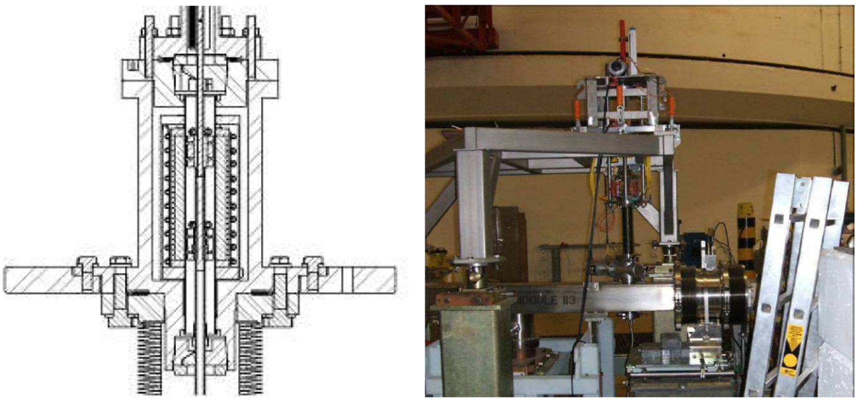

Figure 2.1: MICE Target.

However, as the MICE target pulses into the beam the losses in the region of the MICE target increase. The losses due to the target are thus kept below a set limit.

There is a large mechanical challenge in constructing a proton target which must scrape an accelerating beam, pulsing into a vacuum-sealed beam pipe with minimal beam losses. Several iterations of the target have been developed to account for material being deposited in the beam pipe and fault protection. The control electronics, which relies on an optical measurement of the target depth, has also been upgraded over the course of running. Typically, an identical target to that used in MICE is run in a separate lab at RAL for many more pulses than that in ISIS in order to predict failure rates and investigate failure modes.

2.2.3 Beamline

Particles are transported from the target in the ISIS synchrotron hall to the cooling channel by a chain of quadrupoles, dipoles and solenoids. In the ISIS hall, a quadrupole triplet captures the particles from the target and transports them to two dipoles which bend the beam into the MICE hall; a superconducting solenoid, referred to as the decay solenoid, concentrates the beam into the MICE hall. By this point, most of the pions are expected to have decayed into muons. Two further quadrupole triplets transport the beam to the point at which the cooling channel is expected to begin and match the beam parameters to those required for the cooling channel [44].

The first dipole bends the beam through 60° into a superconducting solenoid with a field of 5 T. The solenoid is 5 m long and results in a second dipole which

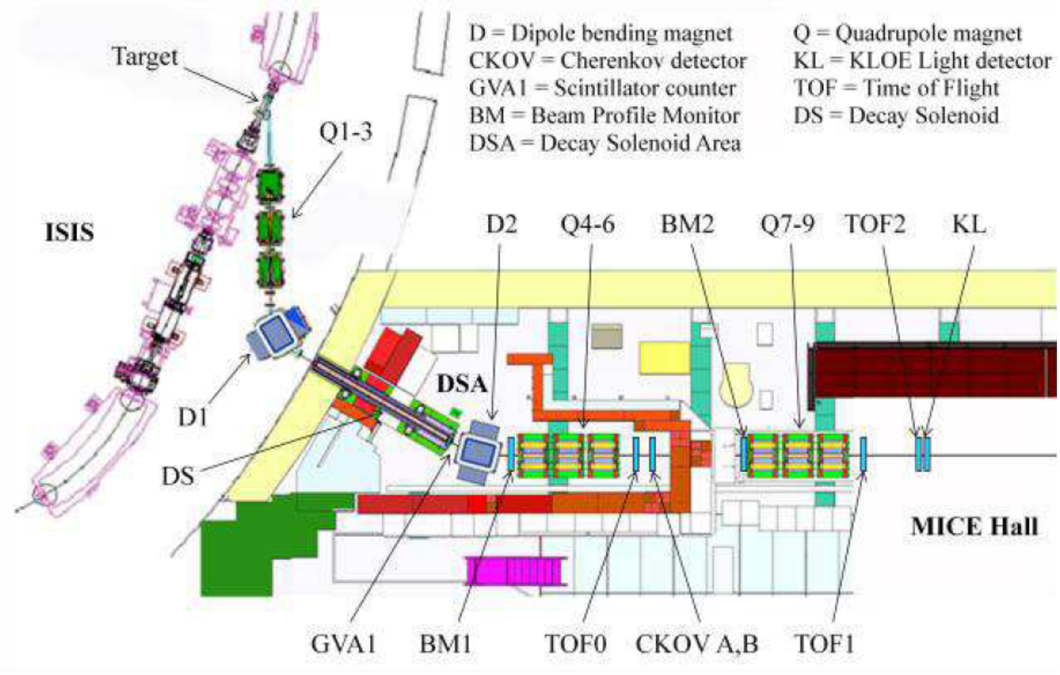

Figure 2.2: MICE beamline from ISIS into the experimental hall.

bends the beam into the MICE hall. This is shown in Figure 2.2. By setting the fields of the two dipoles the momentum and type of the MICE beam can be selected. Given that the momentum of a particle is proportional to the field and angle of a dipole, for a known angle the momentum can be selected by the field. In one mode, the momentum selection range of the first dipole is twice that of the second dipole. This causes muons which have been created between the dipoles to be transported only if their momenta is half that of the initial pion, meaning they had been produced travelling backwards in the pion rest frame. Pions in this mode would not be transported through the second dipole which has a different momentum selection. This is referred to as 'muon mode'. When the momentum selection of the two dipoles is equal only the initial pions should be selected, providing a pion mode to the beamline.

Aside from the decay solenoid, all the magnets in the MICE beamline are normal conducting. The quadrupoles are of two types, the first used in the ISIS synchrotron vault and the second in the MICE hall, both of which have been inherited from the NIMROD experiment previously held at RAL. A summary of the beamline can be found in Table 2.1.

Magnet	Position (mm)	L_{eff} (mm)	Max. B (T)	Max. Gradient (T/m)
Q1	3000.0	854.3	-	2.1
Q2	4400.0	854.3	-	2.1
Q3	5800.0	854.3	-	2.1
D1	7979.1	1038	1.6	-
DS	12210.7	5000	5.0	-
D2	15808.1	1038	0.85	-
Q4	17661.6	660	-	2.3
Q5	18821.6	660	-	2.3
Q6	19981.6	660	-	2.3
Q7	25293.7	660	-	2.3
Q8	26453.7	660	-	2.3
Q9	27613.7	660	-	2.3

Table 2.1: Summary of the MICE beamline elements.

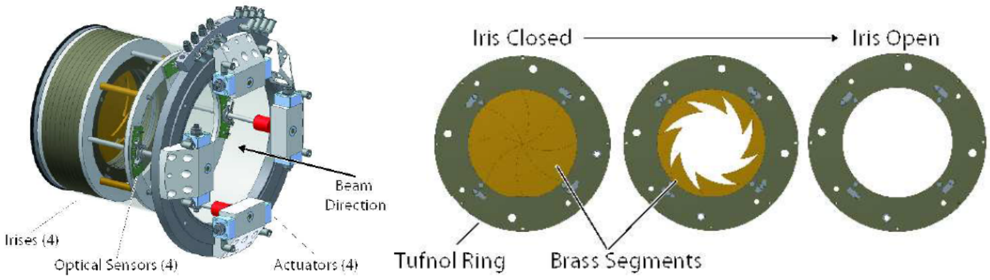

Figure 2.3: The MICE Diffuser. The iris mechanism of changing the thickness can be seen on the right of the figure, with the full structure seen on the left.

Diffuser Plate

In order to deliberately inflate the muon beam's emittance to predefined values, a changeable amount of material can be positioned before the cooling channel. This takes the form of a selection of lead plates which can be combined to give a range of radiation lengths. The plates themselves form the blades of an iris which move into the beam's path and present a varied amount of material to the muons. The diffuser is shown in Figure 2.3. The thicknesses of the diffuser required for generating the target emittance values for each beam is included in Table 2.2.

Proton Absorber

One source of contamination in the beam has been observed to be protons from the ISIS beam being deflected in the MICE beamline. Whilst protons can be easily identified by their long time of flight in comparison to muons they still form an undesirable component of noise which can negatively contribute to detector dead time. In order to reduce this contamination, variable thicknesses of plastic are placed into the beamline upstream of the detectors to absorb the protons but allow the minimum ionising muons to pass through. The optimum thickness for the momentum of the beamline was chosen and shown to reduce the proton contamination [45].

Magnet Settings

The currents of the beamline magnets have to be optimised for several reasons, which include maximising the particle transmission to the cooling channel and matching the optical parameters of the beam to the cooling channel lattice. A zero α function is desired to make a constant measurement of the beam covariance matrix at the measurement planes. With a non-zero α function the physical size of the beam changes with z, which complicates the process of measuring a helix in the trackers.

The upstream (inside the ISIS hall) and downstream (inside the MICE hall) quadrupole settings were computed with a combination of tracking codes [46]. These included TRANSPORT, TURTLE, G4Beamline and G4MICE. Tracking codes typically operate with either a matrix or particle tracking method. A matrix method involves computing the transfer matrices required to transport the optical functions as desired between two points. Full tracking codes, such as Geant based codes like G4Beamline and G4MICE, track individual particles through generated magnetic fields and materials. The benefit of tracking codes lies in a more complete physics model, but the computational cost, especially for large, multi-turn accelerators, can be considerable. Beam-beam effects such as space charge can also be considered in full tracking codes, although space charge will not be an issue in MICE given the low rates. An example of the calculated optical functions for the MICE quadrupole optics can be seen in Figure 2.4. The lines in the figures show the horizontal and vertical beam envelopes generated by the MICE beamline magnets, with the nine quadrupoles focusing and defocusing along with the decay solenoid.

The upstream beamline is intended purely to transport the maximum number of particles to MICE. Using the optical parameters found with studies from a TRANSPORT simulation, a sweep of the upstream quadrupole currents was made and the rates observed in the MICE hall optimised. A scintillating counter placed at the exit of the decay solenoid is used to measure the particle rates. The normalised rates as a function of the triplet current, along with comparison simulation, is shown in Figure 2.5. This shows a reasonable agreement between Monte Carlo and data for finding the upstream magnet settings which maximise the particle rates in the

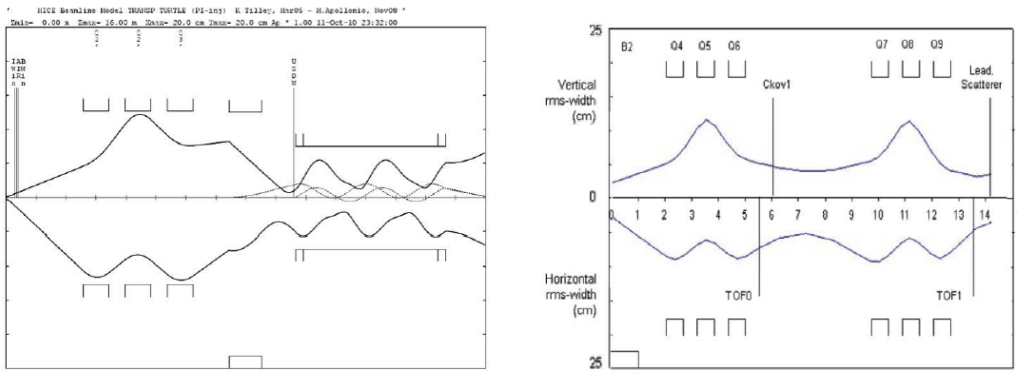

Figure 2.4: Beam envelope of the MICE upstream beamline. The two lines show the horizontal (lower line) and vertical (upper line) RMS of the beam envelope. Quadrupoles 1-3, the first dipole and the decay solenoid are shown in the left-hand figure. Quadrupoles 4-9 are shown in the right-hand figure.

MICE experiment.

Downstream, the optimum current for each setting is found by stipulating the desired Twiss parameters at a point in the beamline (e.g. the diffuser plate) and assuming the nature of the beam originating from the target is described using Monte Carlo studies. The transfer matrices of the intervening magnets can then be found by a suitable fitting method.

In reality, accelerators contain features which are not accounted for in tracking codes. The input beam, which in the simulations may be assumed to have certain symmetries and uniformities, is likely to have suffered many more complicated processes in its interaction with the MICE target. Such inhomogeneities in the beam mean it may encounter non-linear fringe effects in the magnets not accounted for by perfect simulation models. The cumulation of these effects mean that whilst simulations provide a starting point for finding the optimal optics setup, data is still required to complete the process.

The data taking plan of MICE is to measure the emittance change of the beam through the stages of the cooling channel at a number of different momentum-emittance settings. It is intended that the beam will have a set longitudinal momentum and emittance before entering the cooling channel. The initial desired momenta are 140, 200 and 240 MeV/c muons at the tracker reference plane. This will have to include momentum losses in the diffuser and account for the backwards decaying muons, which are selected to improve the longitudinal momentum spread. The key limiting factor in the maximum momentum of the muons depends then on the

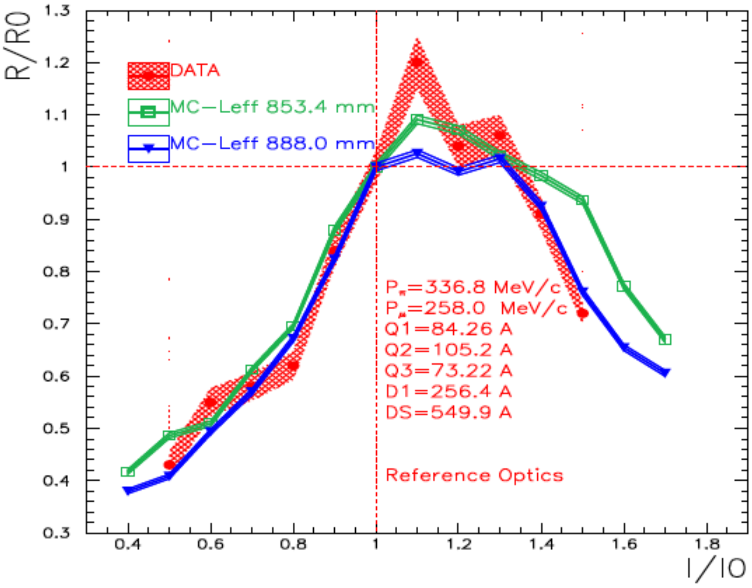

Figure 2.5: Normalised particle rates as a function of Q123 settings. R/R0 and I/I0 refer to the normalised rates and magnet current with respect to the reference values.

desired momentum spread and the achievable bending force of the dipoles.

A data campaign is therefore required to measure the Twiss parameters of the beam at each combination of magnet currents, and optimise according to the desired status of the beam at a given location. As at the time of writing the fibre trackers were not being used in the experiment, this process was performed with Time of Flight (TOF) counters, which naturally involved the decreased resolution of those detectors and some extrapolation to the diffuser plane.

The TOF counters, which give an $x-y$ measurement with a highly accurate time measurement, are placed at three points in the beamline: one prior to the final quadrupole triplet, one after the triplet and another at the end of what will be the cooling channel. By measuring the particle position before and after the quadrupole triplet, and knowing the approximate particle trajectory through the magnets given known momenta, an iterative reconstruction method was developed to find the position and momentum of each particle. The ensemble of reconstructed particles was then used to calculate the beam emittance and Twiss parameters at the second TOF counter. The results are provided in Section 2.7.

The emittance-momentum running modes of MICE found by simulation, along with the momentum at the diffuser and Twiss parameters, are given in Table

2.2.

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

Table 2.2: Optical parameters of the beam for various MICE modes at the diffuser plate, including the thickness of the diffuser t.

2.2.4 MICE Stages

MICE will be constructed and run in several steps, beginning with detector installation and beam characterisation and culminating in the full ionisation cooling channel. The initial operating states were:

- **Step I** - Time of Flight detectors are used to characterise the beam and optimise the magnet settings. This step includes all three TOFs and the Cerenkov detectors.

- **Step II** - The installation of a single tracker enables the accurate measurement of the beam and finalisation of the beamline settings.

- **Step III** - The second tracker is installed and studies of the systematic behaviour of the detectors and tracking can be performed. There is an intermediate step where a small solid absorber is placed between the two tracker solenoids (Step III.i). The Electron-Muon Ranger is also installed at this step.

- **Step IV** - The first absorber and focus coil is installed.

- **Step V** - The first RF cavities are installed. This represents the first full ionisation cooling cell.

- **Step VI** - The final RF cavities and absorber modules are installed to complete the cooling channel.

The implementation of this order is not essential and some of the steps may be combined. At the time of writing Steps II and III will be skipped and operation will commence with Step IV. Characterisation of the tracking detectors will be done simultaneously before data taking with the LH absorbers begins.

2.2.5 MICE hall and infrastructure

The MICE experimental hall is located in a building adjacent to the ISIS target station 1. The beamline from the ISIS synchrotron vault passes through a shield wall and into the MICE hall. The infrastructure necessary for running MICE consists of a number of components and their support systems. Each of the superconducting magnets must have its cryocoolers and helium supplies nearby, as does the tracker detector, and the decay solenoid requires a larger helium tank for its refrigerator.

The hydrogen absorbers generate the need for cooling equipment and additional safety considerations. The spectrometer solenoids, which must be moved between Steps IV and VI, need a support mechanism which allows for the relocating of the second spectrometer. Due to the large magnetic fields present in the hall, magnetic shielding is installed to protect the local control room which is located in a neighbouring hall. The control room itself houses the controls and monitoring PCs, the data acquisition systems and storage devices, and some front end electronics. The target control system is also housed in the local control room.

2.3 The MICE cooling channel

The cooling channel consists of RF cavities placed between absorber modules in the beamline. Whilst the RF cavities are intended to re-supply the energy lost in the absorbers, the final absorber module is not followed by an RF module to protect the downstream detectors.

2.3.1 Absorbers

The energy loss required for ionisation cooling is achieved through an ionising medium. This material must be selected so that the energy loss is comparatively high whereas the scattering (which contributes to emittance increase) is kept comparatively low, and so a low Z material is desirable. The relative cooling performance of differing materials is given in Figure 2.6.

As the preference for low Z materials has been stated, the best option for an absorbing material would be hydrogen. MICE will use liquid hydrogen as the primary absorber type. The liquid hydrogen module is then surrounded by a superconducting focus coil to minimise the β function at the interaction point [48].

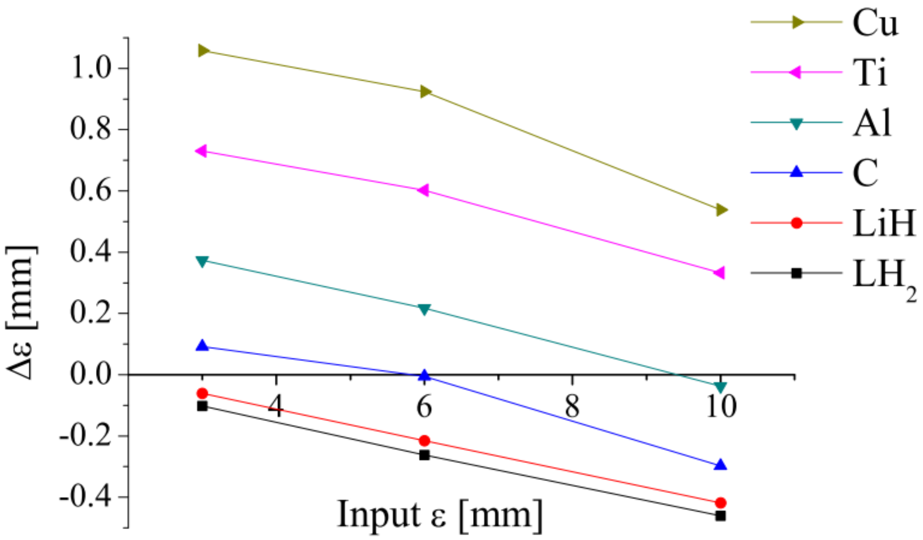

Figure 2.6: Cooling performance for various materials [47]. $\Delta\epsilon$ refers to the change in emittance (mm) between the first and second tracker reference planes. A negative value shows a reduction in the emittance caused by ionisation cooling.

2.3.2 RF Cavities

The re-acceleration part of ionisation cooling will be performed in MICE by two Radio Frequency Coupling Coil (RFCC) modules, with each module being made of four normal conducting RF cavities surrounded by a super-conducting solenoid.

The RF cavities will have a frequency of 201 MHz and a gradient of 8 MV/m. A Neutrino Factory RF system will require an accelerating gradient greater than 16 MV/m; the gradient in MICE is limited by local power constraints. In order to compensate for the lack of power, MICE aims to run with an RF phase of 90° to achieve maximum possible acceleration. However, if this target gradient cannot be achieved the performance of the cooling channel may be compromised [49].

The pulse length will be 1 ms with a 1 Hz repetition rate. The cavity design (a round pillbox design) is based on a prototype built for the MuCool R&D program. The total diameter is 121.7 cm with an aperture and length of 42 cm, where the aperture is covered by a Berylium window with a thickness of 0.38 mm. The total power driving each cavity will be 1 MW. The cavities are moulded in two halves from copper sheets and then joined with e-beam welding. Initial tests on the frequency and tune showed the frequencies of the first constructed cavities to be within 400 kHz of target [50], and during operation mechanical tuners will provide the additional

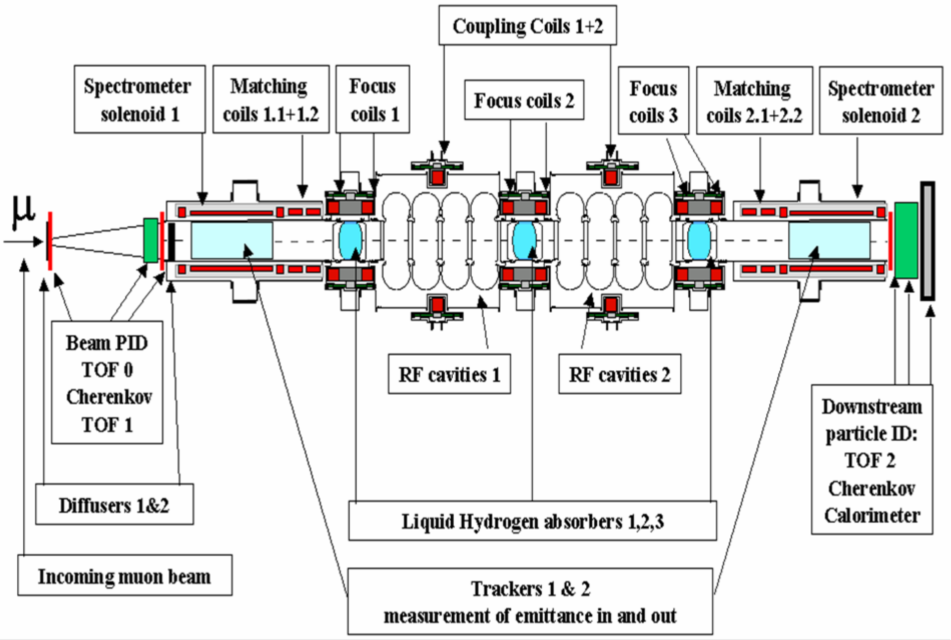

Figure 2.7: MICE Layout: The layout of the MICE experiment including detectors and cooling channel.

modifications.

Test cavities have been constructed and fabrication has begun on the final cavities to be used in MICE. The power system has begun installation in the MICE hall and the required power of 1 MW has been achieved, although with the RF not required until Step V installation will be spread over a longer period than other components.

2.4 Detectors

The MICE detectors are used to measure the particle position and momentum for emittance reconstruction, and to identify the particle type. Particle identification (PID) becomes essential given the possibility of muon decay within the cooling channel. The detectors consist of three hodoscopes to identify particles based on their time of flight; two aerogel Cerenkov detectors to measure the relativistic β versus momentum; a calorimeter to measure the total energy of the particle and to identify electrons; two scintillating fibre trackers to measure the beam emittance before and after the cooling channel for the emittance change measurements. An overall layout of the MICE experiment can be seen in Figure 2.7.

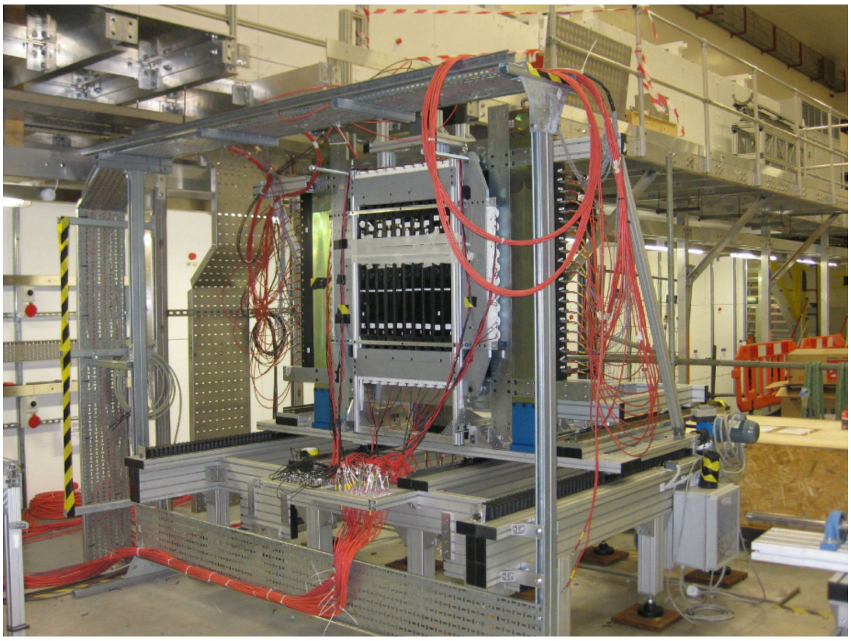

Figure 2.8: One of the time of flight detectors situated in the MICE hall.

2.4.1 Time of Flight detectors

Three hodoscopes, labeled as TOF 0, 1 and 2 in Figure 2.7 (an image of one is shown in Figure 2.8), are installed before the final quadrupole triplet and before and after the cooling channel. The primary purpose of the detectors is to serve as a PID mechanism distinguishing particles by their time of flight (at MICE momenta, muon, pions and electrons are not highly relativistic and have varied times of flight), and to contribute an accurate time component to the emittance measurement. They can also serve as a trigger. Each detector is formed of seven or ten bars of scintillator in both the x and y planes. The scintillator bars are between 4 cm and 6 cm wide.

The scintillator bars are read out at both ends with photo-multiplier tubes; the signal is split between a shaper which elongates the pulse to match the sampling frequency of the flash analogue to digital converter (ADC), and a discriminator whose output pulse is supplied to a time to digital converter (TDC).

The TOFs must be carefully calibrated to take into account the effects of cable length and time walk [51]. Once this calibration procedure (which is a function of the individual bars, as each channel has slightly different relative delays in the cables and electronics) has been completed, the time resolution of the detectors as a whole can be measured. The achieved resolutions are 51 ps, 58 ps and 52 ps for TOFs 0, 1 and 2 respectively, which is consistent with the design requirements. The

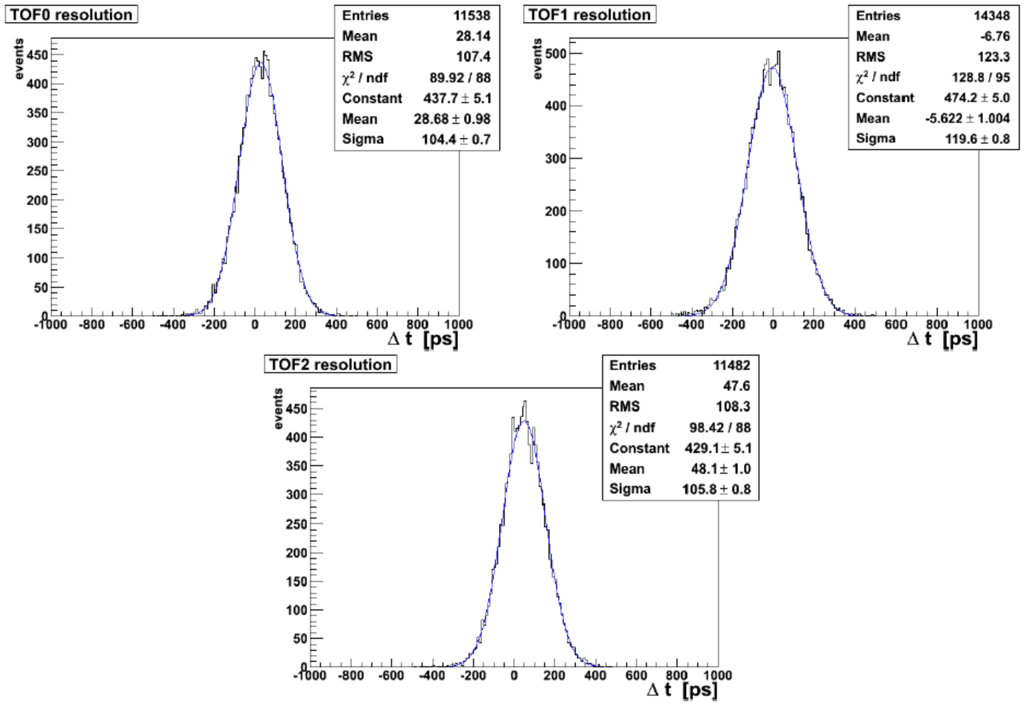

Figure 2.9: Time difference between signals in the x and y planes of each TOF detector.

effect this has on the measured time difference between the planes of the detectors can be seen in Figure 2.9.

2.4.2 Cerenkov detectors

Two aerogel Cerenkov detectors are placed downstream of the second dipole to act as PID detectors. Their refractive indices of 1.07 and 1.12 were chosen so that, at the momentum settings of MICE, the relativistic β of different particles will be distinguishable as above or below threshold. The type of particle can be inferred by knowing the momentum of the particle (by dipole selection) and the β threshold at which the Cerenkov effect occurs in each detector. This corresponds to 220 MeV/c and 280 MeV/c for a muon. The observed light yield with respect to the measured time of flight between TOFs is seen in Figure 2.10. This shows the different light yields produced by different particle types, enabling the distinction between electrons and muons.

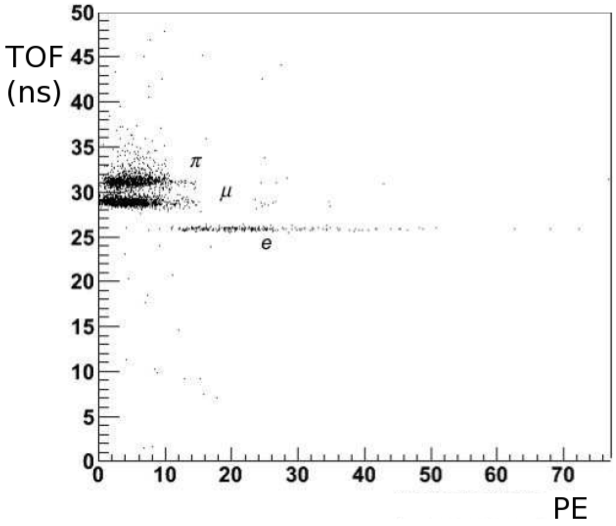

Figure 2.10: The response of the Cerenkov detectors. The light yield in photoelectrons (PE) for the Cerenkov detectors are given with the time of flight between TOF detectors.

2.4.3 Trackers

The MICE scintillating fibre trackers, which are discussed at length in Chapter 3, form the primary detectors for measuring emittance in MICE. Two near-identical, five plane fibre detectors with particle position resolutions of less than 0.5 mm are placed within a 4 T solenoidal field to measure the transverse position and momentum of a particle from a helical radius, and the longitudinal momentum from the dip angle of the helix. Timing can also be obtained from the electronics to provide a resolution of less than 1 ns.

2.4.4 KL

The KLOE-like detector (KL) is a sampling calorimeter based on the KLOE calorimeter [52]. It uses extruded lead foil interspaced with scintillating fibres responsive in the blue region of the electromagnetic spectrum. The fibres are 1 mm in diameter and separated by 1.35 mm, arranged in planes; the planes are separated by 0.98 mm and the fibres in each plane are offset with respect to each other. The total active volume is $93 \times 93 \times 4$ cm^3, which is broken into 21 cells with two readout channels per cell. The output signal is shaped before being digitised by 14-bit flash-ADCs. The response of the KL to two particle species at varying momenta is shown in Figure 2.11.

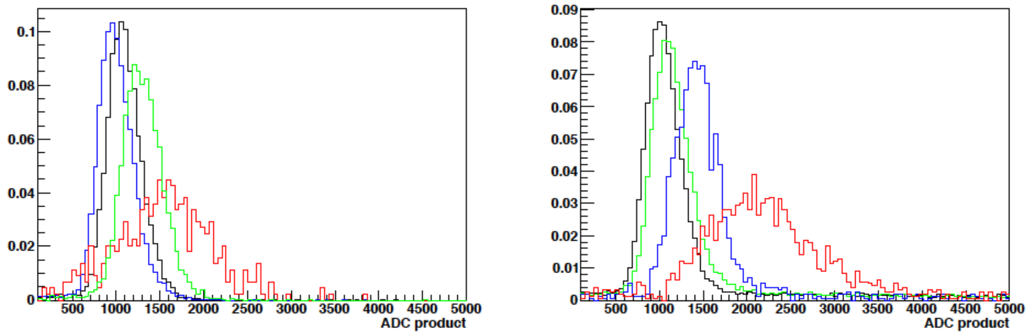

Figure 2.11: The response of the KL detector to muons (left) and pions (right) for varying momenta. The red, green, black and blue lines refer to momenta of 80, 150, 220 and 320 MeV/c respectively for muons, and 135, 165, 240 and 320 MeV/c respectively for pions.

2.4.5 Electron-Muon Ranger

In order to identify muons which have decayed in the cooling channel, an electron-muon ranger (EMR) has been constructed [53], which will distinguish electrons and muons based on their stopping pattern through a volume of scintillator. The EMR is formed of triangular scintillator bars readout by multi-anode PMTs. The bars are arranged in one direction in one plane and in the orthogonal direction in the next plane, forming an $x - y$ geometry. The EMR can be seen in Figure 2.12.

2.5 Software

MICE uses a bespoke tracking, digitisation, reconstruction and analysis software framework based on the Geant4 libraries [54][55] named G4MICE [56]. At the time of writing, a more complete analysis framework and wrapper written in Python was being implemented - the MICE Analysis User Software (MAUS) [57].

2.5.1 Particle Tracking

G4MICE uses a single particle tracking method, which enables the accurate modelling of the momentum dependance of the magnetic field effects, as well as the stochastic processes from multiple Coulomb scattering and energy loss.

In the absence of measured field maps for the MICE beamline magnets[1], the

---

[1] A plan is in place to measure the magnetic fields of the cooling channel magnets

Figure 2.12: The EMR detector placed at the end of the MICE beamline.

coil geometry and current densities are used to generate a field map at run-time. The field map is then used as a look-up table during the Geant managed tracking.

2.5.2 Detector Simulation

In the geometry descriptions, the active areas of each of the detectors (TOF scintillator bars, tracker fibre planes etc.) are defined as Geant sensitive volumes, and so truth information about any particle will be stored at the point of the detector. Each detector has a slightly different MC truth object corresponding to it, depending on the differences between them. For example, the TOF MC truth object (a TOFHit) will record the plane $(x-y)$ and slab number, whereas a tracker MC object (a SciFiHit) will record station, plane and fibre information, whilst all inheriting from a common MCHit object.

In simulation, the digitisation phase takes these recorded truth values and converts them into an object representative of the detector itself. For example, this will involve taking the stored true energy deposition and convert it into a light yield and then into an ADC value. The resulting digit is identical in structure to its real data counterpart.

2.5.3 Real data and calibration

The real data is obtained by front end electronics and stored in different ways for each detector. In the TOFs, for example, the analogue and time values are obtained by separate VME boards, whilst for the tracker these functions are combined. Therefore, it is the task of unpacking software to combine the separate pieces of information associated with a single particle hit.

The stored values, which may be digitised charge values, themselves somehow dependent on observed scintillation light, must then be converted into a physically meaningful quantity such as light yield or energy deposited. The conversion factors can be dependent on numerous features of the detectors and must be individually found. In software, these calibration values are stored and used to convert the real digitised data into the digit object.

2.5.4 Reconstruction

The reconstruction of particle tracks is performed in a series of steps. The digits are first converted from a purely electronics-based object (albeit calibrated from quantities such as ADC into light yields) into a physical representation. The channel mapping is taken into account and the resulting object will represent a physical object, such as a TOF slab. Depending on the detector, the planes are iterated over to combine single plane objects into space points, whose resolution depends on the overlapping region of the single planes.

Once space points have been formed, a tracking algorithm is used to reconstruct the particle trajectory between the points. For the TOFs, this means finding a track between detectors, and for the tracker between planes in a single detector.

The details of the reconstruction methods differ between detectors. The TOF has used an iterative method whilst the tracker uses a full Kalman filter.

2.6 Controls and Data Acquisition

The primary systems used in MICE are the Data Acquisition and Test Environment (DATE) [58] used by ALICE at CERN, and the Experimental Physics and Industrial Control System (EPICS) [59].

2.6.1 Controls and Monitoring

The slow controls of MICE must be responsible for monitoring and controlling the full apparatus of the beamline, as well as a selection of different particle physics detectors. Currents for the beamline magnets, high voltage supplies to photomultipliers and other parameters must be set. This is achieved by various individual systems, the foremost of which is EPICS, which is used for both slow control and logging of monitored parameters. EPICS is a client-server architecture which allows for the monitoring and archiving of set values, and setting via electronics of control variables over networked devices.

2.6.2 DAQ

The data acquisition system of MICE will be required to buffer approximately 600 triggers worth of digitised data, having taken less than 500 ns per digitisation. The buffered data is then read out at the end of the spill, which must be completed in less than 1 s. VME-buses are used as the interface between the front end electronics and the data acquisition PCs. Each VME crate is connected to a single PC via an optical link, and this PC acts as the Local Data Concentrator (LDC) for that detector or part-detector.

DATE then combines the total recorded spill information from each LDC over a Local Area Network in an event building process. The PC responsible for this is labelled the Global Data Concentrator and produces the final output data file. Additional information on the functioning of DATE and how it pertains to the trackers is given in Chapter 3.

2.7 Step I Beam Parameters

Luminosity and Particle Rates

The number of muons available in the MICE beamline is directly affected by the number of particles produced on the target, which is itself a function of the cross section of the target presented to the beam. Ideally, the target would remain in the beam as long as was technically viable. However, given the parasitic nature of the target running there are other concerns.

Losses to the beam in the ISIS synchrotron vault are of particular concern, since particles knocked out of the beam pipe can irradiate and activate accelerator components in the vault, causing damage and creating a safety concern. To mitigate this, the losses induced by the MICE target must be kept within defined limits, where 1 Vms (as defined by the ISIS beam loss monitors) is a typical value.

The dependence of the rates seen in MICE on this limit is therefore of interest. Studies were performed on special high beam loss runs, which showed a clear linear relationship between the losses observed by the ISIS monitors and the reconstructed tracks seen by the TOFs in MICE [60]. This is shown in Figure 2.13.

The design of the target has also been shown to affect both the particles produced and the losses incurred in the ISIS vault, and has been optimised over iterations to maximise the ratio of particle tracks in MICE to the losses incurred in ISIS.

Particle time of flight difference

Using the TOF's level of timing accuracy, the separation of particle times of flight (due to their differing mass but constant momentum as selected by the upstream dipoles) can be used as a method of particle identification. The differing times of flight for electrons, muons and pions based on forward and backward pion decay optics, can be seen in Figure 2.14.

The Time of Flight detectors can therefore be used to give each particle's type based on the time of flight between the detectors. This will enable the tagging of muons for use in ionisation cooling studies.

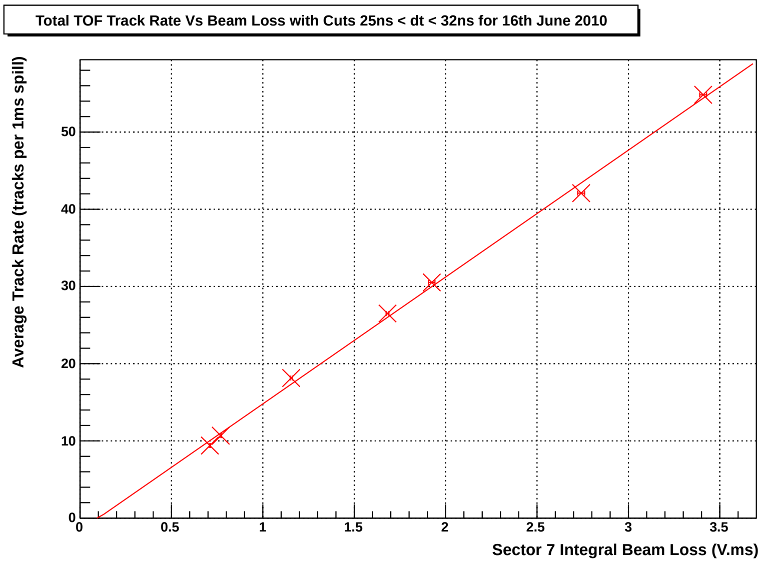

Figure 2.13: Reconstructed time of flight tracks compared to observed ISIS beam loss.

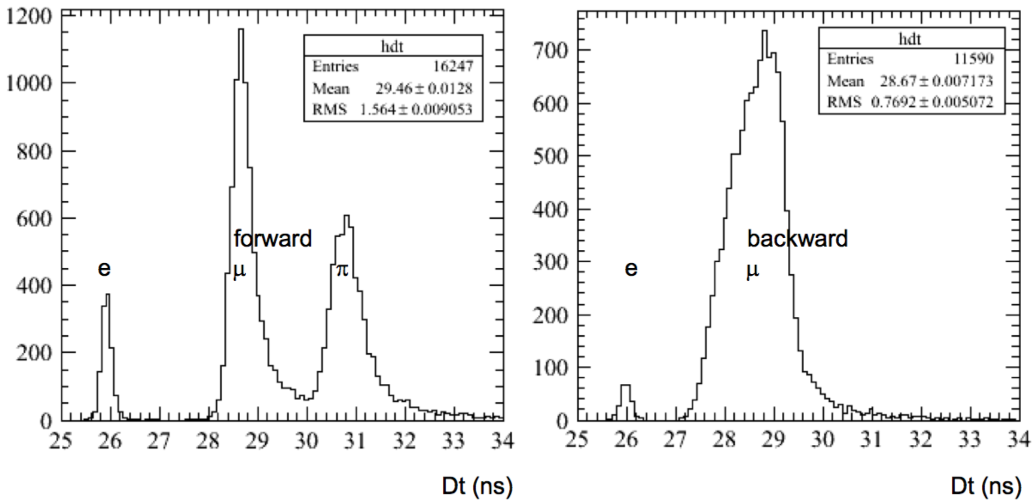

Figure 2.14: TOF detector particle identification: The time of flight between TOF0 and TOF1, where the dipole current has been set to select forward (left) and backward (right) decaying muons.

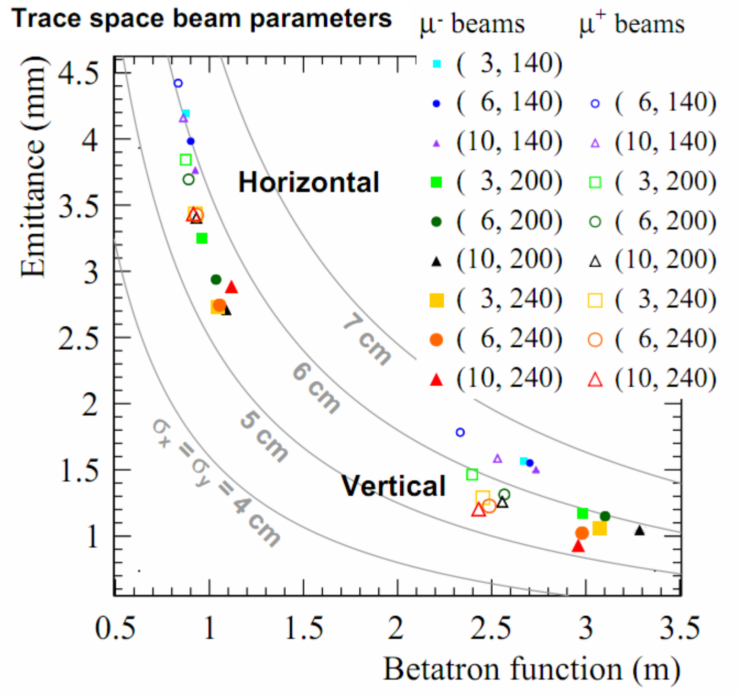

Figure 2.15: Reconstructed emittance and β function for beamline settings corresponding to the MICE operating modes [62].

Reconstructed phase space

Based on the position and time information obtained from the time of flight detectors, the trajectory of the particle between two detectors can be obtained. In MICE, TOF0 and TOF1 are both before the beginning of the cooling channel, and so each particle's trajectory could be measured before the diffuser plate. The intervening quadrupole triplet must be taken into account. The results of using an iterative, single particle transfer matrix method [61] for each of the MICE operating modes is given in Figure 2.15.

The difference between the measured emittance and β values are due to the absence of the diffuser. This shows that the beamline has been commissioned and work can continue towards installation of the cooling channel.

2.8 Outlook

The MICE beamline has been completed and characterised using time of flight detectors. A mechanical target installed in the ISIS proton synchrotron has provided a stable particle rate whilst minimising losses to the host accelerator, and has run for many millions of particle triggers. All of the detectors used in the beamline have been shown to respond according to their design specifications, including a less than 60 ps time resolution on each of the time of flight detectors. In addition to the high time resolution, the TOF detectors have been used to distinguish particle types based on their difference in time of flight at a given momentum (selected by upstream dipole magnets).

Upon completion of the spectrometer solenoids, the trackers will be installed into the MICE hall along with the diffuser plate. This will allow a greatly improved measurement of the beam emittance, including the emittance inflation provided by the diffuser. Production of the absorber focus coils and installation of the hydrogen systems will then allow for data taking of Step IV. Finally the introduction of the MICE RF system will finalise the cooling channel; at this point a precision measurement of ionisation cooling using particle physics detectors will be possible. The complete MICE will then be able to verify the proposal of using ionisation cooling for use in a future muon accelerator.

Chapter 3

Calibration and Performance of the MICE Scintillating Fibre Trackers

The primary measurement of emittance in MICE will be performed with two spectrometers, each consisting of a scintillating fibre tracker [63] placed within the bore of a superconducting solenoid. The magnetic field parallel to the particle motion will force the particle into a helix where the radius is proportional to the particle's transverse momentum. The helical motion of the particle through the magnetic field coupled with the fine granularity of the tracker will enable a precision measurement of both position and momentum, from which the emittance of the beam can be calculated.

Scintillation light induced in the fibres by a charged particle is transferred via waveguides onto photon detection devices which convert the light into a charge. This charge is digitised and the value, along with which channel it belongs to, enable the reconstruction of which channels the particle passed through. The crossing geometry of the channels allow for space points to be reconstructed in each station, and combinations of points along the tracker can be used to fit tracks, and finally, a position-momentum state for the particle to be used in a calculation of emittance.

3.1 Spectrometer Solenoid

The MICE trackers will be placed within a superconducting solenoid, the field of which will cause the particles to move in a helix allowing their momentum to be reconstructed. The design specifications require that the solenoid have an overall field uniformity of better than 1%, and around 0.3% in the detector region. The solenoids are made of two matching coils which can be tuned to match the solenoid field with that of the absorber focus coils, two end coils and a main centre coil. The parameters for each of these sections are given in Table 3.1 [64].

Coil	Turns/layer	Layers	Radius (mm)	Thickness (mm)	Length (mm)	Peak Field (T)
M1	115	42	258.0	46.5	201.2	4.43
M2	114	28	258.0	30.6	199.5	4.01
E1	64	56	258.0	61.7	110.6	5.90
C	768	20	258.0	22.4	1314.3	4.19
E2	64	62	258.0	68.3	110.6	6.37

Table 3.1: Summary of spectrometer solenoid coil parameters.

The coils are made with a Nb-Ti superconductor in a 1:4 ratio with copper, placed within a cold mass filled with liquid helium kept at 4 K by at least three cryocoolers. The solenoid and coil arrangements can be seen in Figure 3.1.

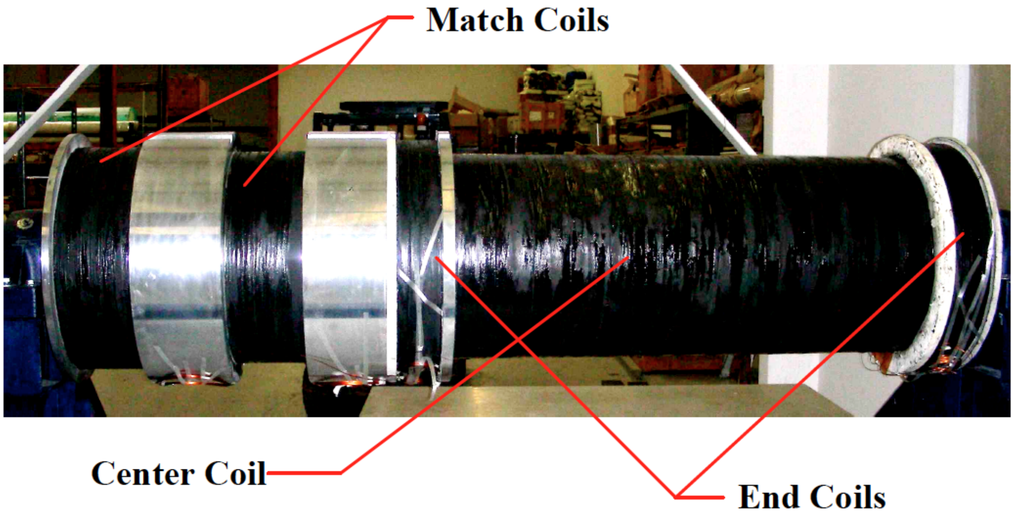

Figure 3.1: Photo of the spectrometer solenoid without the cold mass layer.

3.2 Tracker Body

The warm bore of the solenoid has a diameter of 40 cm, into which will be placed the tracker detector. The stations of the tracker are mounted on a carbon-fibre space frame enabling it to be slid into the centre of the solenoid.

3.2.1 Tracker Stations

Each tracker consists of five stations placed at non-equal intervals. The spacing between tracker stations was chosen to minimise the possible degeneracy in the position measurements of the particles as they precess in a helical motion. The layout of the tracker stations (as they are held in place by a carbon fibre frame) is shown in Figure 3.2, and the distances between stations defined in Table 3.2. The stations themselves are made up of three layers or planes of scintillating fibres.

The scintillating fibres are ganged together into a clear fibre which runs up the inside of the tracker body and is then bundled into waveguides and connected to the photon detection equipment. The attenuation length of the waveguides is approximately 7.8 m, whereas the length of the waveguides is approximately 2 m. The waveguides can be seen in Figure 3.3. Each waveguide connects to the top of the cassette containing the photon detection devices and supplies fibres to a single module. As the number of electronics channels exceeds the number of physical fibre bundles, not every module is used.

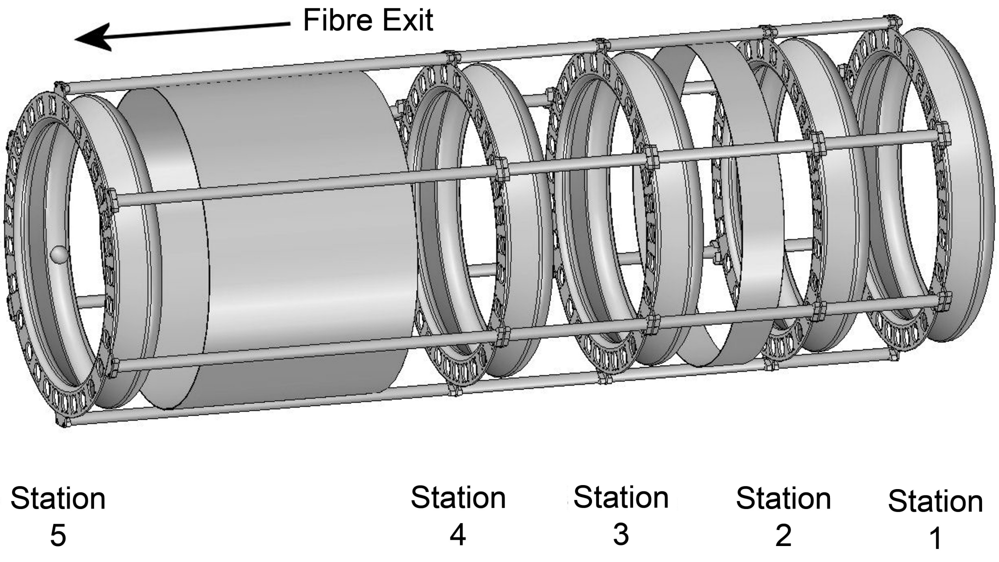

Figure 3.2: Arrangement of the stations within the tracker body, supported by the carbon fibre space frame.

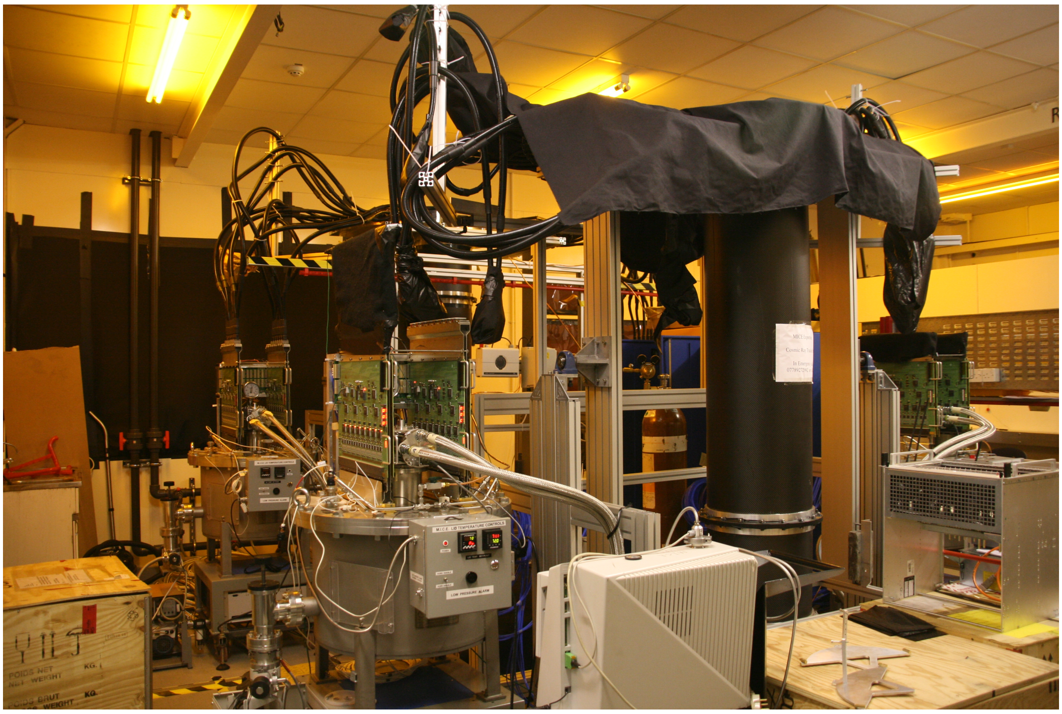

Figure 3.3: Tracker body, waveguides, cryostats and electronics.

Item	Value
Fibre Diameter	350 μm
Station 1-2	20 cm
Station 2-3	25 cm
Station 3-4	30 cm
Station 4-5	35 cm

Table 3.2: Summary of tracker parameters

3.3 Fibres and Photon conversion

The use of scintillating fibres with solid-state photomultipliers is a common detector technology, and the MICE tracker system is heavily influenced by the DØ central fibre tracker [65].

3.3.1 Scintillating Fibres

The fibres are 350 μm thick and arranged in a double layer of seven to form a single readout channel, as shown in Figure 3.4. The arrangement of the fibres produces a pitch (the distance between two fibre centres in the same layer) of 427 μm. The planes themselves are placed at 120° to one another, as shown in Figure 3.5. The overlapping nature of the fibre ribbons ensures that there is no dead space in the tracker active region, which has a diameter of 30 cm and is covered by 212-214 fibre bundles in a single plane.

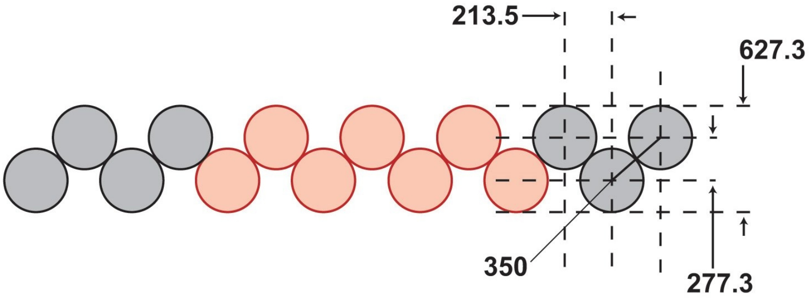

Figure 3.4: Grouping of fibres into channels. The measurements are given in μm.

The fibres contain two dopants, the first to produce the scintillation light and the second to absorb this light and re-emit at a frequency better suited to conversion to an electrical signal. The polystyrene used as the scintillator material is naturally active to the passage of charge particles. The first dopant is present to the level of 1%, which leads to a separation of approximately 1 nm. At this level, the primary dopant is electrically bonded to the polystyrene, creating a very high quantum efficiency emission from the primary dopant.

The primary scintillating dopant is Polythiophene (pT) at 1.25% by weight and the secondary, wavelength shifting dopant 3-Hydroxyflavone (3HF) at 0.25% by weight. As the fibres are only transferred into the waveguides and readout at one end, a mirror (formed by aluminium vapour deposition) is placed at the opposite end. The reflectivity was measured to have a mean value of 0.75. A 25 μm layer of Mylar was placed between each fibre plane and fixed in place with a polyurethane adhesive. The three fibre planes, plus the Mylar and epoxy, contribute a total of 0.0045 radiation lengths.

The fibres underwent a quality assurance program to optimise the dopant levels and check for flaws. A ^{57}Co source was used to fully penetrate the fibre planes and a mechanical rig set up to test the fibres. Due to the low data rate encountered the full active region was not scanned, but rather a path chosen to scan each fibre at least once.

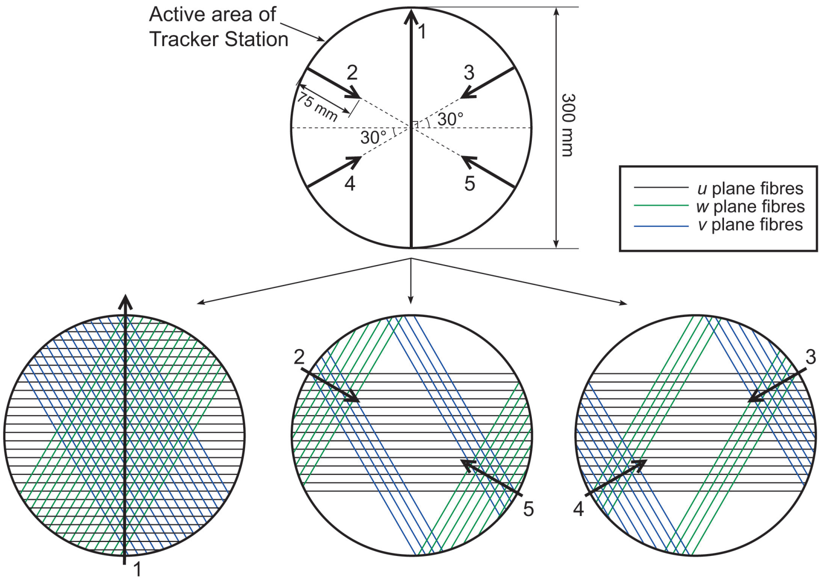

Figure 3.5: Arrangement of fibre planes into a station.

Doublet Geometry

The overlapping nature of the individual fibres combined into a single readout channel, as described earlier, produces a doublet geometry which affects the particle reconstruction. Assuming an isotropic distribution of particles passing through the tracker with no angular dependance, a finite number of particles will pass through fibres from different readout channels, generating two signal hits. The proportion of two channel hits increases as an angular dependance is introduced into the particle distribution.

3.3.2 VLPCs

The photon conversion in the MICE trackers is performed with Visible Light Photon Counters (VLPCs), which are impurity band conduction silicon diodes [66]. They are high gain, high quantum efficiency devices which must be operated at cryogenic temperatures. An incident photon activates an electron and an avalanche is induced, generating a large enough charge to be measured in electronics. The gains achievable by the MICE VLPCs in this mode are 20,000 - 60,000. They differ from solid state photomultipliers only in their response in the infrared region.

VLPCs operate with a large quantum efficiency (QE), typically around 80%, and have been measured as high as 88.2% [67]. Such a high QE is a requirement of single photon detection systems.

VLPCs of similar performance are grouped together in cassettes, and each cassette contains 1024 channels broken into eight 128 channel modules. The modules themselves are made of sixteen VLPC chips. An example of a VLPC chip, where the 1 mm radius of each VLPC can be seen, is shown in Figure 3.6.

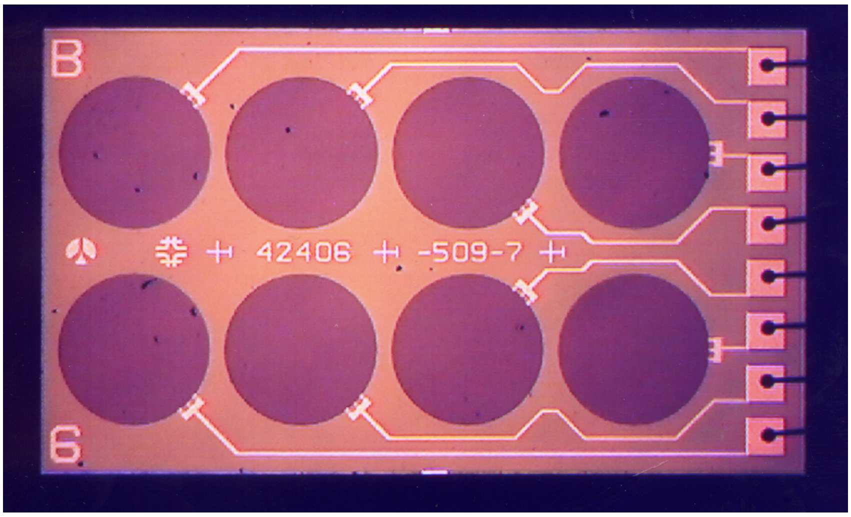

Figure 3.6: A VLPC Chip with 8 channels.

As VLPCs are silicon diodes, a bias must be applied to them. The voltage set across the junction increases the electric field, allowing an avalanche effect. This bias must be chosen to maximise the avalanche (and thus the observed charge per photon) without pushing the diode into breakdown or producing signal regularly on background radiation. The behaviour of the multiplication factor and noise have a clear relationship to the applied bias [68].

The VLPC cassette is a single physical body which is inserted into the cold mass of a cryostat. It is broken into eight isothermal modules with their own temperature monitoring and heaters. The VLPCs in each of these modules (128 channels) are biased and read out in two halves by two different readout electronics modules.

In this discussion, dark count refers to the presence of photon noise within the VLPC cassette when it is light-tight (and thus dark) from thermal activity. Infrared photons from thermal radiation, even with a wavelength far from the optimal operating range of the VLPC, can contribute a noticeable rate. Even at the low operating temperature of the VLPCs, the accidental rate from this process can contribute significantly to the measured charge, depending on the amount of bias applied (the contribution from dark count rises exponentially with the applied bias).

3.4 Electronics

The tracker electronics can be divided into two sections: the front end electronics which interact with the VLPCs, setting parameters and retrieving the data, and the VME electronics which communicate with the readout PCs and data acquisition software. The front end electronics send packaged data for each channel per event via Low Voltage Differential Signal (LVDS) cables to the VME LVDS SerDes Buffers (VLSBs) which are essentially large memory buffers in the VME crate. The DAQ then accesses the VLSBs and writes the data contained in their memory banks to the DAQ stream and eventually to file.

3.4.1 Analogue Front End

The readout electronics for the MICE trackers consists of Analogue Front End (AFE) boards, specifically of the Type-II used by the DØ central fibre tracker [69]. The AFE-IIt boards handle all the necessary controls for the VLPCs as well as the data acquisition. The functions can be summarised as follows:

- **Control** - The AFE board controls the voltages set to the VLPC and to the heaters controlling the VLPC cassette temperature.
- **Trigger, Pipeline and Time** - The pipelining, triggering and time functions are handled by a dedicated chip of the AFE board.
- **ADC** - The conversion from the integrated charge seen by the VLPC to a digitised value is performed by a 10-bit ADC. The ADC also digitises the analogue time pulse produced by the timing circuit.
- **Data packaging** - Field Programmable Gate Arrays (FPGAs) retrieve the digitised values for each channel and package the data to be sent to the VLSBs.

VLPC Bias and Cryo Settings

The VLPCs require a bias voltage to be applied in order for the diode to operate. In the case of the AFE-IIt boards and VLPC layout, this voltage is applied once per group of 64 channels, or module, giving a total of 8 bias voltage settings per AFE board, and 16 per VLPC cassette.

The selection of the correct bias voltage is important to ensure the correct ratio of signal to noise in the VLPCs. If the bias is set too low no signal will be seen; if it is set too high the noise will dominate the lower end of the signal spectrum. The optimisation employed for the VLPCs in DØ was the requirement that the dark count should represent 1% of the pedestal distribution, given a threshold of one photo-electron. Further consideration of the correct bias settings and a method for determining them for use in MICE is discussed later in Section 3.6.3.

The bias set points require calibration both for the VLPC module the bias is applied to, and the AFE electronics applying the voltage. The internal Digital to Analogue Converter (DAC) value is a function of both the cassette and board calibrated values. The board values (which are unique to an individual AFE board) were measured during construction and are taken as is. The relationship between these calibrated values, the set bias voltage and the DAC value is therefore:

$$\text{DAC} = (\text{Slope} \times \text{Bias}) + \text{Offset} \tag{3.1}$$

The requirement for VLPC operation is that they are kept at 9 K. Even with a well set bias, slight variations in temperature can greatly affect the performance of the diodes and generate noise or reduce the observed signal. In order to mitigate this, along with the cryostat, heaters operating in a feedback loop are used with one heater per two modules, controlled by one of the AFE boards attached to the cassette. This enables the temperature to be kept at a level at which the performance of the VLPCs is not affected.

Trigger, Pipeline and Timing

The AFE-IIt electronics, originally purposed for DØ detector, have been modified for use in MICE whilst utilising the built in functionality. Most of the changes have been in firmware for the FPGAs managing data acquisition and the control FPGAs on the board.

The management of the pipeline, the access to a particular bucket in this pipeline as well as obtaining discriminator and timing information is performed by a Trigger, Pipeline and Timing chip (TriP-t), which is used in many high energy physics experiments aside from DØ, such as T2K and Minerva. Each TriP-t has 32 channels corresponding in MICE to 32 VLPCs.

The integration window is defined as a period of time during which the charge produced by the VLPC is considered signal, is integrated, and if coinciding with a trigger, digitised. Given the repetition of collider experiments, this window is periodic, with an *on* period - the window during which the produced charge is integrated, and an *off* period during which the integrator is idle and any signal will be ignored. The size of the integration window is variable; in MICE, this is set to approximately 100 ns on and 200 ns off. The off period is limited to a minimum by the time taken to integrate the charge and move to the next pipeline bucket, and as such, is the trigger dead time. The electronics dead time is dominated by the time taken to digitise each channel in the TriP-t.

The width of the window will have an effect on the observed signal in that only photons arriving within the on period will have their charge integrated. A pulse with any proportion of the charge distribution close to the edges of the integration window (the edges themselves are not firm) will not be fully measured. It is therefore desirable that the window repetition be timed so that any pulses arrive with their maxima at the centre of the integration window, although this may not always be possible given the variations in the time of flight of the particles in MICE.

In addition to the observed charge, the timing of the on and off periods need to be correctly set. Whilst not a collider experiment, MICE is a pulsed experiment with a macro scale corresponding to the target dips or spills, and a micro scale corresponding to the structure of the beam caused by the ISIS RF. Whilst the macro scale is of $\mathcal{O}$(ms) and will encompass many integration windows, it becomes necessary to time the micro scale by ensuring the integration window lines up with the arrival of the (ISIS RF induced) bunches.

This is achieved by inputting a second clock signal into the AFE boards, which is identical to the ISIS RF. The majority of the board, including FPGAs, run off the 53 MHz clock locked to that from the VLSBs. The TriP-ts, and thus the integration windows, run off this second external clock. Therefore, the integration window is synchronised with the ISIS RF period. In the initial tests, this was performed with a constant frequency signal similar to that of the ISIS RF frequency. However, since ISIS is an accelerating machine, this frequency is not constant. It increases until the beam is extracted from the ISIS ring, and despite the abstract existence of a flat top in the proton momentum, the frequency is not constant long enough given the length of the MICE spill gate. Therefore, the particles arrive in the trackers asynchronously with respect to the AFE's internal clocks, and so an externally defined clock derived from the ISIS RF must be used. In addition, the offset between the ISIS RF and the opening of the integration window, generated by the time of flight of particles from the target to the trackers, must be measured and taken into account.

The discrete nature of the integration windows allows for the signal to be pipelined. Given a trigger latency longer than the latency in the photon conversion, it becomes necessary to store the information in each integration window in a pipeline, and then select which bucket in the pipeline to digitise once the trigger has arrived. The pipeline is achieved by a set of 48 capacitors per channel in the TriP-t, and is done for both the analogue (signal charge measured as the ADC count) and time (TDC) pulses. The capacitor whose integrated charge is digitised upon arrival of a trigger, i.e. the integration window expected to contain the signal, is referred to as the pipeline depth and is accessed by defining the first 6 bits in the Pipeline and Gain register in the TriP-t. This depth depends primarily on the trigger latency and must be measured.

The relationship between the integration windows, the signal and the trigger as distributed to the TriP-ts can be seen in Figure 3.7.

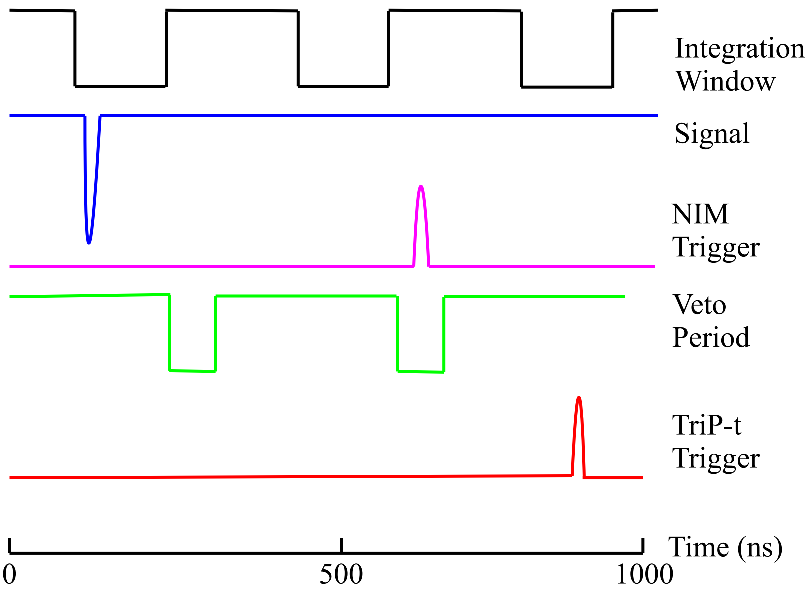

Figure 3.7: Trigger Timing: the integration window is shown in black, the veto period in green, the signal, either from an LED or a particle in blue, the trigger from the logic in purple and the trigger given to the TriP-t in red.

Each channel has a discriminator which is able to discern whether the charge seen within a particular integration window has risen above a certain value, which is set as one value per TriP-t. There are two inputs, one the threshold voltage and the other the signal voltage. If the signal voltage is below the threshold voltage, the output is set to 0. If at some point the signal voltage rises above the threshold voltage, the output switches to 1 and is kept there until the system is reset for the next integration window. The value of this threshold is set as the difference of two programmable registers in the TriP-t, labelled as VReference (VRef) and VThreshold (VTh), each with a default value of 170 decimal counts.

Within the integration window, when the charge rises above a certain level, the comparator output will change and the discriminator will have fired. The range of the discriminator is between 10 fC and 300 fC.

The original testing of the TriP-t chip for use in the MCM-IIt electronics in the DØ fibre detector showed the probability of the discriminator firing as a function of the VTh ADC count rose from 0% at 210 counts to 95% at 216 counts [70]. This would suggest the region in which the VThreshold needs to be calibrated is quite small and requires careful consideration of the effects of even a minor change in the set VTh value.

The timing of the discriminator firing with respect to the time at which the charge rose above threshold is also an issue, both in terms of a direct delay and time walk. Time walk is defined as the difference in times at which two rising charge pulses exceed some threshold, yet originate at the same time. Thus, two signals with a common start will have their discriminators fire at different times. This is not an issue if the only concern is whether or not the discriminator fired during an integration window, as for zero suppression, but it is of concern if the discriminator starts or stops some time dependent function, as in the case of the Time to Digital Converter (TDC) used in the AFE-IIt. This is illustrated in Figure 3.8.

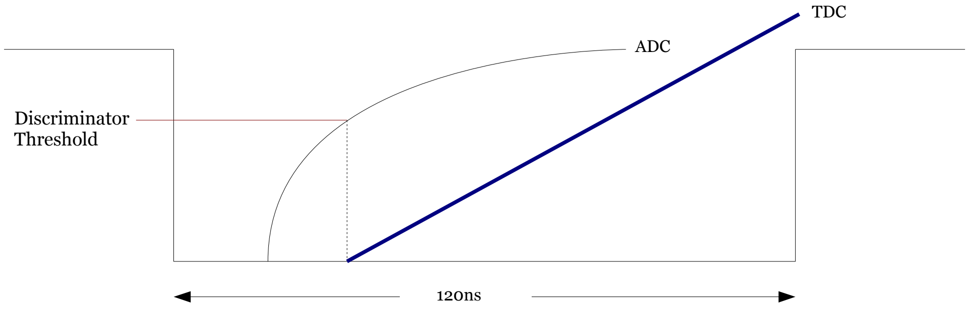

Figure 3.8: Diagram of the ADC, TDC and Discriminator within an integration window.

With the knowledge of whether or not a particular channel saw charge above a definable level during the integration period, a cut on channels failing to see this minimum charge can be made, and the additional information from this channel (e.g. channel number, number of ADC counts) not passed on to the next level of data acquisition. This process is the zero-suppression, as it allows for the suppression (or ignoring of) those channels which did not see signal. Given the absolute minimum number of channel hits for a good space point in each station, the ratio of signal channels to all possible channels is 15:4096. It is clear that not reading out the remaining (signal-less) 4081 channels is of benefit. This is especially true when the data size and speed of data acquisition is of concern.

Timing information is obtained in the TriP-t by use of a Time to Analogue Converter (TAC) which is then in turn digitised by an ADC, the combination of which are referred to here as a TDC. Within a single channel of the TriP-t, when the discriminator fires a constant current is initiated until the end of the integration window, at which point the stored charge is digitised. This allows, given calibration, to measure the time at which the discriminator fired. The value of this current is set as part of the TriP-t initialisation and will affect the observed number of TDC counts for a given time at which the discriminator fired. Given a measured and useful value is now dependent on the manner in which the discriminator fires as a function of time, both time walk and an absolute offset are an issue. The initial testing of the TriP-ts mentioned earlier calculated the time walk as a function of photons received before the discriminator fired, ran from 9.2 ns with 3 photons to 0.8 ns with 16 photons.

The TriP-t also has internal amplification circuits which must be properly set in order to utilise the maximum gain for the signal. This is shown in Figure 3.9, along with the discriminator and timing circuits. The analogue and timing output pulses are pipelined prior to digitisation, whereas the discriminator output is sent directly to the DFPGA.

ClockGen and Charge Injection

The AFE runs off a 53 MHz clock which is internally generated and locked to an external input which is cabled from one of the VLSBs. The internal AFE clock and the externally inputted signal are synchronised. This is different to the externally inputted ISIS RF clock when the AFE board runs in external clock mode. The mode of the ClockGen is defined during initialisation.

Additionally, registers within the ClockGen control the charge injection which can be used for calibration. A charge, set in four levels, can be injected into a specific channel, and is controlled by a charge injection map stored in the TriP-t. By this method, charge can be injected into a single channel or all channels at will. The

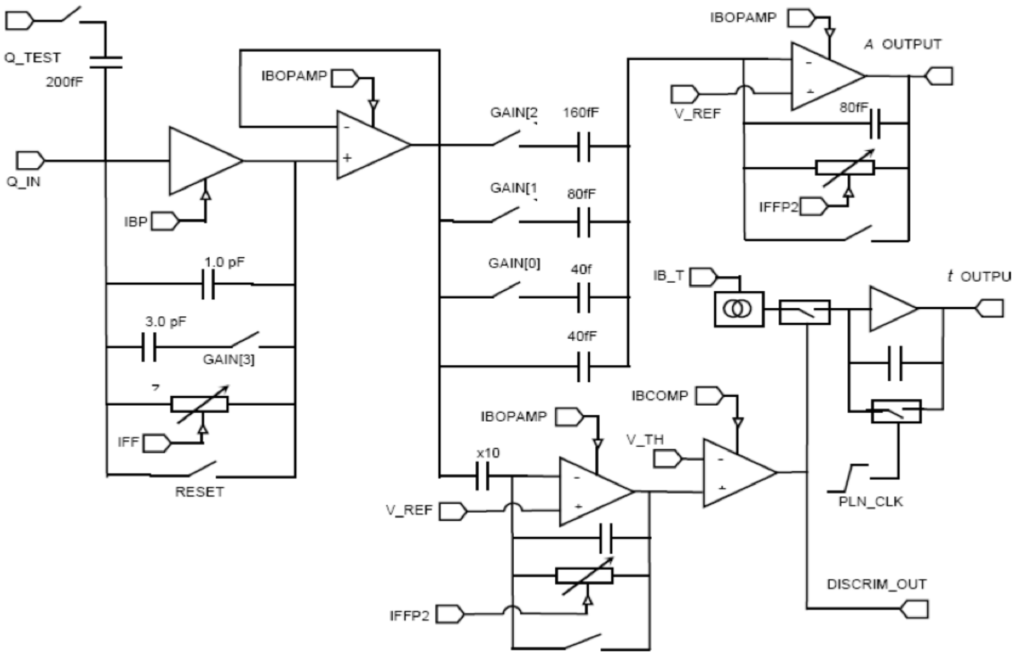

Figure 3.9: TriP-t circuit diagram.

time at which the charge is injected into an integration window can also be defined in internal clock ticks from the beginning of the window.

Analogue and Digital FPGAs

The analogue FPGA handles the packaging of each channel's ADC, TDC and additional data into a single word. The time at which charge is injected into the integration window is also controlled by the AFPGA, and in this case, the integration window must be running on the internal clock to ensure synchronisation. The DPFGA controls the grouping of all channels for readout, as well as the pipelining of the discriminator bit.

The information for each channel at each sample or trigger is packaged by the FPGAs into a single 32-bit word. This takes the format described in Table 3.3.

Initialisation and EPICS

Initialisation, control and monitoring of the front end electronics is performed by the Experimental Physics and Industrial Control System (EPICS). This operates a client-server system where parameters are passed to a server as process variables

Data	Number of bits	Bit locations
ADC	8	0-7
TDC	8	8-15
Channel	7	16-22
Discriminator	1	23
Event Number	8	24-31

Table 3.3: AFE-IIt data format.

and then written by a chosen communication method to the experiment electronics. In the case of the MICE trackers, the communication is done via a VME Mil-1553 interface. By this method, EPICS programmes the AFE FPGAs (using the built in flash memory) and sets desired bias voltages, pipeline delays and other parameters.

The Mil-1553 [71] is a serial bus used by DØ for its extremely low noise. In MICE, a Mil-1553 VME controller communicates with remote terminals on the AFE boards. A back plane connects a single 1553 controller to four AFE boards differentiated by their remote terminal address. Communication takes the form of a 16-bit command/data/status word coupled with 3 synchronisation bits (which distinguish commands from data) and a parity bit. The command word will contain the remote terminal address. In this mode, all AFE boards on a single cryostat receive all communication, but only the board on the correct remote terminal respond to the data.

3.4.2 VLSB

The primary role of the VLSBs is to store the data read out by the AFE boards until the data acquisition writes them to file. In addition, a Master VLSB controls the triggers sent to the AFE in the form of an encoded clock. Upon receipt of a formatted trigger pulse the clock which is locked between the VLSB master at the AFE boards will go through a two-high-two-low cycle; this is the signal for the AFE boards to digitise and read-out.

Each individual VLSB contains four independent memory banks of 128K words (128000 × 32 bits) which are written to by an LVDS interface on the front panel [72]. They are only readable by VME access in the MICE setup. From the LVDS cabling, each VLSB memory bank directly corresponds to one DFPGA on an AFE board, compromising 128 readout channels.

An additional function involves checking the discriminator bit for each word and only storing those with a high bit. Zero suppression is therefore actually implemented in the memory buffers, and not the front end electronics.

3.5 Data Acquisition and Unpacking

After each channel has been digitised and packed by the DFPGA, the data words are sent by LVDS to the VLSB boards, where they are buffered, optionally zero suppressed and stored in memory. The contents of the VLSB memories are accessible via VME. At each spill the data acquisition software must access each VLSB and retrieve the contents of its memory, subsequently writing it to file. The written format is nothing more than the bitstream sent from the AFE boards, and therefore needs to be unpacked for analysis and reconstruction.

Each of the two trackers has a dedicated VME crate to house the VLSBs, which are connected by a CAEN interface board to a PC. Whilst each tracker is readout by its own PC, both trackers can be controlled and monitored simultaneously by a single PC.

3.5.1 DATE

MICE uses the Data Acquisition and Test Environment (DATE) developed for the ALICE experiment. DATE enables many detectors read out by different front ends through many VME crates to have their data combined into a single format. It achieves this by using an object oriented-like framework to create objects representing each piece of real or virtual hardware (virtual hardware is defined as some required VME process which does not directly correspond to a physical board, such as a spill timer). Each VME process is defined as a piece of DATE equipment and must have up to five functions implemented. These are defined as follows:

- **Arm** - The arming phase is performed once per run. Operations necessary to prepare the run are defined here. This might include writing to the registers on the VLSB to prepare for external data.

- **AsynchronousRead** - This function is called for every piece of equipment prior to the EventArrived function.

- **EventArrived** - The event arrived function is called by DATE for each piece of hardware defined as generating triggers. In principle only one trigger may be defined. The function is called continuously until it returns a value, corresponding to the arrival of a spill trigger.

- **ReadEvent** - Upon the arrival of a spill trigger the read event phase begins for each piece of equipment defined as generating data. This involves accessing the memory banks of the buffer boards and writing the contained values to the data stream.

- **Disarm** - At the end of the run the disarm phase performs the required tasks to put the VME boards back into an idle mode, for example, by disabling the receipt of external data.

VLSB Master

The arming phase of the VLSB Master prepares the encoded clock and sets the trigger mode. A single register controls the various modes of the trigger, which includes spill type and duration. The triggers must be enabled after the VLSBs have been armed and are ready to receive data, and disabled before they are put into VME access mode and the data retrieved. This is to ensure all VLSBs receive the same number of triggers. The differing source of these triggers (internal and external) mean different functionality is required of the Arm, ReadEvent and Disarm functions.

In the case of an external mode (where the VLSB Master receives triggers from an external source), the setup parameters are written during the arm phase, but the spill is not enabled until the first calling of the read phase. In the simple case of a single tracker and calibration like data taking (e.g. with an LED pulser), the read phase is called at time intervals governed by a timer. This timer is set to be a trigger which on the first call of the event arrived phase waits a given amount of time and then returns a positive value indicating that a trigger has arrived.

In the case of the full MICE experiment or a two tracker cosmic test, the spill trigger (the event which moves the DAQ into the read event phase) must be managed with hardware in order to reliably distribute the signal across all the VME crates and equipment.

VLSB

The main function of the VLSBs is to store the data arriving from the AFE boards. In order to ensure that no data is written into the VLSBs whilst they are being read out, before any data is retrieved from the first VLSB all boards are taken out of data acquire mode and placed into a VME access mode. Access via the VME backplane cannot coincide with LVDS data entry in any case.

Once the LVDS access has been disabled and the VLSBs can be read by the back plane, the contents of the banks can be read and written to the data stream.

This is either done by accessing the four VLSB banks in turn, or by treating each bank as a unique piece of equipment.

The VME access can either be achieved by single word reads, whereby each memory address of the bank is read to the data stream up to the total recorded amount of data, or the entire volume of data is read by a block transfer. However, the VLSB does not support VME side block transfers (BLT mode), but rather a read cycle whereby the PCI controller sends an array of single word reads without further need for software function calls.

Readout synchronisation

As each tracker is read out to independent VME crates, each with their own PCI bridge, synchronisation is required to ensure that data from spills in each tracker are combined. To achieve this, a CAEN V977 I/O register board was used, which has 16 input and output LEMO pins, the status of which are readable via VME. A VME process was then defined to set an output signal at regular intervals on one of the output channels; this signal was then fanned out and sent to inputs on the original board and the corresponding board on the other tracker. The EventArrived function in each LDC then polled this input pin for the signal to begin the readout process. The readout command then had only one source which was relied upon by both trackers.

In the full MICE setup this process will be used in conjunction with the other detectors. In that case, the ISIS start/target dip signal will be the readout pulse, and not a timer.

3.5.2 Unpacking

An application was written to take the output of the data acquisition and unpack it into a format used by the calibration and reconstruction software. This required an understanding both of how DATE arranges data from multiple detectors and how the VLSBs themselves arrange data from their corresponding AFE boards.

The make-up of each individual word containing the data for a single channel in a specific event has been shown in Table 3.3. DATE writes these words to file in the order they are retrieved from the VLSB memory banks, which is in turn the order the VLSBs received the data from the AFEs. In this fashion, the information for all events for a given DFPGA (128 channels) is retrieved from a VLSB memory bank and written to file before the first event of the next 128 channels. Unpacking refers to the process of re-ordering the data words according to events, as well as retrieving specific information from the bit-packed word. This involves shifting the bits within the word according to the word make-up.

The data stream contains information for all boards and their banks. The division between boards is marked by a selection of header words containing the board ID and the lengths of each of the board's memory banks. From the lengths of the banks and by counting the number of subsequent data words read out, the bank (and hence DFPGA) the data word came from can be ascertained. This provides any subsequent processes with a usable data format.

3.6 Calibration

The calibration of the MICE trackers can be broken into three sections: values which must be set in the tracker hardware to achieve an optimum physical response from the detector; values which must be set in the electronics to gain an optimum efficiency; constants which translate read values (such as ADC and TDC) into meaningful physical quantities.

The first section consists of the optimum bias voltages applied to each module of VLPCs. The correct bias voltage will generate the largest amount of signal for an incident photon whilst balanced against the accidental rate of signal inherent in the system as discussed earlier. The second section refers to settings which are not essential to running the detectors to specification, but which become useful in the event of full MICE data taking. Zero suppression falls into this section; suppressing non-signal channels is not a requirement of detecting particles passing through the detector, but it will be become necessary when the event rates increase and the strain on the MICE DAQ becomes more severe. The final section involves translating the read ADC and TDC values into meaningful quantities such as the number of photo-electrons and the time of the signal, respectively.

Therefore, the calibration of the MICE trackers can be broken up into four stages, which must be completed in order. They are as follows:

1. **Bias** - The range of possible bias voltages for VLPCs lies between 0 and 10 V, where the accuracy of the voltage setting depends on the Digital to Analogue Converter, and corresponds to approximately 40 mV. This is set at one value per module of 64 channels and determines how the VLPCs will behave under noise and signal.

2. **Discriminators** - The discriminator threshold, governed by the set value of VTh, affects the integrated charge (proportional to ADC count, and therefore

affected by the bias) at which the discriminator bit will be set to 1. This is needed both for zero suppression and for the timing circuit.

3. **Light Yield** - Once the bias has been set, the response under known amounts of light can be measured in order to correlate the observed ADC count to the incident light yield.
4. **Timing** - As described earlier, upon the discriminator firing a time-proportional current is turned on and digitised into a TDC count at the end of the integration window. This requires a correctly calibrated discriminator.

3.6.1 Calibration Software

A dedicated piece of software was written for the purpose of calibrating the various aspects of the tracker electronics and photon-conversion. The code was written in Python and is independent from the rest of the MICE software, including readout and reconstruction, and only depends on ROOT for random number and fitting functionality.

The general concept of the software was to be able to read in data in a number of formats (e.g. the raw binary output from DATE; an unpacked text file etc.) and parse the data into classes representing each channel-event, or all the data associated with a single channel of a run. From these items, further information can be calculated and stored in objects representing the higher level electronics, such as a module or the TriP-ts. Manager classes for each calibration aspect can then use these calculated values to find the optimal setting for each electronics object. Support classes for storing configurations were also added.

Electronics Simulation

In order to test the quality of the calibration algorithms a simplified simulation of the electronics was written. The application has a number of levels of complexity to mimic the behaviour of the photon conversion and readout electronics.

The first level of an electronics model requires the conversion from a physical feature into a measured charge, which is in turn digitised. The individual gains, or multiplicity of the electron avalanche, is well documented from the production of the VLPCs. In this model the gain of each channel is controllable. The physical signal, i.e. an incident photon, is multiplied by the gain to produce the charge integrated by the electronics. The conversion of this charge to ADC counts requires a further parameter, defined as the ADC resolution (in fact referring to the resolution of the entire electronics process not limited to the ADC chip). The lowest level of the electronics therefore requires two set parameters, the gain and resolution of the ADC.

In order to model the signal induced by the LED, a few parameters were added onto the existing layer, such as accommodating the pedestal. Assuming the number of photons produced by the LED follows Poisson statistics, the LED itself is defined merely by the mean number of photo-electrons produced. The expected ADC count for this channel can now be calculated using the mean number of photo-electrons and defined gain parameters. This is then smeared according to an assumption of the ADC resolution. For a collection of photons about a given mean, the ADC distribution is therefore modelled, and an example is shown in Figure 3.10.

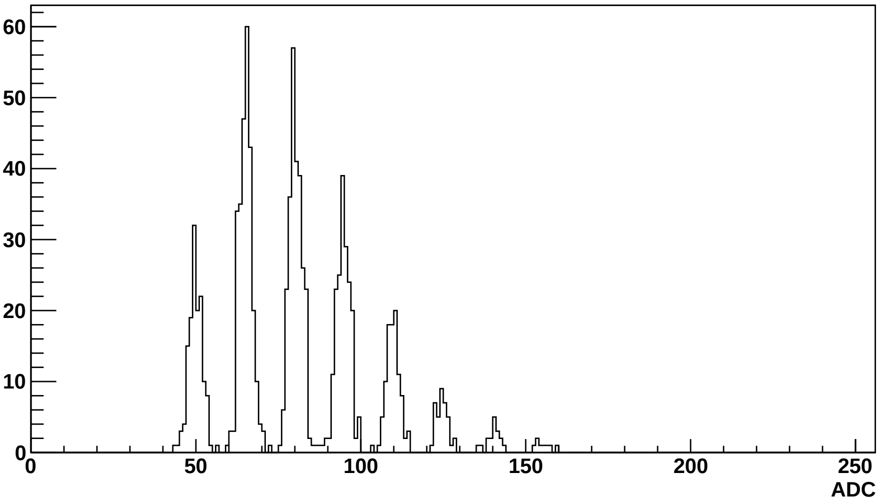

Figure 3.10: Electronics toy Monte Carlo modelling an LED with a mean PE of 2, gain of 15 and pedestal of 50.

At the simplest level, a model of the discriminator can be constructed by comparing the calculated signal charge to a set threshold; if the charge is higher than the threshold the discriminator is set to one. In this case, no model of the pulse shape is considered.

The timing is related directly to the discriminator in that the time current is active from the moment of the discriminator firing. Ignoring time walk, the start time of the current is the same as the time of the signal, if the discriminator fired. The time remaining to the end of the integration window is multiplied by a current, and added to the TDC pedestal. There is an additional layer in that the current produces a charge which is subject to the digitisation resolution previously mentioned. As the time-charge-TDC conversion is completely linear, the current can be defined to convert time into TDC counts.

Channels and Run information

For each channel, at every event, the ADC, TDC and discriminator information is stored in histograms. Over the course of a run dedicated to taking non-zero suppressed data, the ADC counts combine to provide an image of how the channel is currently behaving. At low bias only the pedestal would be seen, whereas at a high bias a large amount of noise and eventually breakdown is visible. These *pedestal* histograms are the primary tool for developing the optimum bias and calculating the VLPC gain.

A peak finder was implemented to look at each individual channel and find the pedestal and multi-photo-electron peaks. Once the peaks had been located further information was optionally stored: the average distance between the found peaks is used as an estimate for the channel gain; the integral between and above each peak is used in the bias calibration, along with the ratios of these integrals, as is discussed later. The peak finder class also provides the option to make a simple Gaussian fit to each peak, with the mean centred on the peak and limits of the fit constrained by the average distance between peaks. The peak finder also takes into account the possible emergence of false peaks produced by firmware complications.

3.6.2 LED and Triggering

To supply test signals to the VLPCs, two types of LED pulser were used during the calibration: a single module and a single cassette LED. The single module provides light to 128 channels and the full cassette LED to 1048 channels. The generator used was capable of delaying the light pulse with respect to its trigger pulse, which enabled the light to be correctly timed with respect to the electronics trigger.

The delay applied to the trigger was set to match the delay inherent from the cosmic ray trigger and logic. The use of the LED for timing allowed the difference between the generation of light (and arrival of signal into the VLPCs) and the arrival of a trigger to be matched to the integration window clock and veto. The light-to-trigger delay, as measured by the LED, is then also the delay applied to the veto with respect to the start of the integration window. A diagram of the triggering system is shown in Figure 3.11.

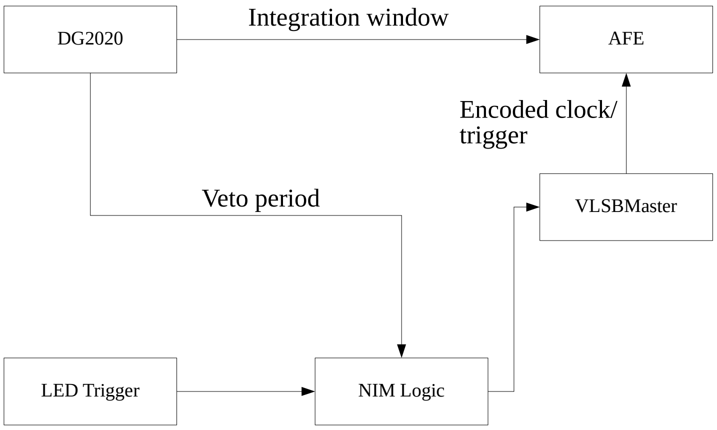

Figure 3.11: Trigger logic for calibration data taking, using a DG2020 digital function generator to define to integration window and veto periods.

The optimal pulse amplitude, width and delay with respect to the trigger were found to avoid saturating the system with light, and scans in time were made to find the position at which the light was centred in the integration window, and to find the optimum phase of the veto period. An example of the ADC distribution for a single channel with light from an LED can be seen in Figure 3.12.

3.6.3 VLPC Bias Settings

In order to automatically and repeatably calibrate the VLPC bias, an optimally biased channel must be quantitatively defined. DØ, which has an identical electronics set-up, uses the definition that the number of ADC counts above the most-probable value of the 1 PE (singles) peak should be less than 1.5% of the whole ADC distribution for a given channel. This single photon noise, which is different from the pedestal, is defined as the dark count. As MICE has different fibres as well as an alternately defined integration window (which will affect the ADC counts seen) this has been variously selected as between 1.7% and 2%. In order to accommodate this variation the optimum dark count ratio is a parameter of the calibration software.

The 2% requirement comes from the expectation that 2% dark count will correspond to the minimum bias at which the maximum efficiency of the VLPC is

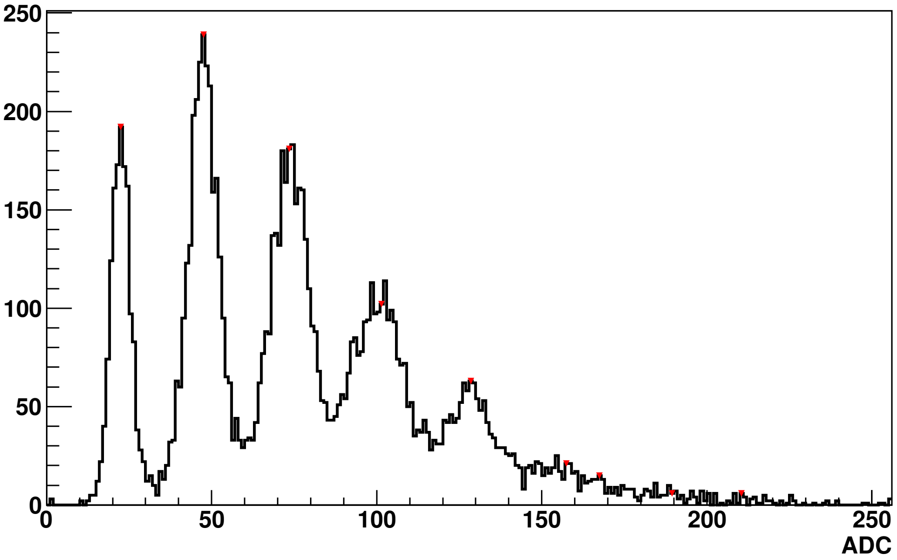

Figure 3.12: ADC distribution for a single channel with signal from an LED. The red marks are the output of the ROOT peak finder used by the calibration software to calculate the VLPC gain.

achieved.

Given that the bias is applied only once to each group of 64 channels, there is the added consideration that not all channels will behave identically under the same bias. It then becomes necessary to find the correct median across the 64 dark count ratios, and a number of possible algorithms have been considered.

- **Module Average** - The simplest method is to take the average dark count ratio of all the channels in a module and find the bias whose average is closest to the optimum value. However, this could be heavily skewed by those channels which are either extremely noisy or dead, and a single channel may negatively alter the bias for the remaining channels.

  $$\langle DC \rangle = \sum_{i}^{N} \left(\frac{\int_{1PE}^{max} ADC}{\int_{0}^{max} ADC} \right) / N_{chan} \tag{3.2}$$

  This mean dark count $\langle DC \rangle$ becomes the quality factor used to calibrate the groups of VLPCs.

- **Module Average minus outliers** - To counter the effect of noisy or dead channels, the channels with the highest and lowest ratios could be removed from the averaging process.

- **Weighted Quality Factor** - A further algorithm involves weighting the deviations from the optimum ratio. The difference between the measured dark count ratio and the optimum D is calculated for each channel and averaged over the module. This average is then divided by the ratio of channels which contained a 1PE peak; this avoids tending towards those channels with no 1PE peak, which cannot deviate by more than the optimum value itself. This algorithm is defined by:

$$QF = \frac{1}{S\langle DC\rangle} \tag{3.3}$$

where S and DC are defined as:

$$S = \sum^{N} \delta_{1PE} \tag{3.4}$$

$$\langle DC\rangle \to \sum_{i}^{N} \left(\frac{\int_{1PE}^{max} ADC}{\int_{0}^{max} ADC} - D \right) / N_{chan} \tag{3.5}$$

Examples of a single channel under biased, optimally biased, over biased and in breakdown are shown in Figure 3.13.

Calibrating Bias purely from noise

As the definition of optimum bias depends on the location of the single photo-electron peak only, it is possible to calibrate the bias settings using only the natural noise in the VLPCs. Finding the singles peak, especially for low biased or low gain VLPCs, is an issue.

In order to investigate the viability of using only the noise to calibrate the bias, a full sweep of possible biases was applied to a single module and the ADC data taken at each value. The runs taken at each step were then subjected to the algorithms described earlier, and the results shown in Table 3.4.

As the single photo-electron peak is not always clearly visible in the signalless ADC distribution, an additional method is to use an LED pulser to supply signal to the VLPC and locate the position of the 1PE peak. This position is then used to calculate the dark noise ratio from the ADC distribution taken at the same bias but without the LED emitting light.

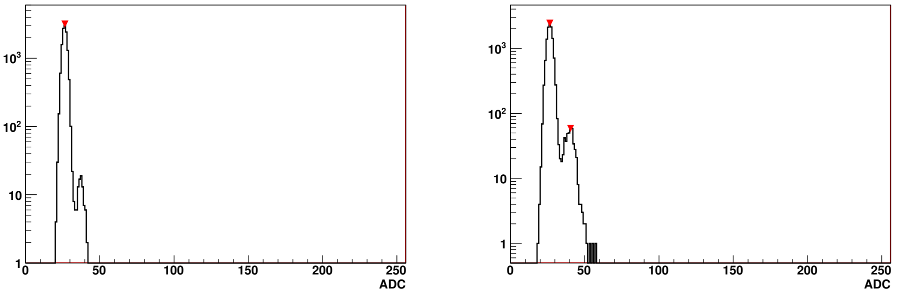

(a) Under biased (b) Optimally biased

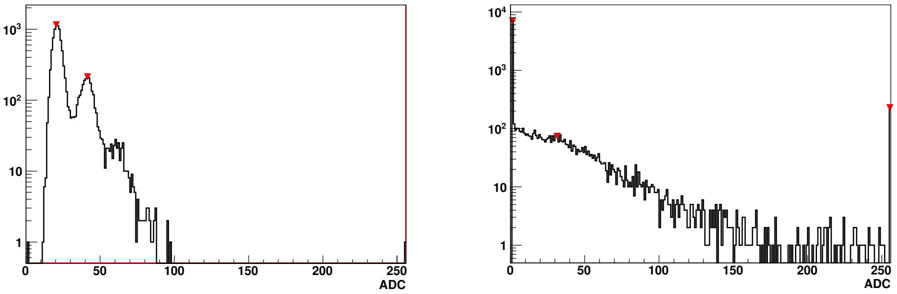

(c) Over biased (d) In break-down

Figure 3.13: ADC distributions for a single channel at varying biases.

Efficiency Curve

Whilst a direct measurement of the quantum efficiency is difficult using a statistical light source such as an LED pulser (di-photon sources are ideal for this measurement), some function of the quantum efficiency is obtainable. By defining a new efficiency as the average light yield achieved with an LED minus the average light yield at the same bias without the LED, the efficiency should rise and then plateau with the applied bias. It is then possible to choose a bias which lies on this plateau but produces minimal noise. An example of how the efficiency changes with the applied bias can be seen in Figure 3.14.

The definition of when the efficiency has plateaued can be a function of the gradient of the efficiency with bias. Once this gradient falls below a defined value that bias is chosen, regardless of what its corresponding noise ratio is. Development of this form of calibration is a feature yet to be added to the software suite.

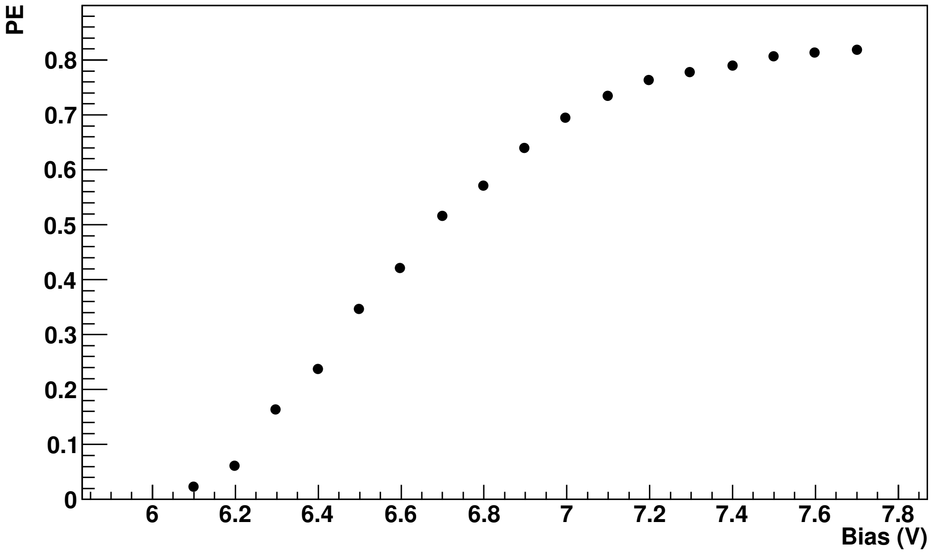

Figure 3.14: Change in difference of the mean PE between a module with the LED on and off as a function of the applied bias.

Results of the bias calibration

A complete scan of the possible bias voltages was performed for each cassette of one tracker using an LED pulser. This pulser is large enough to fit the fibre connections for an entire cassette, providing signal to 1024 channels. The bias voltage was swept from 6.0 V to 7.5 - 8.0 V (depending on the point at which the VLPCs went into breakdown) in steps of 100 mV.

By plotting the dark count ratio against the applied bias, the behaviour of the VLPCs in each region can be observed. Additionally, an optimum dark count ratio can be extrapolated to a corresponding bias. An example of this curve can be seen in Figure 3.15, and is summarised for all modules (using the averaging method) of a single cassette in Figure 3.16. The exponential fit used refers to the behaviour of the dark count as a function of bias [68].

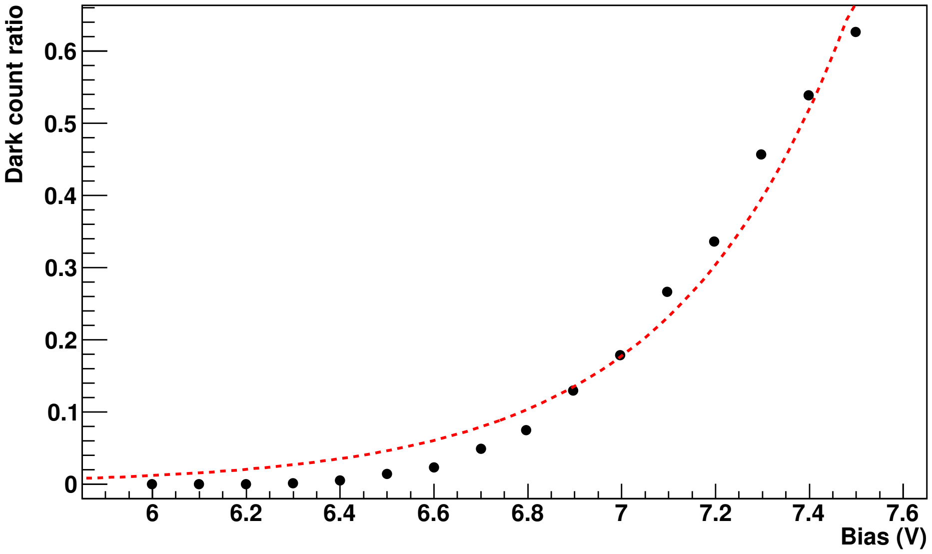

Figure 3.15: Dark count versus bias for a single module. The points are shown with an exponential fit.

Given the various options for quantifying the contribution from dark count for each module, the method of selecting the best value also has a number of options: the highest bias with an average ratio less than optimum; the bias which produces an average ratio closest to the optimum, even if it is higher; by plotting the dark count as a function of bias and fitting it, as is discussed later. For various single modules in the same VLPC cassette, the results of these algorithms are presented in Table 3.4.

As data was taken at a range of biases with and without the LED, the method described in Section 3.6.3 can be employed. Plotting the efficiency against the dark count ratio at that bias can provide a check of the behaviour of the VLPC at the suggested 2% noise region. An example for a single module is given in Figure 3.17.

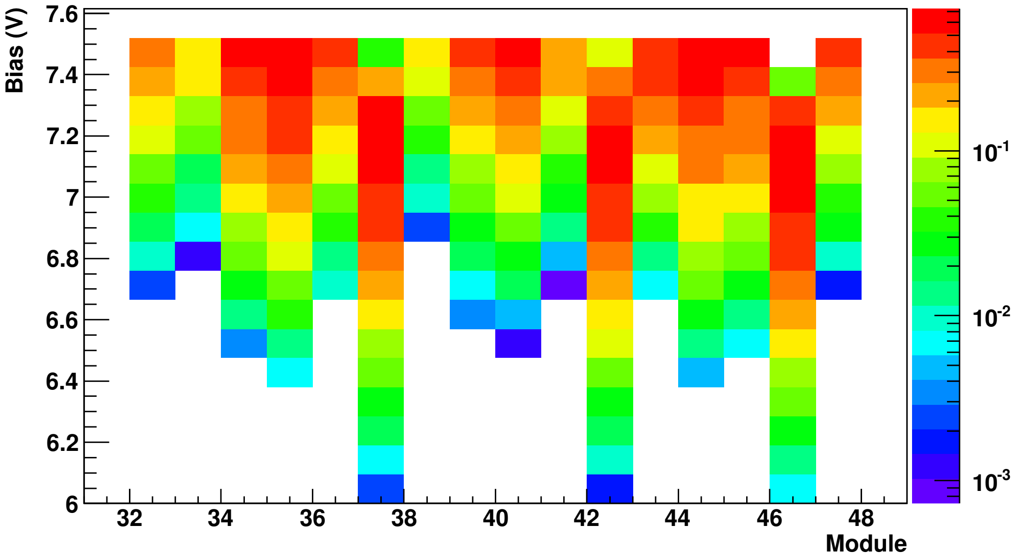

Figure 3.16: Change of dark count noise with bias for all modules of a single VLPC cassette.

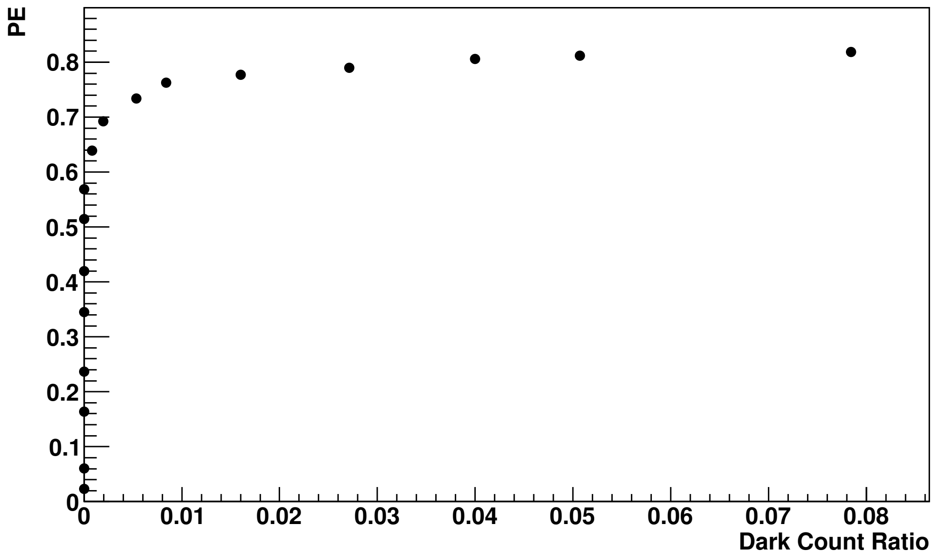

Figure 3.17: VLPC efficiency versus the dark count at a range of biases.

The final optimum bias found for selected modules based on the different algorithms suggested can be seen in Table 3.4. This table to shows the different bias voltages for individual modules returned by each of the calibration algorithms, and the corresponding ratio of noise that these voltages produced. The abbreviations are Highest Average (H.A.), which corresponds to the highest bias whose average dark count was below the desired value, Closest Average (C. A.), which refers to the bias whose average dark count was closest to the optimum, Fitted (F.) which gives the bias corresponding to the fitted value at the optimum dark count ratio, and Efficiency (E.), which corresponds to the bias whose efficiency value is approximately at the shoulder of curve. This last value is not well defined. N stands for the dark count ratio at that bias.

Board	Module	H. A.	N.	C. A.	N.	F.	N.	E.	N.
6	0	7.1	1.5	7.2	2.0	7.15	2.0	7.1	1.5
6	4	6.3	1.5	6.3	1.5	6.3	1.5	6.3	0.4

Table 3.4: Summary of optimum biases (V) found by the different methods for a selection of modules, with the corresponding dark count (%).

Since the method of finding the bias which produced a noise ratio closest to some value was itself only an approximation to the efficiency method, and the issues in this approximation have been highlighted, the efficiency method is naturally the selection of choice. However, it is possible that by using only the efficiency parameter and ignoring the noise, one might select a bias which produces the maximum efficiency but at the cost of too much noise. For this study, and to produce the biases used in the subsequent cosmic ray test, the method of the highest average bias which produced noise less than a defined constant was used.

3.6.4 Pedestal and Photo-Electron Peaks

Given a signal of ADC counts adc_i from a channel with corresponding measured pedestal ped_i and gain g_i, the measured number of photo-electrons is:

$$pe_i = \left(adc_i - ped_i\right)/g_i \tag{3.6}$$

The expected number of photo-electrons for an LED pulser from each trigger should follow a Poisson distribution and the distribution of ADC counts about the mean charge produced by the total light is expected to be Gaussian. Using the LED signal at an optimum bias, it is possible to fully characterise the VLPC response to signal as a Poisson of Gaussian distributions, where the number of photons is distributed according to a Poisson distribution, but the number of ADC counts for each photon is distributed according to a Gaussian distribution. The final distribution will therefore be a sum of Gaussians, with the height of each Gaussian defined by the Poisson statistics of the photons. A function can therefore be defined and a fitting package used to find the pedestal and gains of each channel, including the spread in ADC as the light yield increases. Alternatively, a simplified method could use the peak finder to calculate the gain as the average distance between peaks.

3.6.5 Discriminators and Zero-Suppression

The method for finding the optimum discriminator threshold voltage is similar to that of the bias. The possible values must be swept through and data taken at each step. The quantity of interest is the percentage of times the discriminator fires at a given setting under a given amount of charge, typically the noise from the optimally calibrated bias.

The threshold voltage is set once per TriP-t, meaning that it controls the threshold of the discriminator for 32 channels. Any calculated quantity must then be considered across the chip in the same way as the bias calibration. The figure of merit (the percentage of hits causing the discriminator to fire) is also subjected to the same algorithms: average; average minus noisy and dead channels; a weighted average.

The optimum value is chosen to be equivalent to the optimum bias; the percentage of times fired on noise should be less than a chosen value, typically 2% for MICE. In this case, the threshold voltage corresponds directly to the ADC count at the most probable value of the 1PE peak.

In order to set the threshold of the discriminator two registers must be appropriately defined in the Trip-T, which control the voltages VRef and VTh (as seen in Figure 3.9). The threshold is primarily a function of VTh. However, in initial testing it was found that VTh alone could not provide the necessary range, and so VRef must also be set differently from its default value. As the discriminator threshold is inverted with respect to VTh and normal operating with respect to VRef (which is to say, the higher VTh, the lower the threshold, and vice versa with VRef), and the added difficulty of calibrating a multi-parameter setting, VRef was set low enough such that the range of the discriminator threshold could be controlled by VTh alone.

A scan was performed for a single AFE board by initially choosing points in the region expected to see the maximum gradient of discriminator firing. Once the region had been defined a finer scan was performed. The optimum is found by choosing the value of VTh which should give, for a single TriP-t, an average discriminator firing on noise, at optimum bias, closest to 2% (or that noise ratio corresponding to the optimum bias found by other means). The behaviour of the discriminator by varying the values of VTh with a set VRef can be seen in Figure 3.18.

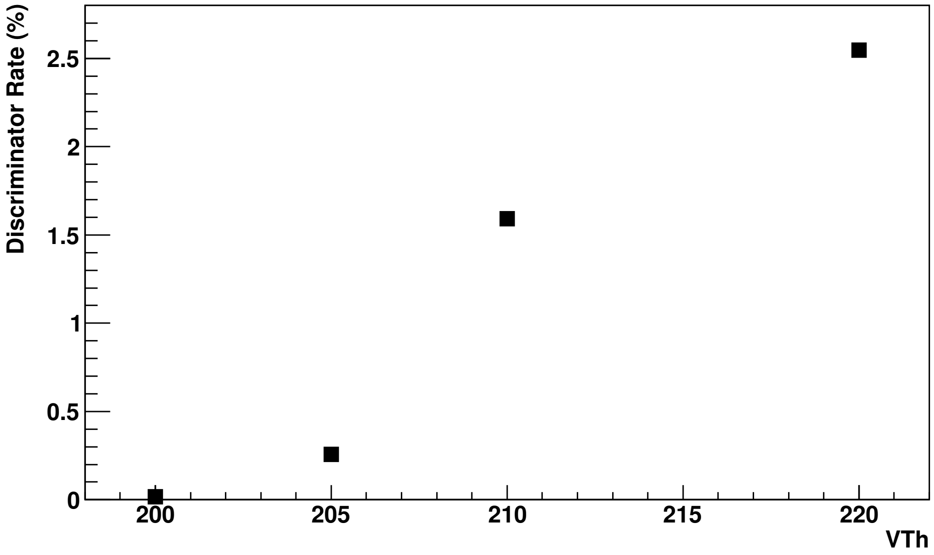

Figure 3.18: Rate of discriminator as a function of VTh for a noise-only system.

Discriminator cross-talk

Given the close proximity of the discriminator circuits on the TriP-t, cross-talk is of concern, as is the behaviour of the circuits at the extremes of the possible threshold voltage. Experience of the TriP-ts on DØ show that at the minimum possible threshold, at which one would expect all the discriminators to fire, none do. In order to investigate this, along with finding the optimum threshold voltage based on the firing ratios, the behaviour of neighbouring channels should be looked at. This would require significantly more calibration data than simply finding the optimum values of VTh for the discriminator, and sufficient time was unavailable prior to writing.

Discriminator-Analogue cross-talk

Since the discriminator circuits turn on within the TriP-t, an effect begins to emerge on the analogue channel in the form of noise, blurring the distinction between photoelectron peaks. The difference between a channel without discriminators and a channel in which the discriminator has been set very low (and so the majority of neighbouring channels are active) is shown in Figure 3.19. This is attributed to the proximity of the analogue and discriminator circuits within the chip. It is a factor of the production and can only be mitigated by minimising the firing rates of neighbouring channels. Given the percentage of channels which will see signal in the trackers, even in the presence of backgrounds, is quite low, this is not expected to be an issue in MICE. However, it is an issue for calibration.

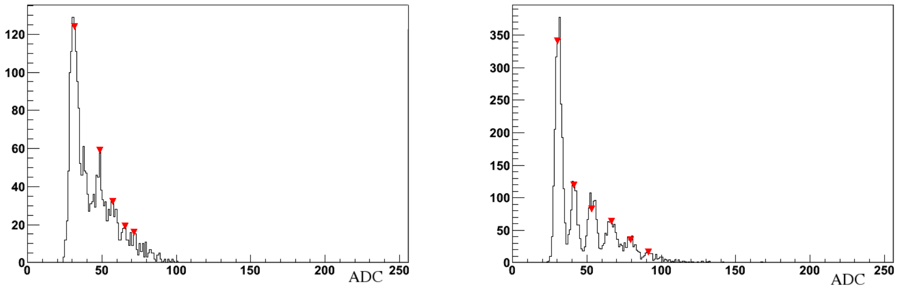

(a) With discriminators

(b) Without discriminators

Figure 3.19: Discriminator-Analogue cross-talk.

3.6.6 Timing with Time to Analogue and ADC

The time at which the discriminator fired can be calibrated from the measured TDC value by injecting charge at known intervals. For a constant time current source, the TDC value is directly proportional to the injection time, so a series of points can be fitted to get the gradient of the current in TDC counts. The interval at which charge can be injected is limited to the frequency of the board clock, however this still allows for 5-6 injections within an integration window.

Once data has been taken with charge injected at each time interval, the calibration software fits the TDC distributions for each channel at each run and plots the fitted mean value against the injection time. The gradient and offset (which should be the ADC pedestal) are then extracted from the fit. The gradient of this slope is governed by the time current source, IBt, which is a set parameter of the TriP-t. It is desirable that this should be large enough to fully maximise the number of TDC counts per nanosecond, yet not saturate the 8-bit range of the TDC at any point.

A preliminary time calibration has been performed on one AFE board. The board was set to an internal mode (with the integration window defined by its own internal clock, and triggers generated by the VLSB master) and the charge injection mode used. Firstly, the time current source should be set so that the maximum recorded value of time does not saturate the TDC. In order to accommodate this, charge should be injected at $t = 0$ (which would incur the largest TDC value) and the time current source adjusted until the TDC no-longer saturates. In this initial test, the scan in time was performed prior to the level optimisation, generating the saturated point shown in Figure 3.20. The input charge was set to maximum and the discriminator constant to ensure that any time-walk was minimised. The TDC value as a function of the time of injected charge can be seen in Figure 3.20, where the saturation of the ADC is shown by the $t = 0$ point not matching the fit line. The error on the TDC value is given by the RMS of the TDC distribution for that injection time. Whilst the resolution needs improvement to match the 0.5 ns resolution shown by the bench tests, the timing circuit of the MICE trackers can give information about the arrival time of the particles.

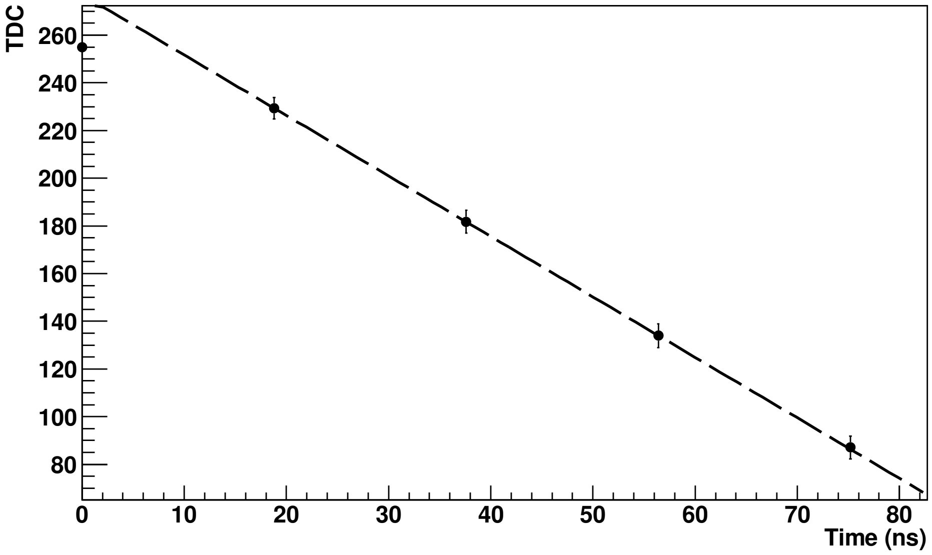

Figure 3.20: TDC Resolution taken for a high current, which would normally saturate the TDC at the minimum time setting.

Cross-talk is also an issue between the TDC and ADC, since the TDC current can contribute to the measured ADC count. The conflict between neighbouring discriminators must also be taken into account. As such, the charge injection map feature is used to avoid injecting charge into neighbouring channels, and each of 32 channels within the TriP-T are calibrated in four sets of eight.

Measuring Time walk

As the point at which the timing circuit activates depends on the time when the integrated charge rises above threshold, variations in the rise time of pulses will cause variation in the signal, all occurring at the same time, to produce different TDC values. In order to measure this effect, the differing TDC values for a channel which has had charge injected at a precise time, but with different charge levels, can be compared. The discriminator level and injection time are kept constant, but the level of charge injected is varied. However, given the level of charge injection on the TriP-t is very large, this is an inefficient method. Preferably, by varying the discriminator level for a given charge level and injection time, the integrated pulse shape can be measured as a function of time, which characterises the time walk in each channel.

Time-Analogue cross-talk

The proximity of the analogue and time channels on the ADC chip create the possibility of cross-talk between the two signals. To investigate this, the ADC value for a constant input charge could be compared for varying levels of the TDC (achieved by varying the time at which the discriminator fires). The issue is that the method of precise time charge injection is not accurate when using analogue values. The charge is often so large as to saturate the ADC. An accurately timed LED pulse could be used in place of the charge injection. This must also be handled by calibrating in groups of non-adjacent channels, but is not expected to be an issue in MICE data taking.

3.7 Particle Tracking

Reconstruction of the particle trajectory through the trackers requires a combination of the detector response, channel geometry and magnetic fields. Based on the known performance of the photon conversion and electronics, cuts can be placed to ensure a channel is only considered if the probability of it containing a signal hit is high.

The geometry of the channels, confirmed during the tracker construction, is then used to provide a position in space for this signal hit. Its timing is obtained from the electronics. Combinations of channels with signals that overlap can form space points, and collections of space points combined to form tracks. For track fitting, the trajectory of particles as they travel in a magnetic field must be optimised with a filtering algorithm.

3.7.1 Simulation

The simulation of the MICE trackers is achieved with a Geant4 based package, where the materials and geometry of the tracker and solenoid are accurately represented. The solenoid and other MICE magnets are modelled based on their current and coil configurations. The field maps for the magnets are generated at the start of the simulation run. A number of different algorithms can be used to generate the map from geometry, or an externally generated or measured field map can be used. The electronics is modelled with a simple approximation to the relationship between light yield and ADC counts, depending on general values for the fibre attenuation, quantum efficiency, gain and other parameters. Once the light yield in the fibre has been estimated from the energy deposition, it is converted into an ADC count and smeared with a Gaussian distribution.

3.7.2 Reconstruction

The simulation and reconstruction framework for the trackers, G4MICE, arranges this reconstruction process in layers. The initial state involves unpacking the data into a primary hit class which stores the raw information obtained from the AFE boards, called a VLPCHit. The hit stores the electronics channel information, along with the data obtained from the front end (ADC, TDC, discriminator). The equivalent step in simulation is handled by a SciFiHit which stores the Geant step energy deposition along with time, momentum, position and detector information such as the fibre through which the particle passed.

Digits

Reconstruction of digits from raw AFE information is performed with the DigitRec function. In order to ascertain which channels observed a large enough light yield to be considered as signal candidates, their numbers of photo-electrons must be calculated. To convert the recorded ADC count for this channel into photo-electrons, the pedestal and gain for this channel are retrieved from a calibration class. The method of obtaining these calibration values was described earlier in Section 3.6.4.

Once the number of photo-electrons in a channel has been calculated, a cut is made on it. This is an arbitrary cut designed purely to limit the presence of noise and was initially implemented in the absence of zero suppression. With the discriminators calibrated and active, such a cut may no longer be required.

The mapping of the electronics channel to the physical fibre is also performed at this stage. There is a direct correspondence between a single electronics channel, which is defined by three integers signifying the AFE board, VLSB bank and DFPGA channel number, and the fibre which is also defined by three integers signifying the station, plane and fibre number. The mapping of electronics channel to physical channel is stored externally and retrieved with a cabling class. The resulting SciFiDigit therefore records the physical channel, time and light yield for a given event, and is kept if its light yield is above threshold.

Clusters

Given the overlap between fibre ribbons it is possible that a single particle can pass through two channels, generating two separate signal hits. In order to avoid this degeneracy, any two digits which are adjacent, above threshold and in the same event are combined into a cluster. Any remaining digits without a neighbour are also converted into clusters. The initial implementation of this process did not consider the possibility of electronic or optical cross-talk, and therefore only a maximum of two adjacent clusters were ever combined. An investigation into cross-talk at the reconstructed level is presented in Section 3.8.3.

Next, the physical position of the cluster centre is calculated and stored in the SciFiDoubletCluster object. With the central fibre number being a defined parameter at $x = y = 0$, and the width of each channel and its orientation also defined, the position of the centre of the cluster is calculated by its deviation from the central channel and a coordinate transformation from the orientation of the plane.

Space Points

The overlapping of clusters in different planes provides a 2D region through which a particle may have passed. The triplet active regions for each station in tracker one can be seen in Figure 3.22, where the positions have been found by Monte Carlo and cut for one half of the active region. The three planes within each station provide a number of combinations in which clusters can overlap; in the instance in which all three possible clusters produce a signal above threshold, and all three are combined to form a space point, this type of point is labelled a triplet. In the event of only two clusters overlapping, the space point is termed a doublet.

The clusters are iterated over and possible combinations are passed to an algorithm which decides whether the clusters should form a triplet or doublet, and keeps or rejects the candidate space points based on the sum of the triplet's constituent channel numbers. Beginning with all possible triplets, defined by clusters in the same station but different planes, the channel numbers are added and tested to see if they fall within a certain range.

The channels within each plane are numbered in such a way that if a point is taken anywhere in the active region of a station, the sum of the channel numbers covering that region will equal the sum of the central channel numbers. The vector sum of this addition is given in Equation 3.7, where u_i, v_i and w_i are the channel numbers of the clusters in the triplet, and u_0, v_0 and w_0 are the channel numbers at the centre of the plane, as shown in Figure 3.21. This is well proven from data from the first cosmic ray test, where this algorithm was not used (a cut on the internal residual, or separation of the clusters, in combination with higher light yield cuts was used instead).

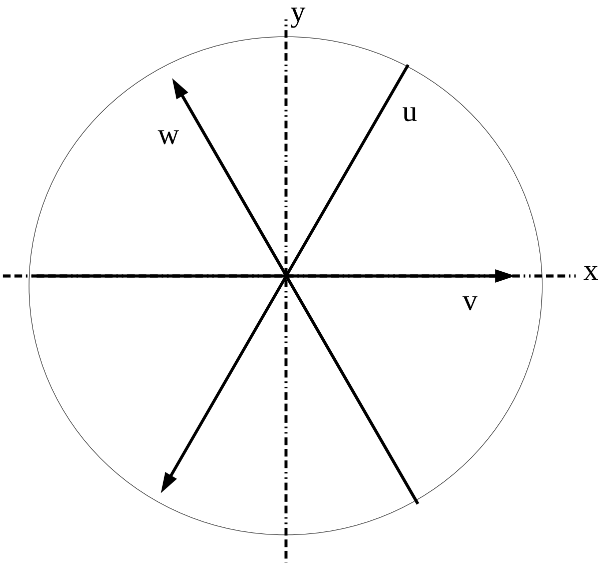

Figure 3.21: Station plane channel ordering, with a constant vector sum of 0 for triplet points.

$$tr_i = (u_i + v_i + w_i) - (u_0 + v_0 + w_0) = 0 \tag{3.7}$$

Once this conjecture had been proven from data (as shown in Section 3.8), it could be used as a method of verifying the quality of a space point made from three clusters. In this way, it forms the major component of the space point pattern recognition. All clusters are sorted into station and plane and ordered by their light yields (low light yield clusters had a detrimental effect on the space point finding, as is discussed later). Clusters in the same station but different planes are then combined to form triplets and tested to see whether the sum of their channel numbers lies within an acceptable range; if so, the point is kept and the constituent clusters removed from the list.

Once all possible combinations of triplets have been considered, the remaining clusters are formed into doublets and tested to see whether the resulting space point lies within the tracker active volume; if so, the point is kept but the clusters remain within the search.

The intrinsic error on each space point depends on the physical area covered by the overlapping fibres. Since the tracker stations contain three planes at 120° to each other, the active region can take many shapes, and hence the errors will differ depending on these shapes. In order to fully test the treatment of the fibre combination in the reconstruction and simulation, a high statistics simulation was performed without multiple Coulomb scattering or energy loss considered in the physics list. This allowed an investigation into the effects of the tracker geometry on the space point error without other effects. Using the results of the MC residual with the position of the space point, the $x-y$ errors of each type could be considered, albeit without the added complication of the overlapping regions. A full treatment would ideally consider all possible combination of overlapping regions, or ideally keep the unbiased $u-v-w$ reference frame and perform a track fit with the original clusters.

By plotting the MC position of a particle in each station, and making requirements on the sum of the channel numbers, the effect of this requirement and the overlapping region for a triplet space point can be seen. This is shown for tracker 1 in Figure 3.22.

Pattern Recognition

The set of accepted space points is then used by the track reconstruction for pattern recognition and fitting. The pattern recognition algorithm finds sets of points likely

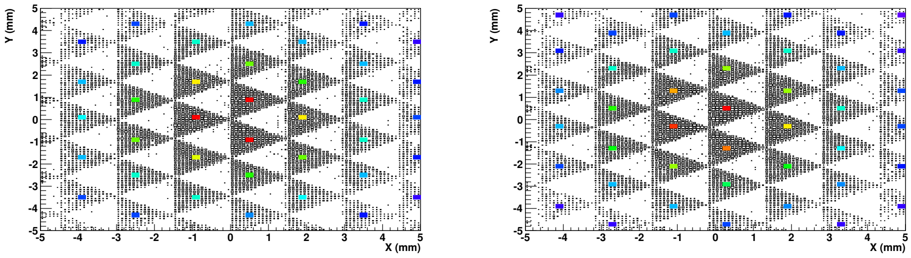

(a) Station 1 active regions (b) Station 2 active regions

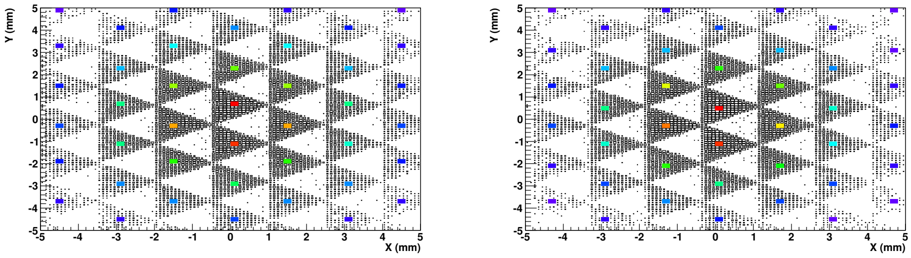

(c) Station 3 active regions (d) Station 4 active regions

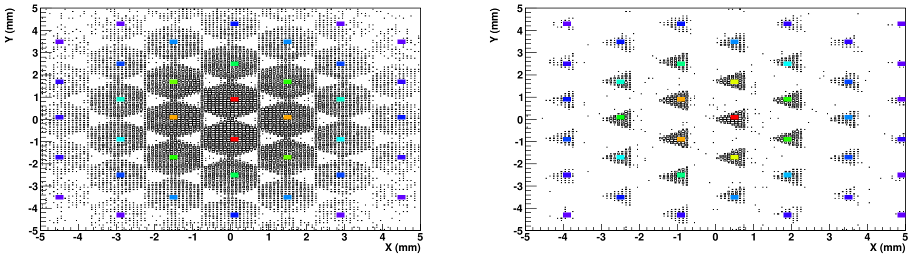

(e) Station 5 active regions (f) Station 5 active regions

Figure 3.22: Tracker 1 active regions drawn from Monte Carlo.

to have been caused by a particle passing through the detector (given the expected particle motion) and passes them to the track fitter. The standard method is to define a path through which the particle is expected to have passed and to select all space points which lie on that path.

Given the known time resolution of the tracker TDCs, it becomes possible to define a likely track as those points which fall within a finite amount of time (defined by the time of flight of the particle through the tracker), rather than those which match an expected trajectory given the magnetic field. Space points are first sorted into tracker and station, then, using the measured resolution of the TriP-t, all groups of five space points whose calibrated time lies within one unit of variance of the combined TDC resolution and expected time of flight through the tracker are kept as pattern objects.

Track Fitting

Pattern objects are then given to a track fitter. A number of track fitting options have been considered for their accuracy, efficiency and portability given the requirements of the trackers and their position within MICE. They are:

- Analytical Helix
- χ^2 Minimisation
- Kalman filter

The Kalman filter [73] was chosen as the default method for its documented performance when fitting tracks with stochastic noise [74, 75].

3.7.3 Kalman Filter

The measurements provided independently by the tracker stations, in the form of space points, can be considered to be measurements in a noisy system. A prediction can be made for the location of each point, which depends simply on the state of the particle and the system (in the form of the magnetic field) and the noise in two forms: firstly, the noise on the measurement comes from the uncertainty of the particle's true position given the area of the space point active region; secondly, the noise of the prediction produced by multiple Coulomb scattering and energy loss within the tracker.

A commonly used method for estimating the state of a noisy system is a Kalman filter. The original implementation of a Kalman filter within the MICE and tracker software was an external package, Recpack, developed as a general purpose reconstruction framework [77]. Recpack allows for transporting particles through many different media and magnetic fields, combining measurements from different detectors to fit a particle state. Given the more simplified nature of the MICE trackers (a single, highly uniform magnetic field, limited material effects and a low number of measurements per particle), and the high processing cost of using Recpack, it was not felt to be the best option for tracker-only reconstruction. Both Recpack and a bespoke Kalman filter have been used in the analysis of real and simulated tracker data.

Implementation

A Kalman filter provides a method for incrementally updating the state of a system based on measurements which depend on that state, and includes the possibility of errors and system dependant noise on both the measurement and the state.

Given an initial state, a prediction is made based on some understanding of the system. In the case of particle tracking these are the equations of motion for the particle. This prediction is made at the next step at which a measurement is available, and will include the propagation of the state covariance and any noise on the state induced from the previous step. The gain is then calculated, which provides a quantitative method for updating the state based on the measurement now available. Loosely, the gain depends on the errors of the prediction and the measurement. If the measurement is more accurate, the state will be updated to one which is closer to producing that measurement; if the prediction is more accurate, the update on the state may be less severe.

The process of the filter can be summarised as follows:

- **Initial Seed**

  A seed state must be supplied to the filter in order to make a prediction of the first measurement. This is taken from the output of the pattern recognition. All points which have passed the pattern recognition are used to calculate the transverse and longitudinal momentum from the equation of a helix.

- **Propagation of state and covariance as a prediction**

  The propagation of the state from one measurement to the next can be achieved by using the equation of a helix, or the transfer matrix for a drift space in the absence of a magnetic field.

- **Calculation of the gain**

The update of the state depends on the accuracy of both the state prediction and the measurement. The gain quantifies how much the prediction should be trusted over the measurement when the update of the state is made. It is defined in as

$$K_k = \frac{P_k^- H^T}{H P_k^- H^T + R} \tag{3.8}$$

where K_k is the gain of the k^{th} point, H is a matrix which maps the state into a measurement, P_k^- is the covariance matrix of the predicted state and R is the covariance matrix of the measurement.

- **Update of the state**

  The state is updated at each iteration based on the measurement and the weighted state prediction. The state then becomes:

  $$x_k = x_k^- + K_k(z_k - H x_k^-) \tag{3.9}$$

  where x_k is the state at the k_{th} point, x_k^- is the predicted state and z_k the measurement. Finally, the state covariance changes as

  $$P_k = (I - K_k H) P_k^- \tag{3.10}$$

  where I is the identity matrix.

These steps are repeated for each measurement in turn. In a single iteration, the covariance on the state at the last measurement will be better than the first. In order to smooth the quality of the state over all the measurements, additional iterations of the process are required. This is achieved by reversing the filter and stepping backwards towards the first measurement. This process can then be repeated an arbitrary number of times, however this is of little benefit to the state covariances.

A collection of C^{++} classes was written to provide the necessary functionality of a Kalman filter for the MICE trackers. The code was broken into sections with classes to handle the state, propagation, updating and overall management.

Performance

The Kalman filter was tested under conditions similar to those expected in MICE by generating randomised data in patterns similar to the conditions under scrutiny. These consisted of straight tracks which would be created by cosmic rays, and helical tracks produced within the tracker solenoid. A transfer matrix approach was taken where the matrix used to generate the data was also used to reconstruct it, although the original data had been randomised and the particle state left unknown. For straight tracks, a simple drift space transfer matrix was used. For helical tracks, non-linear transport equations are used, precluding the matrix approximation.

The evolution of the state covariance matrix for a single system state can be seen in Figure 3.23, where the covariance of the state improves with each iteration.

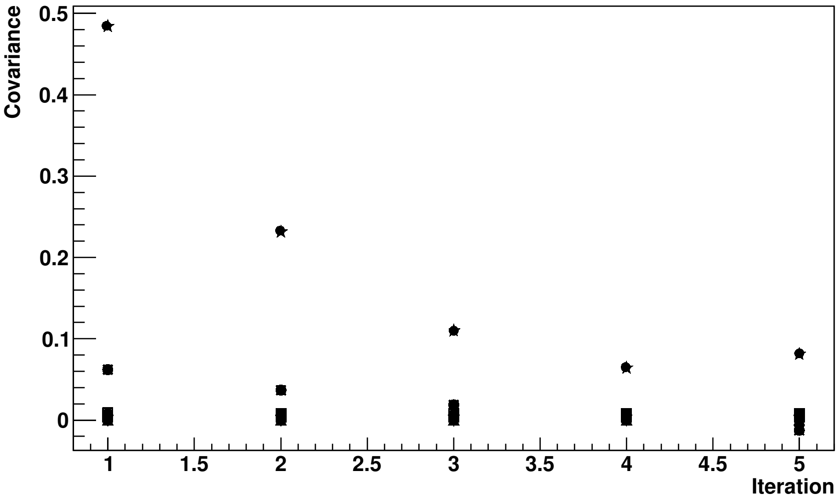

Figure 3.23: Evolution of the state covariance matrix for set initial covariance, measurement error and state prediction noise. Each marker corresponds to an element of the covariance matrix, where many of the elements are overlaid in the figure.

With the Kalman filter shown to work in simple testing, it could then be used in conjunction with other sections of the tracker reconstruction, including the space point reconstruction and the pattern recognition by timing.

3.8 Cosmic Ray Test

The first cosmic ray tests of the MICE trackers were performed in 2008-2009. The calibration was obtained in a different manner to that already described above. Whereas the method described in this chapter used the ADC noise rate as a method of finding the optimum bias, the previous method relied on optimising the induced photo-current. The data acquisition also used a different system, in that the complete initialisation and readout was performed with an Excel-based Visual Basic package used for AFE-IIt debugging. The functionality was identical to the system described above. The reconstruction relied on an external Kalman filter provided by the reconstruction toolkit RecPack.

The trackers were tested separately using scintillating panels above and below the tracker body to act as a trigger; a coincidence between the two panels would initiate the trigger. Approximately 10 cm of lead was placed between the bottom of the tracker and the lower scintillating panel to act as a momentum veto, which is taken into account in the simulation.

3.8.1 Light Yield Correction

The dynamic range of the ADC means a high light yield event could saturate the ADC count. A method was developed in order to compensate for this effect on the measured light yield, which would observe a lower than true mean value given the rejection of saturated (i.e. very high light yield) channels. For each cluster, at each time it was hit, a histogram of light yields was filled up to the maximum possible number of PE for that cluster. This could include one or two channels, and the binning per hit includes the probability of the cluster being hit (in complete tracks, the central regions of the stations are more densely populated). An example based on real data is shown in Figure 3.24. The clear distinction between regions above and below approximately 20 PE is due to the inclusion of two-channel clusters. The sharp drop shows the difference between high and low gain cassettes, which saturate between 10 and 20 PE.

As this distribution represents the probability of a hit being accepted at a given light yield, in order to correct for the saturation effects the measured light yield histogram is divided by this probability distribution. The effect of this correction is included in the light yield plots for each tracker shown later.

3.8.2 Results

Position by Station

The reconstructed space points are shown for each station in Figure 3.25 for tracker 1. Space points that are not used in a track are considered noise and discussed later. For each station, the reconstruction coverage is evident, proving the viability of the tracker active regions.

Triplet quality without cuts

The quality of a triplet space point depends on the sum of its constituent channel numbers. It is expected that the sum will peak around the sum of the central fibre

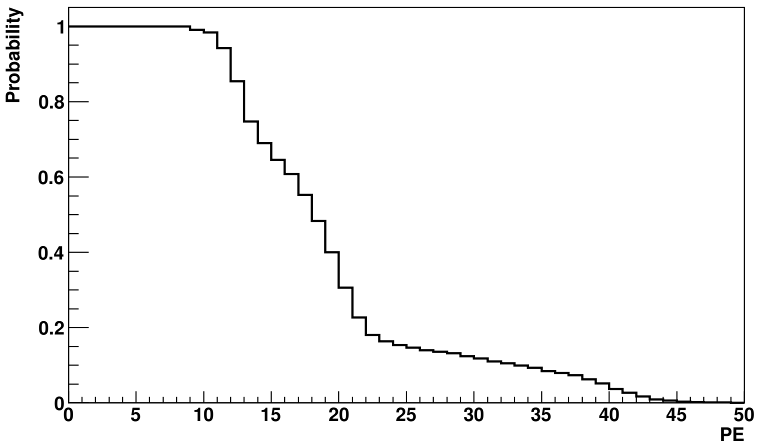

Figure 3.24: Tracker 1 non-saturation probability, showing as a function of light yield, the probability that a hit in a cluster having this light yield had not saturated the ADC. The histogram is also weighted towards the frequently hit clusters. The sharp gradient shows the range of gains for a single channel, while the shallower gradient shows the combined effect of two-channel clusters.

numbers, and will contain tails equivalent to the overlapping of the fibre bundles which will be proportional to the ratio of the overlapping region. The sum of the channel numbers for all triplets used in a track is shown for each station for tracker 1 in Figure 3.26 and for tracker 2 in Figure 3.27.

By limiting this calculation to triplets made of only single-channel clusters, a measurement of the quality of the fibre plane overlapping can be made. This is shown for trackers 1 and 2 in Figures 3.28 and 3.29, respectively. According to the geometry, a perfect overlapping would create equal sums either side of the sum of the central channel numbers, given the symmetric arrangement of the fibres. Any deviation in the mean would signify an internal misalignment of the stations. The interpretation is not a translational misalignment with respect to the tracker coordinate system, but rather the overlapping of the fibre bundles to produce the triangular active regions, and so may represent either a misconstruction of the fibre bundles in each plane, or a misalignment of one plane with respect to another. The offset in the mean indicates the internal misalignment of each station, and thus the offset in each plane is $\delta_{\text{station}}/3$. The measured deviations for each station are given in Table 3.5.

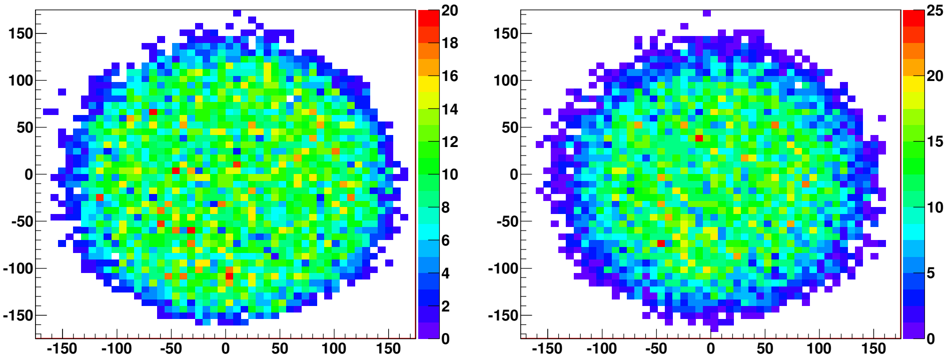

(a) Station 1 Space point x-y position (mm) (b) Station 2 Space point x-y position (mm)

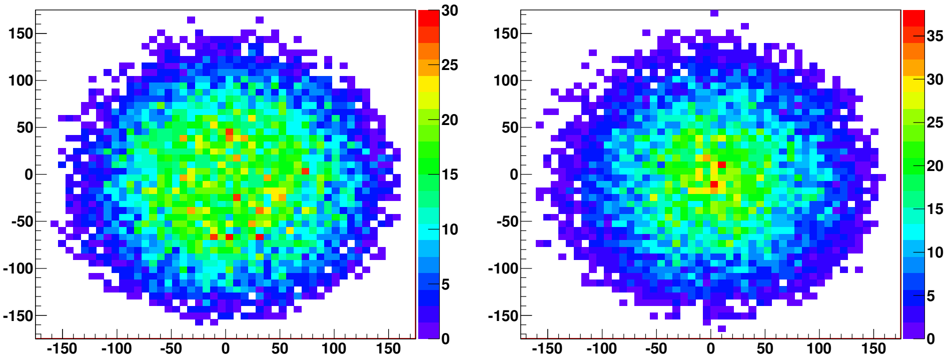

(c) Station 3 Space point x-y position (mm) (d) Station 4 Space point x-y position (mm)

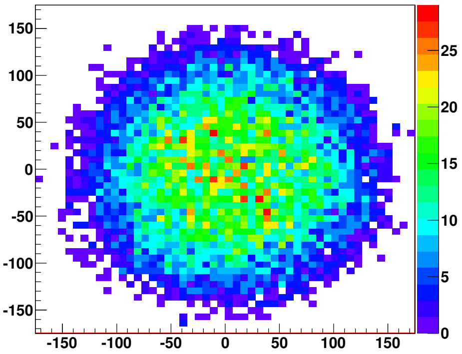

(e) Station 5 Space point x-y position (mm)

Figure 3.25: Tracker 1 space point position by station.

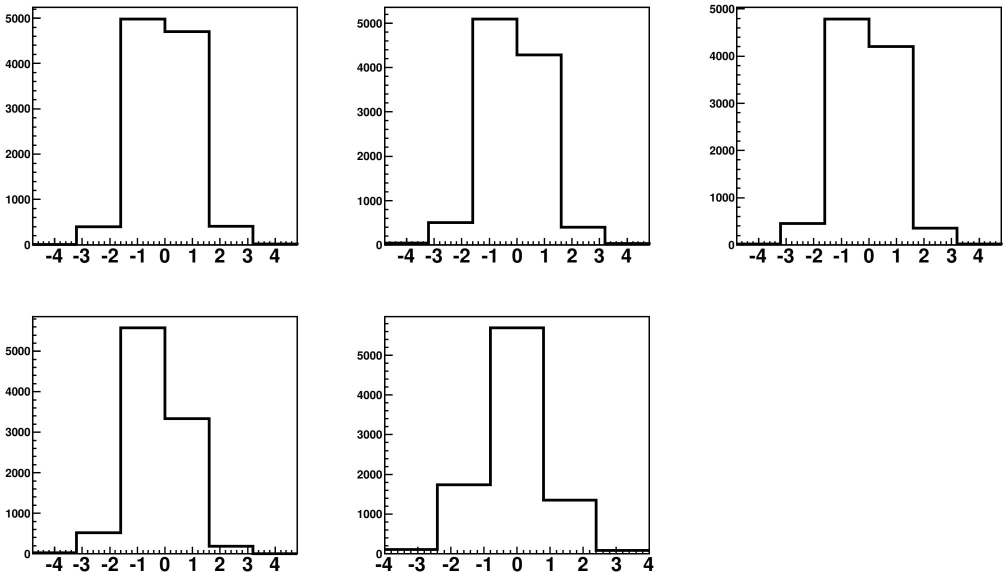

Figure 3.26: Tracker 1 Triplet channel number sum minus the sum of central channel numbers multiplied by the channel width (mm) for stations 1-5 in order.

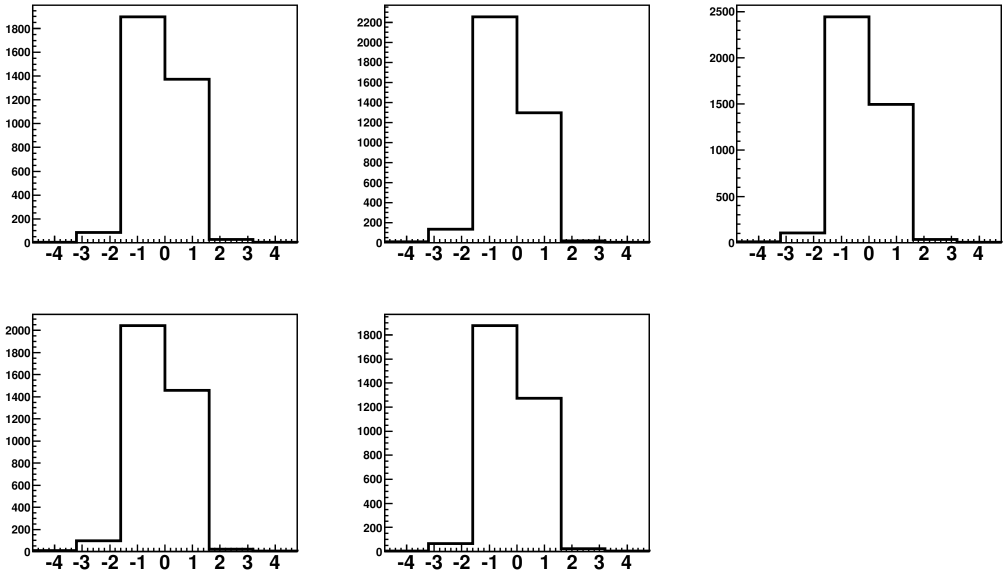

Figure 3.27: Tracker 2 Triplet channel number sum minus the sum of central channel numbers multiplied by the channel width (mm) for stations 1-5 in order.

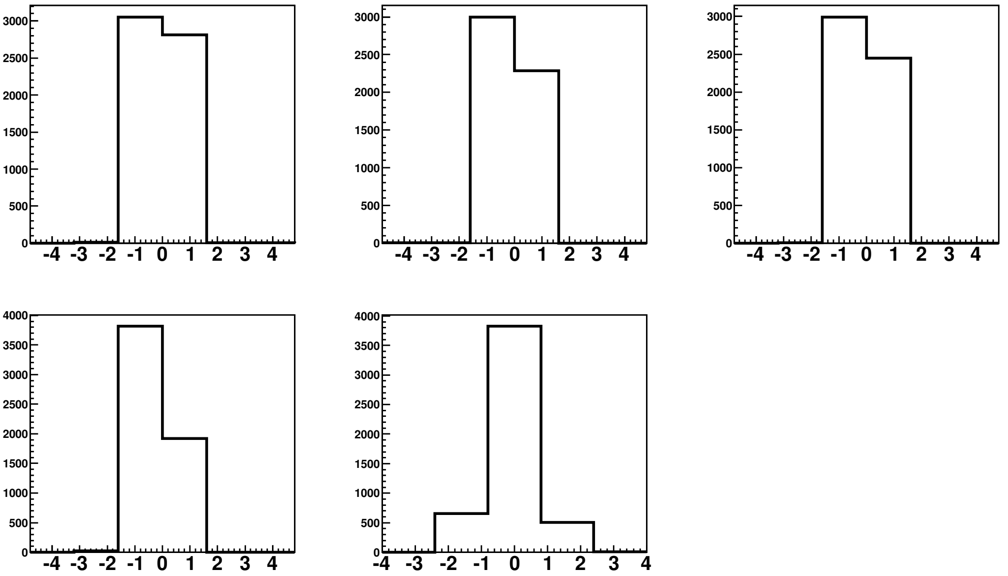

Figure 3.28: Tracker 1 triplet channel number sum minus the sum of central channel numbers multiplied by the channel width (mm) based on single channel clusters for stations 1-5 in order.

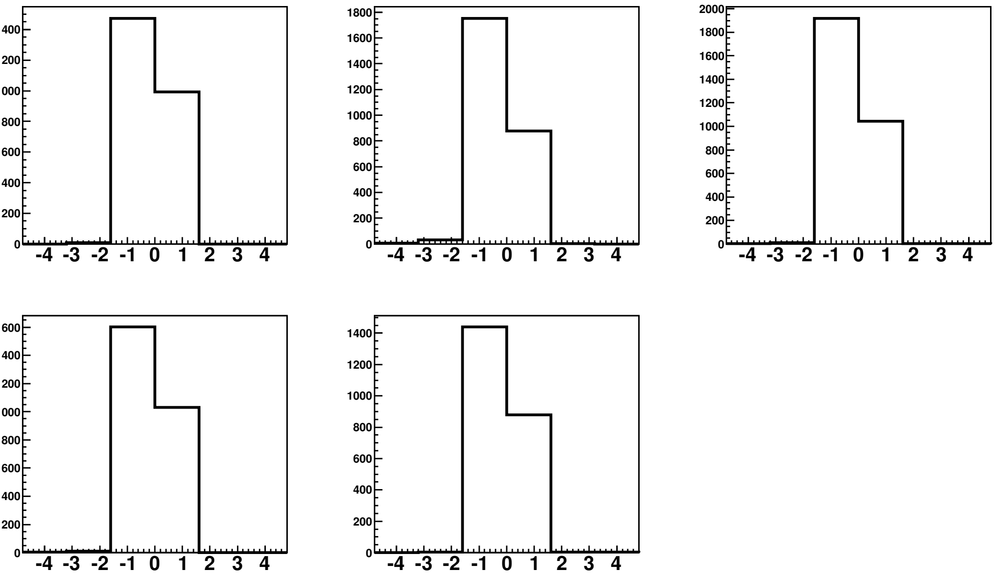

Figure 3.29: Tracker 2 triplet channel number sum minus the sum of central channel numbers multiplied by the channel width (mm) based on single channel clusters for stations 1-5 in order.

Tracker	Station	Offset (mm)
1	1	0.032
1	2	0.114
1	3	0.085
1	4	0.280
1	5	0.047
2	1	0.170
2	2	0.303
2	3	0.248
2	4	0.188
2	5	0.197

Table 3.5: Tracker station internal offsets

Track Residuals

The track residual is defined as the perpendicular distance from a reconstructed track to the centre of the fibres which produced the signal contributing to the track, and where the space point formed from those clusters is not used to make up the track. A five point track will therefore have five sub-tracks, each containing four points. It is the distance between this sub-track and the clusters of the missing fifth point which constitutes the track residual. Given the triangular symmetry of the fibre arrangement, each view is analogous to each other, and is therefore combined into the same histogram, as are the residuals from each station in the trackers.

A Monte Carlo study was performed by generating a beam with a cosmic ray-like momentum spectrum. The track residuals are shown in Figures 3.30 and 3.31, with the real data as points and the simulation as the histogram. By selecting only those space points made from single channel clusters, the plot removes the exaggerated residuals generated by mis-reconstructing the space point in many channel combinations. This occurs when the highest light yield channel is one channel removed from the region which overlaps with the clusters from the other planes, pulling the point from its likely true location (in a two-channel cluster, the centre of the cluster was defined as the centre of the highest light yield channel). This can be improved by weighting the space point's location to the overlapping region rather than the binary charge weighting employed by this reconstruction arrangement.

The sub-millimetre track resolution shown here is sufficient to meet the tracking requirements of the detector, and is likely to be improved with a more complete treatment of the combinatorics of the overlapping fibres.

By extrapolating the track to each station and determining whether or not a space point should have been found, the efficiency of each station can be measured

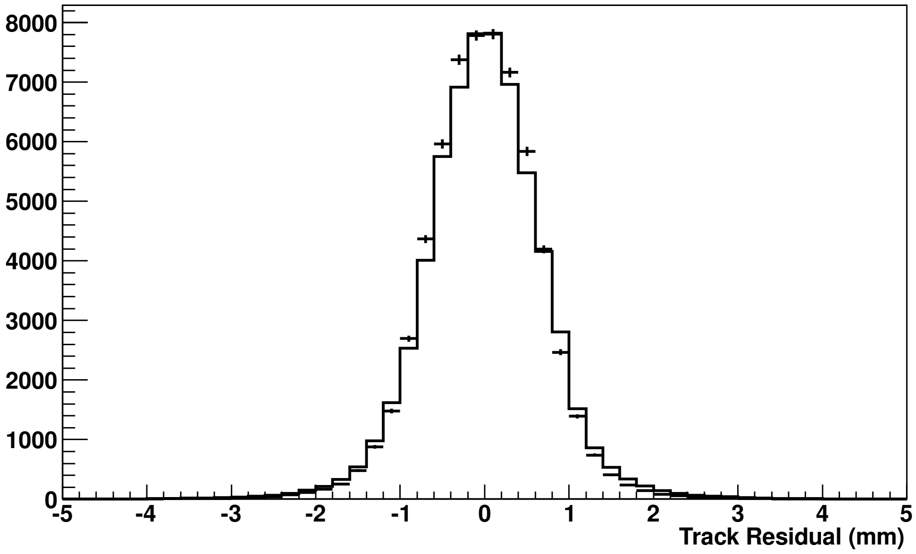

Figure 3.30: Tracker 1 Track Residual. The real data values are shown as points with a $\sigma = 661 \pm 2$ μm and Monte Carlo as the histogram with $\sigma = 718 \pm 1$ μm.

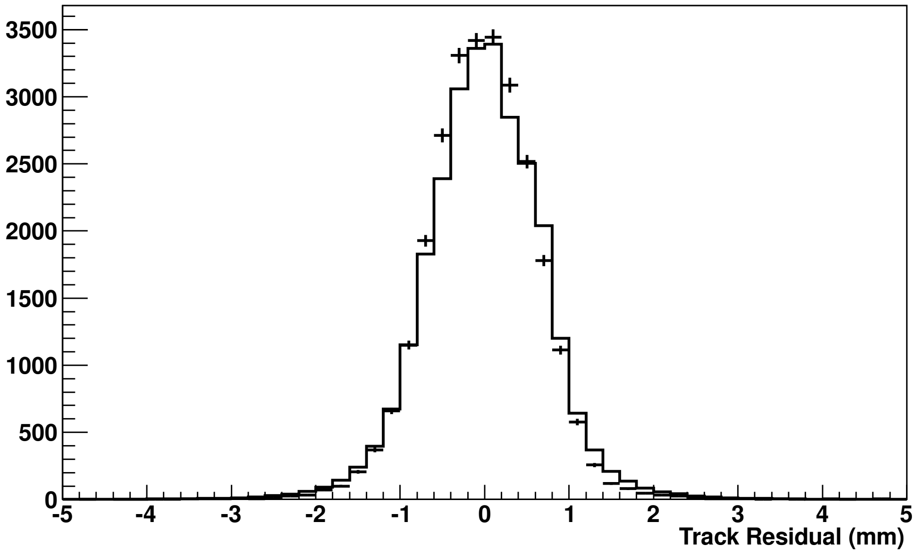

Figure 3.31: Tracker 2 Track Residual. The real data values are shown as points with a $\sigma = 641 \pm 3$ μm and Monte Carlo as the histogram with $\sigma = 706 \pm 1$ μm.

and was found to be above 99% for each tracker.

Light Yields

The light yield for each cluster used in a track (subject to the saturation adjustment described in Section 3.8.1) is shown for trackers 1 and 2 in Figures 3.32 and 3.33, respectively. Light yields are calculated based on the channel-by-channel calibrations and summed for those channels forming a doublet cluster.

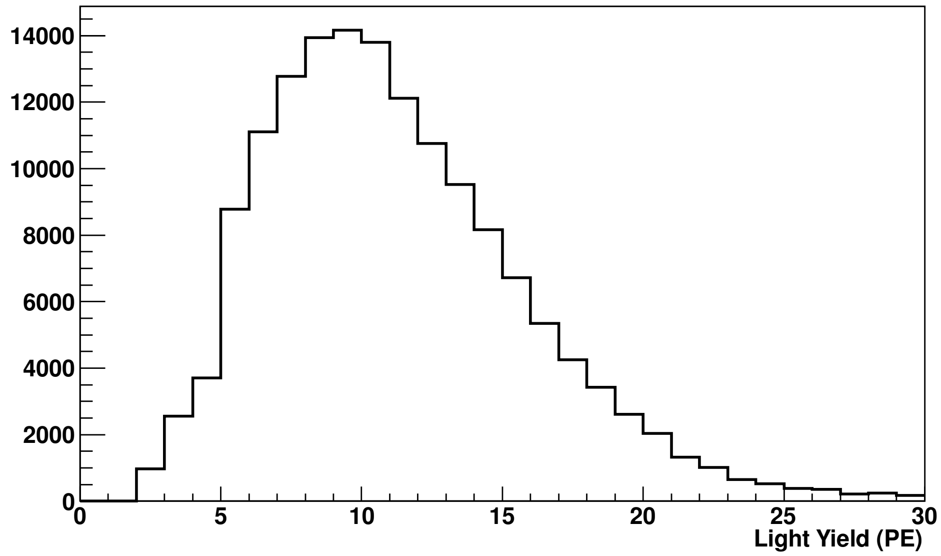

Figure 3.32: Tracker 1 light yield. The mean value is 11.23 ± 0.01 PE.

3.8.3 Analysis of Noise and Cross-Talk

From the analysis of the cosmic ray test, it was found that clusters not corresponding to a track, but nonetheless having signal above the 2.5 PE threshold, were interfering with the reconstruction procedure. The character and source of these clusters were subsequently investigated for optical/electronics noise and cross-talk.

Noise

Initially, there was the possibility that noise could be induced either by an external light source or in the electronics. The incidence of each cluster, plotted by channel as a function of station and plane, is shown for those clusters used in a track, a triplet

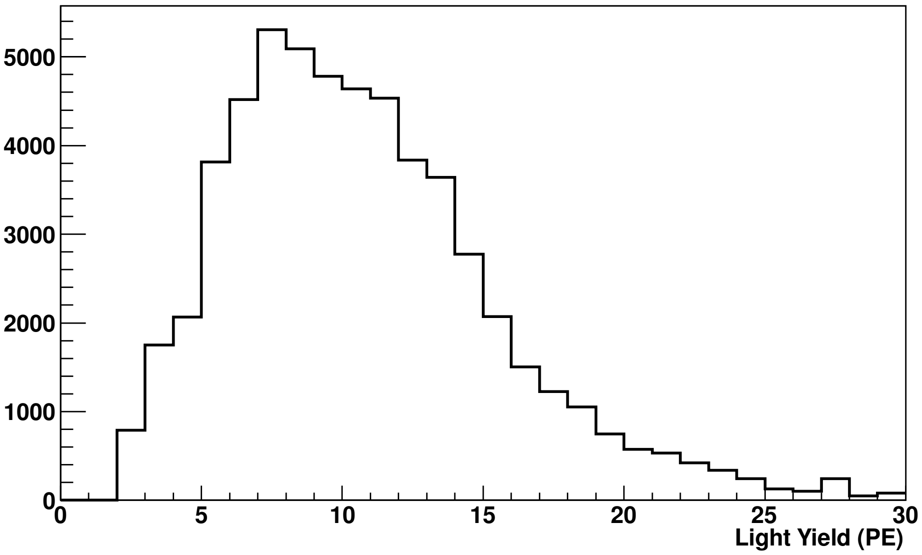

Figure 3.33: Tracker 2 light yield. The mean value is 10.73 ± 0.02 PE.

point, a doublet point and not used at all in Figures 3.34 and 3.35 for Tracker 1. The lack of any large scale abnormal signal in the fibre number space precludes the possibility of a light leak in the tracker body. Observations during the calibration procedure showed that if the light-proofing of the waveguides or VLPC cassette connectors failed, entire modules of channels showed high, above threshold signal for every trigger; this was not the case from the cosmic ray data. However, it is clear that a very small number of channels observe signal at a high rate, whereas all channels (excluding those known to be dead in electronics space) observe at least one instance of having a signal above threshold yet not being used in a track.

Re-binning this histogram in the electronics space (plotting each channel by its electronics channel, board and module, rather than station, plane and fibre channel) will show any correlation between noisy channels in the electronics, and again was not observed on a large scale. It has been seen that abnormally high signal in an electronics module can be due to a failure of the temperature control (although this is spread across a VLPC module and can be confused with a light leak) or a mis-application of bias. From the plots, there is some indication of a pattern between modules, where groups of 128 channels show marginally higher rates of hits than others. This could be attributed to a bias setting which produced a higher dark count rate than neighbouring modules.

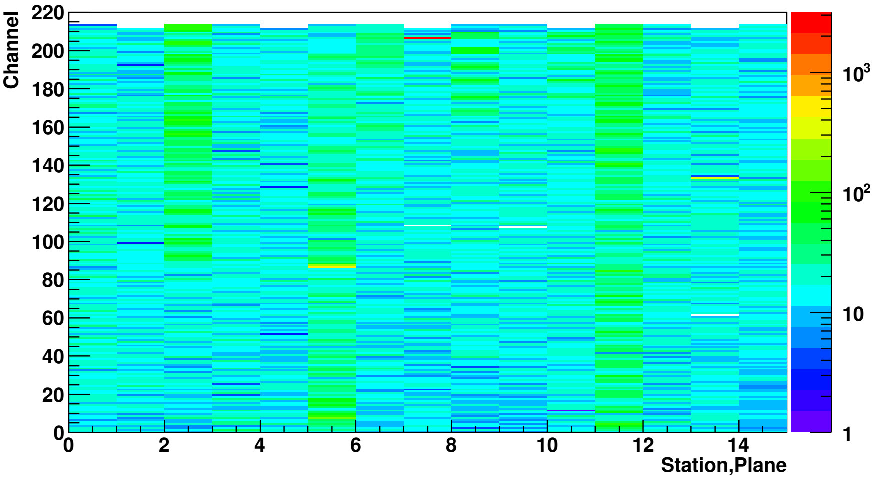

(a) Clusters not used in a space point

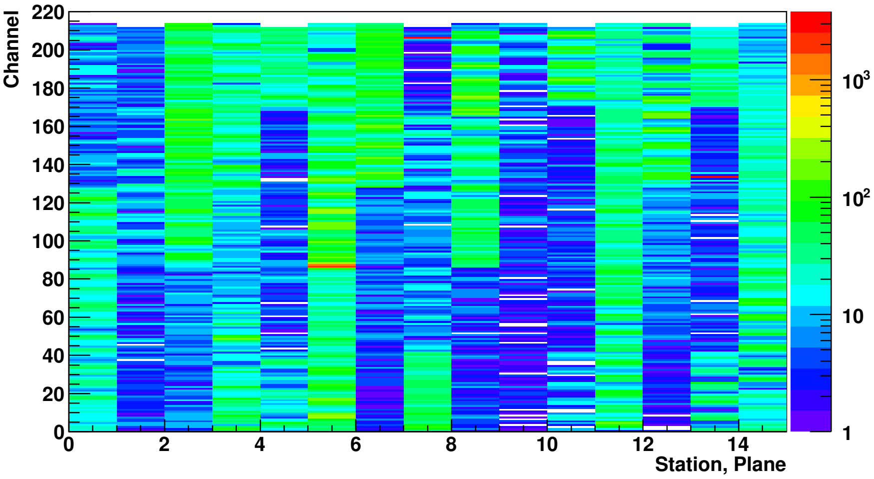

(b) Clusters used to form a double space point

Figure 3.34: Tracker 1 noise indications for single clusters and doublet space points. The X axis shows the plane number from 0 to 14 with three planes in each station. The Y axis shows the fibre number in each plane.

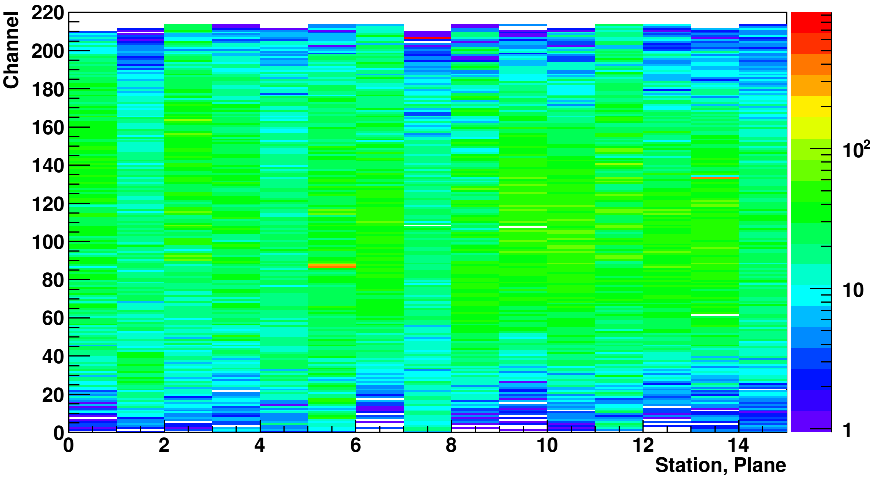

(a) Clusters used to form a triplet space point

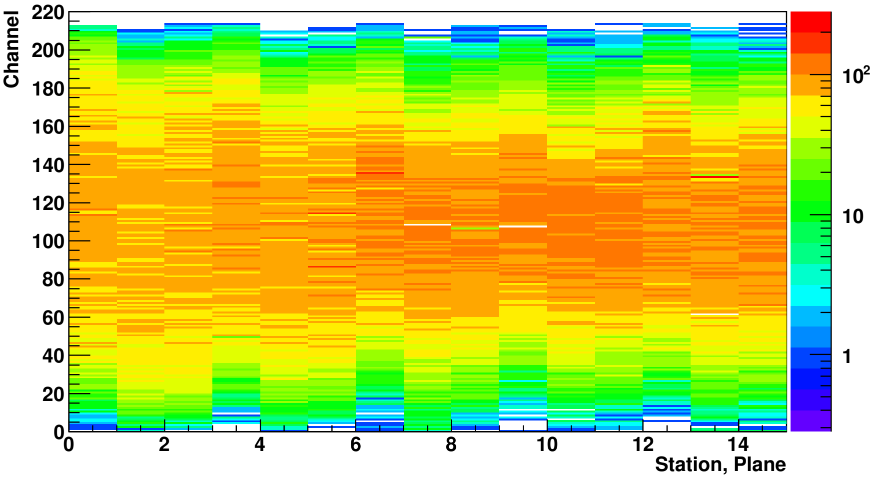

(b) Clusters used to form a track

Figure 3.35: Tracker 1 noise indications for triplet space points and tracks. The X axis shows the plane number from 0 to 14 with three planes in each station. The Y axis shows the fibre number in each plane.

Cross-talk

In addition to the possibility of noisy channels, signal may well be transferred, either optically or electronically, between neighbouring channels. This is termed cross-talk. The possibility of electronics cross-talk has already been discussed in the calibration section. Optical cross-talk would be caused by light leaving one channel and entering a neighbour, and subsequently being propagated to the VLPCs. Given the angle of internal reflection, this is not anticipated.

To investigate the possibility of optical cross-talk, events which had two or more clusters (which themselves can be made from, and are limited to, two channels) adjacent to each other are binned in Figure 3.36. It is clear that there is an incidence of adjacent clusters. The number of channels found to be neighbours in a single event is shown in Figure 3.37, along with the light yield of the cluster. The mean of the light yield distribution for clusters involved in cross-talk is significantly lower than those included in a track. There is some overlap between the channels contributing strongly to noise and cross-talk.

The inclusion of cross-talk-like channels in a track adds uncertainty to the position resolution of that point, in an exaggerated manner similar to that described for two-channel clusters earlier. As the reconstruction algorithm forbids clusters with more than two channels, this effect cannot contribute per se, but can generate clusters with the incorrect channels, thus lowering the efficiency.

Since the original reconstruction method gave no preference to high light yield clusters, or made any requirement on the channel sum of the triplets, the effect on the reconstruction performance was significant (to the level of a few % in the efficiency of each station). However, the modifications to the reconstruction routines nullified the effect of the non-track clusters, and their presence (at their current level) would not interfere with the performance of the trackers in MICE.

When a VLPC enters breakdown, it emits photons which can be received in neighbouring channels. Such channels contribute both high instances of noise on their own account and cross-talk. The calibration procedure described earlier found several channels in breakdown affecting their neighbours in this way. The high individual instances of noise and cross-talk in the cosmic ray tests are therefore attributed to this effect. In order to counteract this process, the connections on the AFE boards to the VLPC cassette were masked for an individual channel in breakdown. This removed the bias applied to that channel and thus made it inert. This nullified both the noise and cross-talk.

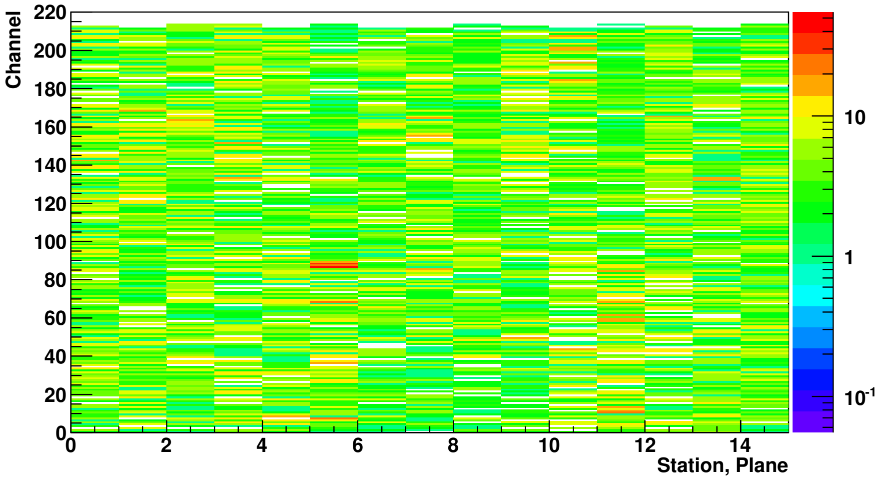

(a) cross-talk clusters not used in a space point

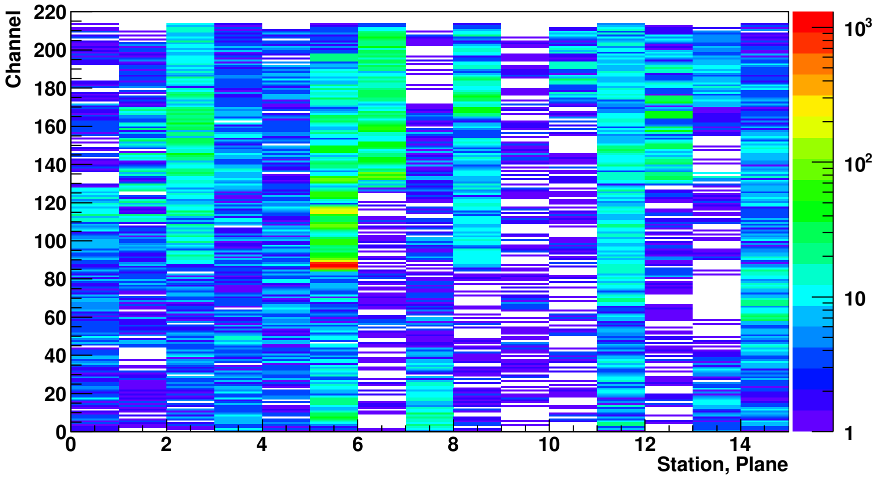

(b) cross-talk clusters used to form a space point

Figure 3.36: Tracker 1 cross-talk indications. The X axis shows the plane number from 0 to 14 with three planes in each station. The Y axis shows the fibre number in each plane.

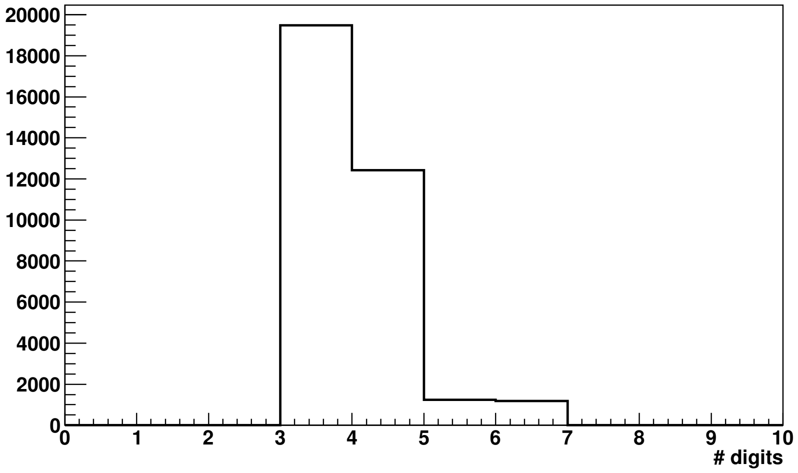

(a) Number of digits found in adjacent clusters

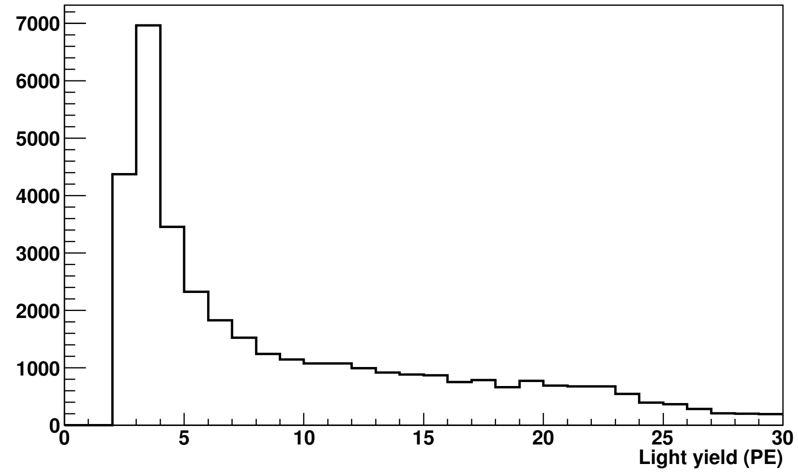

(b) Light yield of clusters considered as cross-talk

Figure 3.37: Tracker 1 cross-talk indications. The mean number of digits in each instance is 4.04 and the mean light yield is 9.4 PE.

3.9 The simulated performance of the MICE trackers

To measure the resolutions of the trackers in each of the phase space coordinates (x, y, P_x, P_y, P_z) a Monte Carlo study was performed, employing the reconstruction methods and track filtering previously described. A beam was generated conforming to the emittance-momentum properties of a MICE operating mode, with its covariance matrix defined by Ref. [43].

3.9.1 Pattern Recognition with Timing

Using the measured time resolution of the TDCs found in Section 3.6.6, and an approximation to the time of flight of a muon through the tracker, a slice in time can be defined during which all track-related signal should reside. The aim of using time information in pattern recognition is primarily to veto noise originating in the MICE RF cavities.

Taking these factors into account, a slice with a variance of 10 ns (an estimation of the muon flight time, rise times in the fibres, cable lengths and TDC resolutions) was chosen. Any space points from different stations (limited to one point per station), which as a set had an RMS of less than 10 ns, were combined into a pattern object and used to generate a seed. Using this method in a full simulation of the MICE trackers (including muon decay and other physics processes), the efficiency of using timing alone was measured around 90%.

The seed state is defined by taking three points and calculating the centre and radius of the circle they form. The longitudinal component of the helix, corresponding to Pz, is found by the angle ϕ advance between the known dz of the stations.

3.9.2 Resolutions

The helical tracking was tested using an input beam with a transverse emittance of 6 mm, β function of 333 mm and total momentum of 200 MeV/c. Using the aforementioned pattern recognition and Kalman filter, the $x - y$ residual between the track state at a station and the Monte Carlo truth at that station is shown in Figures 3.38 and 3.39 for tracker 1, with all stations combined into the same plot. This is in fair agreement with the design specifications and initial simulations [76], with some variations expected given different approaches to defining the space point resolutions.

The MC - Reconstructed longitudinal momentum residual was also calculated with an RMS of 5.8 MeV/c and a mean offset of 1.4 MeV/c. The mean offset is

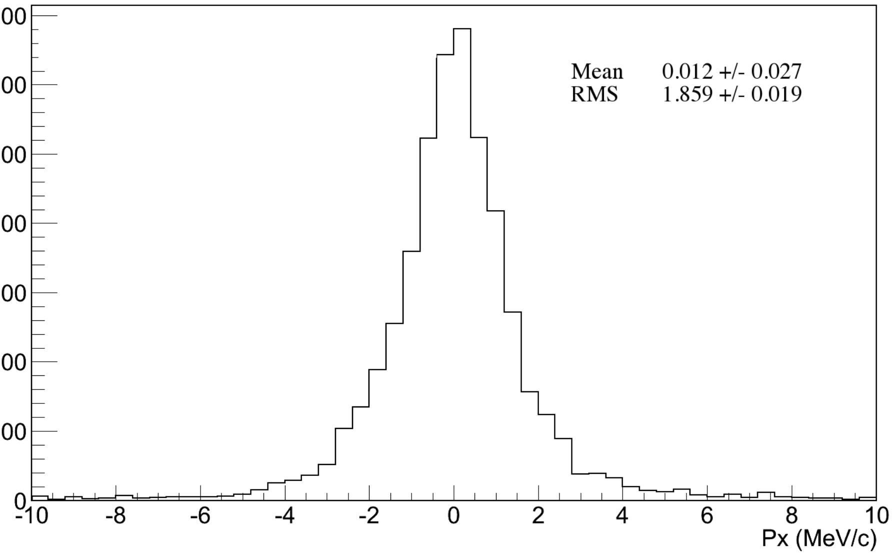

Figure 3.38: Tracker 1 x momentum residual.

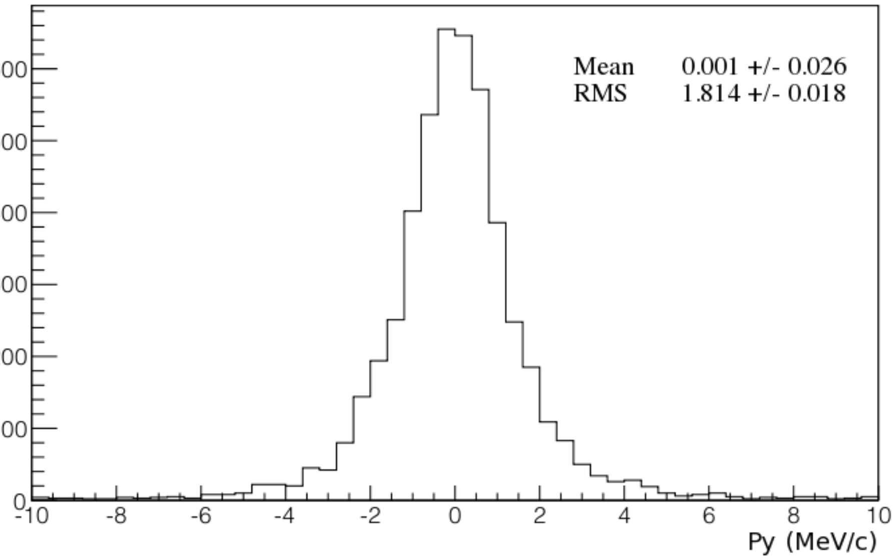

Figure 3.39: Tracker 1 y momentum residual.

accounted for by the energy losses through the tracker stations not being taken into account in the simple track fit.

3.10 Conclusions

The purpose of the MICE fibre trackers is to provide a high efficiency, high granularity detector which, in combination with a well defined solenoidal field, can provide an accurate measurement of particle momenta. During a cosmic ray test, the efficiency and space point resolution has been determined and shown to be reproduced by the simulation. A calibration procedure has been developed and shown to optimise the fibre light yield whilst minimising noise, which was an issue during the cosmic ray test using a different calibration algorithm. Overly biasing the modules produced noise which in some instances contributed to electronic cross-talk. The calibration algorithm which minimises the noise is equivalent to optimising the efficiency, and so solves this issue in both respects.

Modifications to the reconstruction algorithms have allowed for the rejection of fake space points, which, in conjunction with time-based pattern recognition, will be essential in compensating for the expected RF induced noise in latter MICE stages. As part of this modification process, the differing performance of the stations was analysed which showed the appearance of minor misalignments. With the minor misalignments, noise and cross-talk analysed from the cosmic ray test, the as-built features of the trackers are understood, as is their contribution to the overall performance.

The behaviour of the front end electronics is now well understood along with the required integration of the data acquisition systems into the MICE infrastructure. The speed of the DAQ is predicted to be sufficient for high particle rates once the zero-suppression has been implemented.

The calibrated time circuit has been employed in a pattern recognition role, and can contribute strongly towards maximising the muon content of the tracks passed to the filtering algorithm. Track finding using a Kalman filter has been written and used to show the $x - y$ momentum resolution is consistent with design requirements. Therefore, the completed and tested trackers are expected to meet the requirements for a precision emittance measurement in MICE.

Chapter 4

Particle weighting with Voronoi Diagrams

4.1 Introduction and Motivation

Due to the mechanical constraints of the beamline, MICE data has a limited number of control parameters, and is subject to often large-scale variations in the source particles which are either undesirable or detrimental to analysis. Given the lack of a horn or other focusing of the particles produced from the target, or bunch rotation or chopping, the MICE beam varies greatly from that which would be expected to enter the cooling channel of a full neutrino factory. This generates a number of reasons why weighting is both necessary and useful:

- **Particle Time versus RF Phase** - The population of the MICE spill gate (itself governed by the period during which the target scrapes the ISIS beam and DAQ limits) is neither uniform nor quantised. Whilst the density of triggers increases with time, the triggers themselves could occur at any time. Conversely, the phase of the MICE RF cavities will be set once per spill. As the particles arrive at the RF cavities with their wide time distribution they will fill the entire period of the RF wave; some particles will be on crest as desired, but others will experience the decelerating parts of the wave. Whilst this is advantageous in longitudinal bunching, MICE aims to measure maximum longitudinal re-acceleration, and so only particles on crest are desired; this can be achieved by either selection or weighting.
- **Mismatched Beam** - In order to achieve maximum emittance reduction the betatron function must be minimised during the period of interaction with the

absorber; to achieve this, the MICE beamline is designed to match the betatron function in the absorber and the trackers. If the matching is incorrect, or the state of the real MICE beam such that matching is difficult to achieve, the maximum emittance reduction will not be observed (heating may even occur). In this case, each particle could be weighted such that the full weighted beam does observe maximum emittance reduction (or rather, is in such a state where maximum emittance reduction is expected).

- **Missing Data Points** - The range of emittance-momentum data points from the MICE beam is largely governed by the initial state of the beam (for the minimum achievable initial emittance) and the absorber (the corresponding maximum). It would be desirable to weight the particles such that the total beam distribution can then fill intermediate points. As long as the relevant phase space is adequately occupied (absolutely, rather than in relation to other spaces) the beam could be weighted to appear like any well defined distribution.

A method has been developed using Voronoi diagrams [78] based on the phase space occupation of the beam to statistically weight each particle. The effectiveness of using Voronoi diagrams to weight particle beams is shown and the quality and efficiency of a number of algorithms tested. The weighting method has been used to solve a selection of problems facing MICE, such as the incongruity of the beam time distribution and RF phase. The method can also be employed to use real data to model the effects of different input beams (such as varying emittance).

4.2 Voronoi Diagrams

A Voronoi diagram is defined by dividing an N-dimensional region of points into cells, with one cell per point, in which each cell contains that volume of space closer to the enclosing point than any other. It is closely related to Delaunay tessellation [79], since the vertices of a Voronoi cell are formed from the centroids of the Delaunay cells. An example of a 2-D Voronoi diagram can be seen in Figure 4.1.

4.2.1 Use as a Weighting Algorithm

It has been shown that Voronoi diagrams can be used to apply a statistical weight to a particle based on its phase space location [80]. The weighting method is best illustrated with the one dimensional case in which no approximations are involved. The one dimensional Voronoi diagram can be formed by taking a 1D histogram and

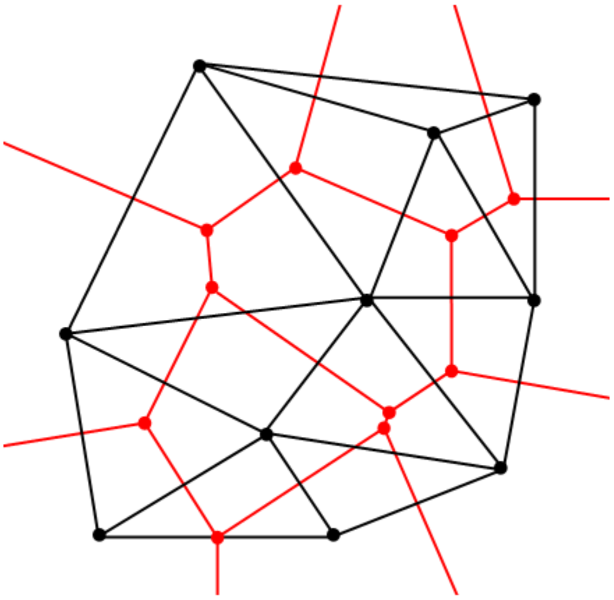

Figure 4.1: A 2-Dimensional Voronoi Diagram with Delaunay tessellation. With the black points joined together to form the Delauny triangles, the red points indicate the corners of the Voronoi cell.

varying the bin widths until each bin contains only one entry. The height of each bin then becomes $1/dx$, where dx is the bin width. If a probability density function (PDF) is overlaid on the histogram it can be seen that each bin can be increased to match the PDF $f(x)$ by multiplying by a factor w, which is calculable as shown for a continuous case in Equation 4.1, from which the discrete case follows. This is illustrated in Figure 4.2.

$$w = \int f(x)dx \tag{4.1}$$

Integration of a target PDF across the region of a Voronoi cell generates the weight which, if applied to the point contained within that cell, will transform a discrete distribution of points into one possessing the target probability density function. This remains true for higher dimensions, although the problem of integrating over the regions becomes more complicated. From Figure 4.2, the weight is the integral of the PDF line over each bin, applied to the particle in that bin.

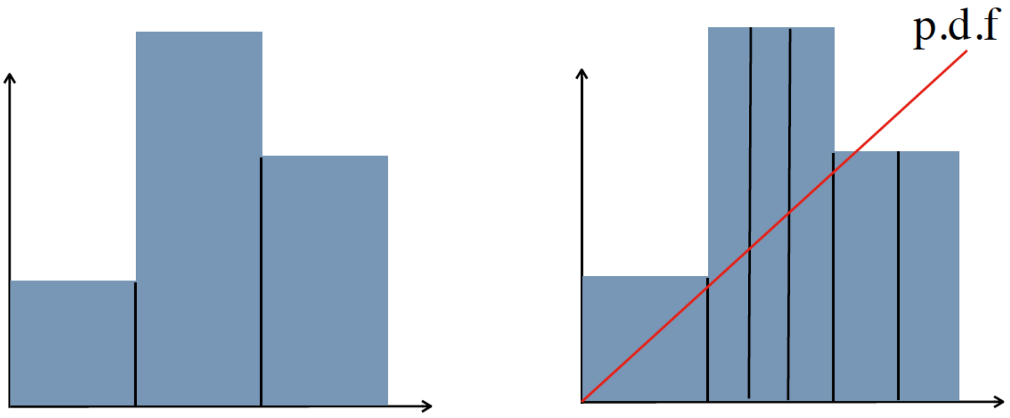

Figure 4.2: A 1D Voronoi diagram with a target p.d.f. overlaid. The right-hand figure shows how each bin would be divided to achieve the bin value of one, which could then be integrated to the overlaid p.d.f.

4.2.2 Methods of generating a Voronoi Diagram

A varied selection of algorithms have been developed for efficiently generating Voronoi diagrams:

- **Bowyer-Watson Algorithm**
  The Bowyer-Watson algorithm is an incremental method which uses the circumcircles of triangles to accommodate each new point. Upon insertion, a search is made to find these triangles. Any vertices of the triangles within the circumcircles are removed and new vertices added from the new point to the points on the triangles. This algorithm has the possibility of generating degenerate sets, and has complexity $\mathcal{O}(n \ \mathrm{ln(n)})$.

- **Fortune's Algorithm**
  Fortune's algorithm [81] uses a sweeping line method which moves across a set of points and draws the Voronoi diagram as it passes them. The running time has a worst case of $\mathcal{O}(n \ \mathrm{ln(n)})$ requiring $\mathcal{O}(n)$ space.

4.3 Approximate Voronoi Diagrams

In order to reduce the complexity of the algorithm, owing to its inherent processing time and the memory required to store all the vertices for an N-dimensional set of points, approximate methods of producing a Voronoi diagram have been developed. These consist of analytical methods of subdividing the region between points.

4.3.1 Approximate Diagrams by Sub-division

A fast analytical approach involves bisecting the global region between two points, and then repeating in the newly generated sub-regions for all subsequent points. A simplified example can be seen in Figure 4.3. The algorithm operates as follows:

1. **Create Major Region(s)** - The initial regions should contain the full set of points from which the Voronoi diagram will be drawn. The boundaries for this region are obtained depending on the criteria of the analysis. For example, a diagram drawn from the points in phase space found in the MICE trackers might take the tracker solenoid fiducial volume as the starting global region. The region sizes are largely inconsequential to the calculated weights, except for those points at the physical edge of the distribution. It is also possible to define several starting regions.
2. **Insert initial point** - Since this method draws the diagram around each point in turn, the order of the points is of concern. The effect is explored later.
3. **Insert additional point**
4. **Find region containing two points** - This method of drawing a Voronoi diagram naturally produces a tree which can be used to find the containing sub-region.
5. **Bisect the region** - Once the region containing both the new point and an old point is found, it is divided into two new regions, each containing a single point. The distance between the points is found and a mid-point used as the dividing line between the two new regions.
6. **Store new regions** - The new regions, along with pointers to their parent, are stored and form part of the diagram tree.
7. **Repeat steps 3-6 for all remaining points**
8. **Strip the outer points** - The cells whose borders are not provided by other cells, but by the initial major regions, will have a size dependent on that of the major region. This could be arbitrary and may not reflect the true size of the cell (if this could be defined using this method, which it cannot). All cells with a border on that of the major region could then be rejected.

9. **Remove high weights** - As an alternative to stripping the outer regions a cut can be made on the calculated weights, excluding the high outlying values.

Upon completion of the algorithm an approximate Voronoi diagram is drawn along with a tree for nearest neighbour searches. For weighting, the limits of the integration are the only pieces of information required, and thus are the only pieces produced in the generation and searching of the cells.

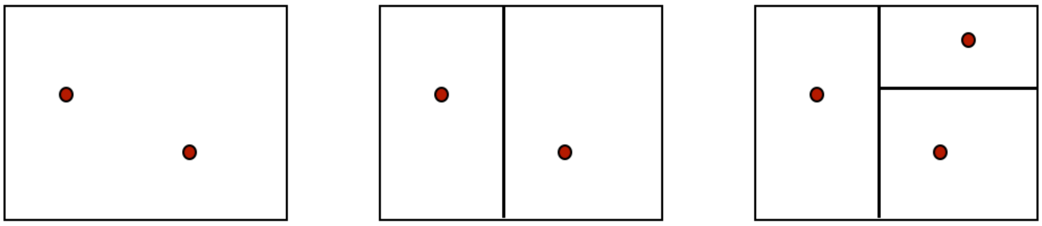

Figure 4.3: Creating an approximate Voronoi diagram.

The total amount of information stored for each region, measured in floating point numbers, in this form of the Voronoi diagram is then $2N_D$, where N_D is the number of dimensions, with smaller items such as integer indices and status bits. The total number of regions produced, including the intermediate steps (for example, the regions in the second step in Figure 4.3) is $2n - 1$, where n is the number of points.

Search methods

One particular area of inefficiency lies in step 4, in which the existing region in which a new point lies must be found. As each region gets divided a new level is created in the tree structure of the regions, with the initial region at the top and most recently created regions at the bottom. As the number of lowest level regions increases with the points, the time spent searching for the region increases geometrically. The simplest brute force method is to loop over all existing regions, and if they are of the lowest level test to see if the new point lies within them.

For a limited set of points the efficiency of this method is comparable to other areas of the algorithm. However, for larger data sets, the time spent searching through every possible region, which themselves scale as $2n - 1$, may become untenable. The scaling in time taken for a range of points and dimensions is given in Table 4.1. The memory required follows $2N_D(2n - 1)$. For the analyses presented here, and for the majority of MICE running, the number of events is within the acceptable time spent on the region search.

Points	1	2	4	6
10^2	0.4	0.4	0.4	0.4
10^3	0.4	0.4	0.4	0.4
10^4	2.4	2.4	1.8	1.3

Table 4.1: Times taken in seconds to generate Voronoi diagrams for a range of dimensions and numbers of points, based on a single processor core on a virtual machine. The host processor was a 2.66 GHz dual core with hyper-threading.

At the rates expected in MICE, and the total number of particles required for emittance reduction measurements, a brute force method is acceptable for both online and offline running.

4.4 Weighting

Once the approximate Voronoi diagram has been produced, the weight for each point is calculated by integrating a target function over the containing cell. In the case of a multivariate Gaussian, this is only achievable analytically if the covariance matrix is diagonal (it will be shown it is possible to first diagonalise the covariance matrix and draw the Voronoi diagram in the diagonalised reference frame). In the case of a mono-variate Gaussian, the integration simply takes the form of calculating the Error function, which can be calculated computationally as a Taylor expansion. For a multivariate, but diagonal, Gaussian distribution, the weight is the product of the Error functions in each dimension.

4.4.1 Testing

A selection of tests were performed in order to verify the method and measure the quality and effectiveness of the weights . The transverse beam distribution is almost entirely a multi-variate Gaussian distribution, as is the longitudinal component. In addition, the sinusoidal nature of the accelerating RF wave gives reason to test the viability of using trigonometric functions as the target probability distribution.

Starting with a bi-variate Gaussian distribution without covariance terms, as shown in Figure 4.4, the effect of the weighting can be seen in the variances shown in the plots. In this case, the accuracy of the method decreases with the number of dimensions. It is also possible to weight towards functions such as a sine wave, shown in Figure 4.5.

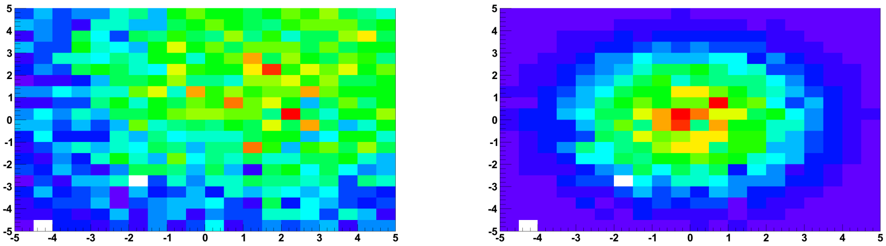

(a) Unweighted distribution (b) Weighted distribution

Figure 4.4: Weighted Gaussian distributions. The unweighted distribution is shown with a mean of 2 and a σ of 4.5. A Voronoi diagram is drawn and each cell weighted according to a second bi-variate, diagonal Gaussian distribution, with a σ_x of 2 and a mean of 0. The means and σ_T of the weighted distribution are 0 and 1.9 respectively.

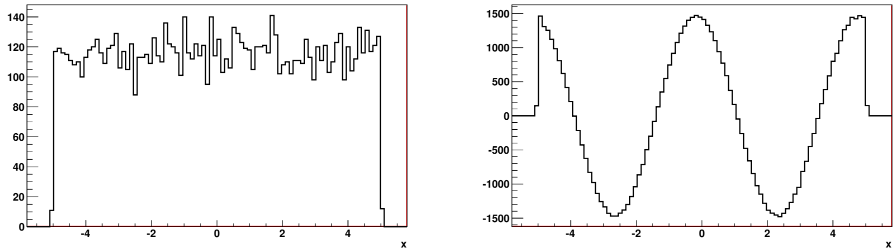

(a) Unweighted distribution (b) Weighted distribution

Figure 4.5: Example weighting from a uniform distribution to a sine function. The target frequency in the weighted case was 0.2 Hz.

One source of systematic error that could occur in the calculation of the weights is caused by the order in which the diagram is drawn. As the points are added to the diagram one at a time, there arises the possibility that a different ordering will create different cell volumes and limits, thereby generating different weights for the same point. The effect of randomising the order of the particles can be seen in Figure 4.6, where 10000 particles have been randomised in order 50 times and their weights calculated. Whilst the absolute variation of the weight is small, the low values of the weight itself can have a large fractional variation, as high as 50%. This will not affect the result of a weighted analysis, but will control how well the weighted set resembles a target PDF.

4.4.2 N-Dimensional Gaussian

In the case in which the target function has a non-diagonal covariance matrix, it becomes necessary to first rotate the reference frame of the points into a basis in which the matrix is diagonal. This is achieved by first diagonalising the covariance matrix and then using the diagonalising matrix to rotate the set. This process is defined as follows:

1. **Diagonalise the Target Covariance Matrix** - By first calculating the eigenvalues and eigenvectors of the covariance matrix, a matrix P is formed from the column of eigenvectors. A diagonalised form the covariance matrix is then found by $D = P^{-1}AP$. As a real, symmetric matrix, a Gaussian covariance matrix is always diagonalisable.

2. **Rotate the basis of the points** - Using the previously found matrix P, each point x can be rotated into the new basis as $x_d = Px$.

3. **Draw the Voronoi diagram** - An approximate Voronoi diagram is then drawn in the new reference frame.

4. **Weight each cell** - Using the diagonalised covariance matrix, the weight now becomes the integral sum of univariate Gaussians, using the eigenvalues of the original covariance matrix as the variances.

5. **Re-draw** - Once weighted, the distribution can be redrawn in the original basis to monitor the effects of the weighting.

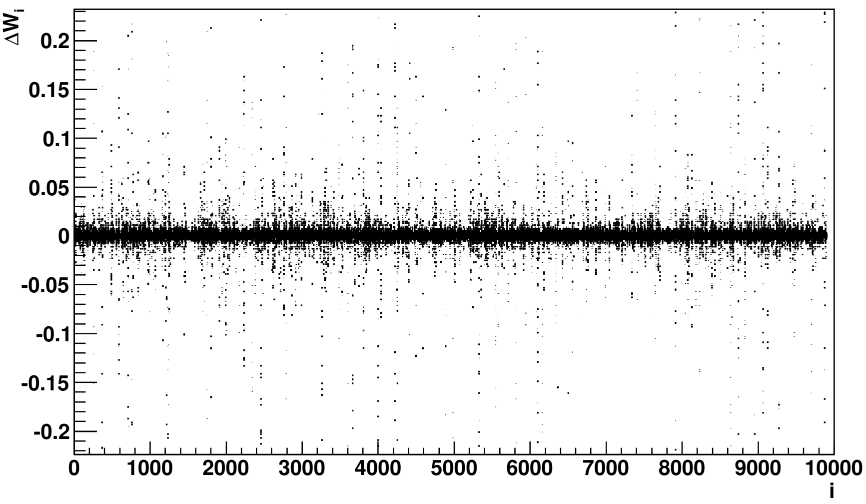

(a) Change in weight by Voronoi ordering.

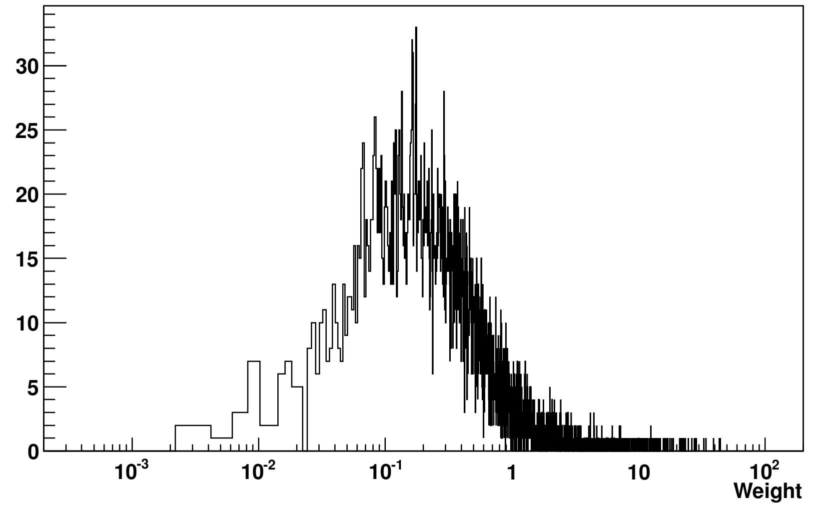

(b) Absolute value of weights for a sample distribution.

Figure 4.6: Variation in the calculated weight based on randomising the order in which the diagram is drawn.

4.5 Weighting Analysis Method

The analyses presented in this section are based on a simulation of MICE steps IV and V. Unless otherwise stated, all data points are taken from Monte Carlo truth values at the tracker reference planes, which are within the trackers at points closest to the cooling channel. The beams are generated with a covariance matrix defined by Equation 1.27 but have no time dependance unless stated (such as the Step V simulation). The longitudinal momentum spread is also kept minimal. A particle is required to induce hits in both the virtual detectors producing the MC information, and be likely to produce two five point tracks in the trackers, although no reconstruction is performed to verify this to avoid introducing reconstruction biases. If the preceding cuts affect the input emittance, a cut on the particle amplitude is made to ensure the emittance at the first tracker reference plane is not greater than desired. Due to scraping, the emittance may be lower than expected.

The measurement is then defined as $(\epsilon_1 - \epsilon_2)/\epsilon_1$, where ϵ_1 and ϵ_2 are measurements of emittance in the first and second tracker reference planes, respectively. A reduction in emittance is then given as a positive value.

4.5.1 Error analysis

With the use of weighted points additional care must be given to the calculation of the measurement errors. Since the MICE data consists of measurements in two detectors of the same ensemble of particles, the initial emittance measurements are correlated. This leads to an adjustment of the standard statistical error by a correlation factor which leads to a reduction in the expected statistical error [82]. Due to the x and y planes being treated as independent measurements, the error scales as

$$\sigma_f \propto \sqrt{\frac{1}{2N}} \tag{4.2}$$

where N is the number of particles, and would be adjusted by a term corresponding to stochastic effects. The calculation of the weights is normalised such that the sum of the weights is equal to the number of particles. However, with selection criteria that could remove some weighted points, this statement no longer holds and the statistical error must be adjusted such that $N = \sum w_i$, as given by

$$\sigma_f \propto \sqrt{\frac{1}{2\sum w_i}} \tag{4.3}$$

Due to the correlations, it is likely this error is an overestimate. Whilst this represents the statistical error on each individual point, ignoring detector effects, when combining many weighted points together their correlative effects must be considered. As each weighted set is made from the same particles, albeit with different weights, there is a strong correlation between them, which is maximised by sets which have been weighted to similar PDFs. This can have two effects: firstly, when the true ensemble of points suffers no severe statistical variations, the weighted sets will follow one another and the fit will have a very low χ^2, giving the impression that the errors have been over estimated (although it is difficult to separate this from the measurement correlations). Secondly, fluctuations in the ensemble will be expressed by neighbouring points due to their similar weighting and give the impression the errors have been underestimated or there is some other effect involved. Examples of the former and latter effect can be seen in Figures A.1 and A.3, respectively.

4.6 Weighting Due to RF Phase

The RF in MICE is designed to operate with a phase of 90° (on crest). This is due to power constraints in the MICE hall, and so the bucket is timed such that particles receive the maximum acceleration (the peak of the wave coincides with the arrival of the particle in the cavity). This will partly make up for the difference between the 8 MV/m MICE RF and a Neutrino Factory RF system.

The phase of the RF will be set once per spill with respect to the ISIS RF phase. During each millisecond spill, there will be approximately 3000 bunch interactions with the target, with each bunch being 100 ns long. This 100 ns bunch is further inflated by material effects, momentum deviations of the original pions, path lengths through the beamline optics and momentum differences caused by the pion decays.

The summation of these effects is likely to produce a muon beam with a time spread larger than the RF bucket. Particles which lie outside of an RF bucket are subject to inherent instabilities and will move further from the target phase of 90°, inflating the longitudinal emittance. As the goal of MICE is to measure sustainable emittance reduction, particles which fall out of the RF bucket should not be included in the measurement of emittance reduction, as they will eventually leave the beam completely.

The total energy change for a set of particles passing through Step V can be seen in Figure 4.7. The beam was uniformly populated in time, allowing particles to arrive across the full period of the RF wave. Whilst the true MICE beam will differ in the time structure to this example, the effect on particles outside the RF bucket is highlighted. The corresponding effect on the longitudinal momentum is shown in Figure 4.8. There is a clear relationship between the time of the particle at the reference plane and the acceleration it receives.

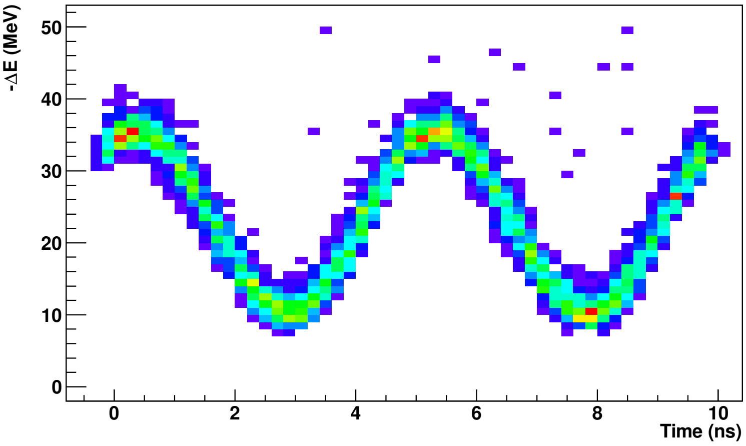

Figure 4.7: Loss in energy for particles passing through Step V, based on a 6 mm, 200 MeV/c beam with a uniform time distribution of 10 ns. The effect of the 201 MHz RF can be seen.

A number of options based on weighting or selection criteria could be used to mitigate this effect:

- **Selection requirement on longitudinal momentum change** - As the purpose is to maintain the longitudinal momentum of the muons, a simple selection requirement can be made on the change in longitudinal momentum through the absorbers, requiring that only particles with less than a certain amount of longitudinal energy loss are included in the emittance calculation. However, this will be shown in Section 4.6.1 to introduce a bias to the emittance loss measurement.

- **Remove particles outside the RF bucket** - Once the time of the particle with respect to the RF bucket is known, those particles which fall outside can be removed. Given the ISIS particle bunch will appear almost uniform at the scale of the MICE RF, this could create an artificial bunch structure with unrealistic properties.

- **Weight particles based on their phase** - Given the time of the particle in the RF wave, a weight can be applied to bias the measurement towards those

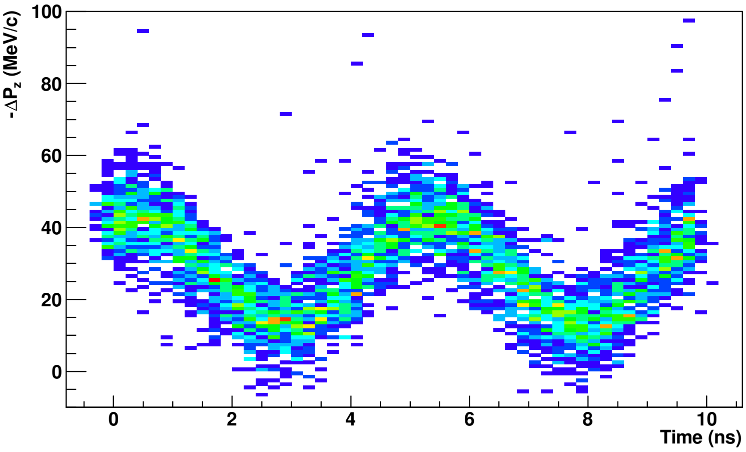

Figure 4.8: Loss in longitudinal momentum for particles passing through Step V, based on a 6 mm, 200 MeV/c beam with a uniform time distribution of 10 ns.

particles within the RF bucket, without excluding those particles which will not be optimally accelerated.

4.6.1 ΔPz Selection requirement

As the desired measurement of emittance reduction is made on those particles which sustain longitudinal momentum, it might be desirable to cut on the longitudinal losses. However, given the expected level of energy losses due to $\langle \frac{dE}{dX} \rangle$, any minimisation of the P_z losses will be transferred to the P_T losses, thus increasing the observed transverse emittance loss, and generating a bias. The changes in P_z and P_T are thus correlated.

This is best observed in the absence of the RF cavities. The effect of such a cut on the change in longitudinal momentum, $\Delta P_z = P_{z(in)} - P_{z(out)}$, is shown in Figure 4.9 for a 6 mm 200 MeV/c beam passing through Step IV.

There is a clear correlation between the cut and the increase in transverse emittance loss, which is interpreted as selecting those particles where the ratio of the energy loss to the multiple scattering (caused by random processes) to be high. Any arbitrary selection in this manner, which does not depend on some knowledge of the particle relationship with the RF wave, must be rejected.

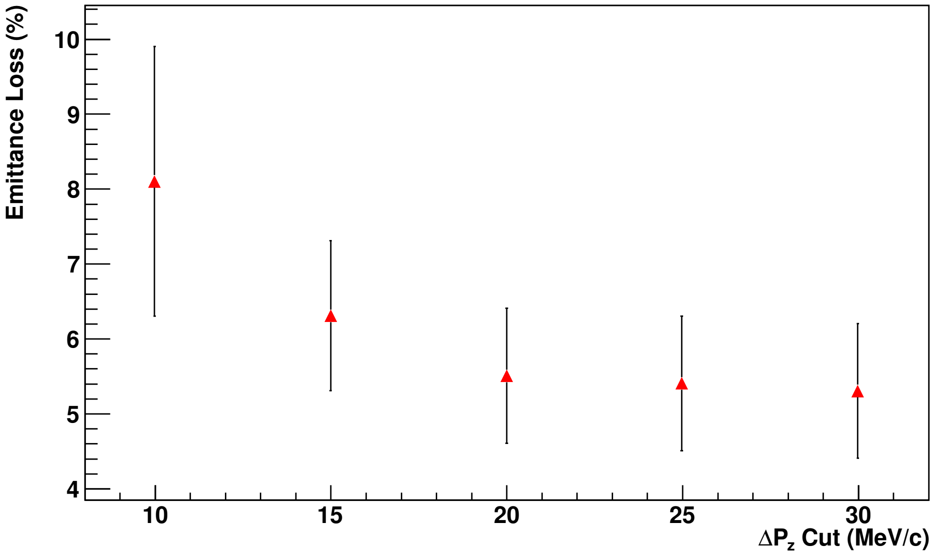

Figure 4.9: Emittance change bias from a P_z selection requiring the loss in the longitudinal momentum to be less than the the value on the x-axis, based on a 6 mm 200 MeV/c beam passing through Step IV.

4.6.2 RF phase selection

An additional selection which could be made is to reject all particles which fall outside the RF bucket. This would involve a measurement of the particle absolute time at a given plane (in general, a detector) and extrapolated to the RF cavities. The phase of the RF would also have to be calculated at this point. If the particle upon arrival lay within the RF bucket it would be kept, otherwise it would be rejected. Although rejection criteria are useful for discrete properties of a beam, such as the particle type, selecting on phase space coordinates will create a structure to the beam which does not match any real scenario. Therefore, in addition to the statistical increase, weighting enables a more realistic particle ensemble to be analysed.

4.6.3 RF Phase Weighting

If the relationship between the particle time and RF phase is known, a weight which corresponds to the proximity of the particle to the target phase can be employed using the Voronoi method. In this case, the target function would be a sine wave with the same period as the RF wave. The Voronoi diagram is drawn in (at least) the time dimension, and the t=0 of the weighting function matched to that of the reference particle (the particle travelling at such a velocity and phase that it is always on the crest of the RF wave).

There are a number of considerations which must be made in the case of sinusoidal weighting which are not an issue in the transverse, Gaussian case. Firstly, the target function is periodic, and so any integration over a complete period (which a single Voronoi cell may cover) will self-cancel. Secondly, the integral of a sine function will produce negative weights.

The first issue can be overcome by either defining a set of initial regions, each of a size of one half-period, to ensure the maximum cell size does not create a zero weight, or by requiring the phase space to be sufficiently dense as to reduce the probability of this occurring.

Further issues arise calculating the phase of the weighting function with respect to the RF. In the simulation, the time structure of the beam is calculated at some reference plane, different to the plane of RF acceleration. The phase of the RF is set once per spill in MICE, and correspondingly once per run in the simulation, and is set with respect to a particle originating at t=0. If the time structure of the beam is known to be uniform about 0, the phase offset at the measurement plane is equal to the mean time of the particle ensemble, representing the mean time of flight from the source of the beam to the reference plane. Then, the weights can be applied based on this offset at the measurement plane and allowed to propagate through the RF cavities (if the phase is set to 0° the mean of the sinusoidal PDF is set to t=0 plus the time of flight to the reference plane). However, if the original time structure is asymmetric about 0, the structure must be extrapolated back from the measurement plane to the source point of the beam, and then weighted according to the location of the t=0 particle. This would be the case if the time between bunches changed significantly due to the scaling of the ISIS RF.

Detector timing considerations

A key component of practically implementing these cuts or weighting is the accuracy to which the particle can be placed on the RF wave. This is a function of both the detector time resolution and the predictability of the particle motion through the optics and material which lies between the detector and the RF.

The trackers are closest to the RF cavities with the least intervening material. However, the time resolution of the tracker has been shown to be of the order of 1 ns. For a 201 MHz wave this does not provide an adequate time resolution. The TOF counters have been shown to operate with a resolution of less than 60 ps, or ∼

1% of the RF period, although they are at a greater distance from the RF cavities, and so some extrapolation and estimation of energy loss effects on the time of flight will have to be included. The energy measurement in the first tracker will help with this.

There are electronics issues involved with digitising picosecond scale events across a millisecond spill, and so the simplest option may be to digitise the RF waveform at each trigger. Since the RF is a periodic function, only the phase of the particle, and not its absolute time, is required for either a cut or a weight to be applied.

4.6.4 Generating the MICE time structure

The ISIS synchrotron has six radio frequency cavities which act on a shifting frequency. This ranges from 1 MHz to approximately 3.1 MHz at extraction. Whilst there is a flat top to the ISIS acceleration, during which the MICE target is dipped, the frequency is still not constant in this period. This is the frequency at which the ISIS bunches hit the MICE target and form a pulsed microstructure to the MICE beam. The ISIS bunches are approximately 100 ns long whilst in the synchrotron, expanding in MICE due to material effects and decays described earlier.

Taking this 100 ns pulse, an approximation of the time structure of the MICE beam over a spill was made. There are further considerations as the frequency of particle triggers is observed to increase as the target dips further into the beam. However, it is neglected here as the number of particles required to simulate this effect whilst keeping a Gaussian shape to each of the micro bunches is very large. The bunches are also considered to be equal in size and uniform in time. An example of this generated beam structure can be seen in Figure 4.10.

The variation of the ISIS RF frequency during the spill is likely to affect the arrival time of the bunches in MICE. The effect this will have on the tracker electronics has been discussed in Chapter 3. It is also likely to affect the phase relationship of the particles in each bunch within the MICE RF system. However, as the ISIS bunch is significantly larger than the MICE RF bucket, subtle phasing relationships are also ignored.

4.6.5 Results of Step V with a realistic bunch structure

With the real ISIS bunch structure defined, the effect of the RF period in MICE on re-accelerating the beam and how this leads into cooling performance was studied. For convenience, multiple bunches were combined into a single distribution as the

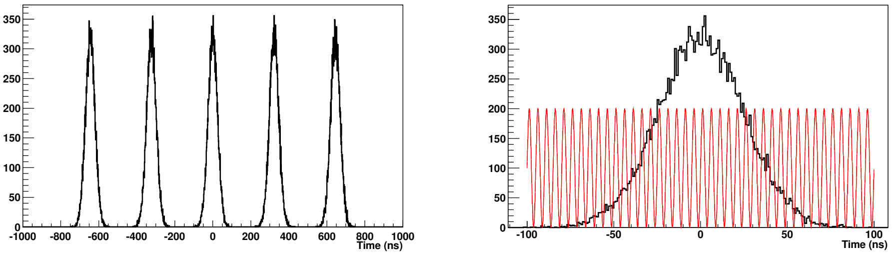

(a) Multi-bunch structure

(b) Single bunch detail. The MICE RF period is overlaid.

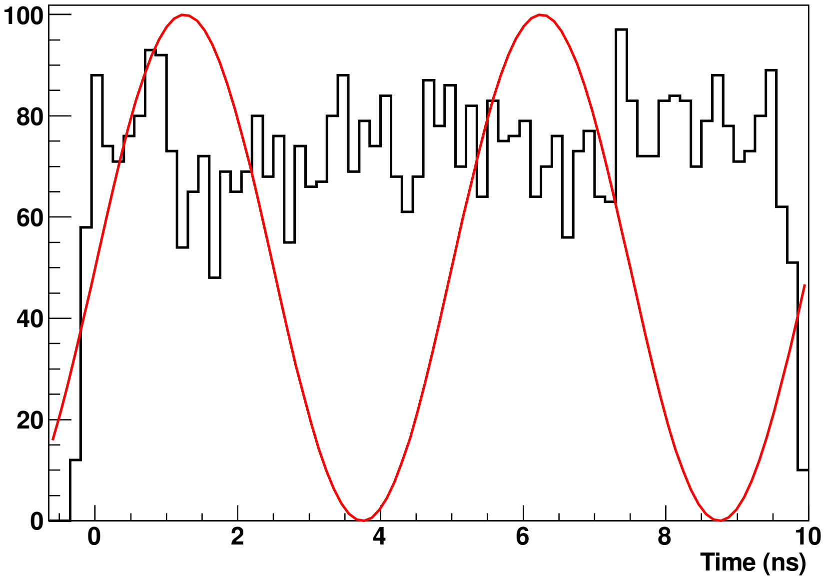

(c) Single bunch reduced to a uniform distribution across the MICE RF period

Figure 4.10: Generated MICE beam structure according to a 3.1 MHz ISIS repetition rate and a 100 ns bunch length.

variation in the timing between bunches is negligible. In the case of real data, these bunches will also be extremely sparsely populated (MICE achieves the order of a few hundred particles spread over many thousands of ISIS bunch crossings).

Ignoring the effect of the ISIS RF ramping, which is absorbed in the size of the bunch in real terms, the ISIS bunch and what is actually observed in MICE are quite different. With several thousand 100 ns bunches creating only a few hundred particles in the MICE hall, the nature of the bunch is irrelevant. The only scale of importance is the period of the RF wave, or roughly 5 ns. When broken into 5 ns slices the Gaussian distribution of the ISIS bunch (combining the particles from many target crossings) will average to a uniform distribution. In these simulations, when considering Step V, the time structure of the beam is assumed to be uniform across one RF period.

Two weighting scenarios were considered: maximum acceleration and a shorter bunch length. In the acceleration scenario, only particles which are expected to be accelerated (whose phase is in the accelerating period of the RF wave) are preferred. To achieve this, each particle is given a sinusoidal weight tuned towards the maximum acceleration of 90°. In the bunch length variation the particles are weighted to a Gaussian distribution with a bunch length similar to that defined in the Neutrino Factory scoping study [26] with a varying mean phase.

To illustrate the effect of the maximum acceleration, a sine function was used to weight the particles in the ISIS bunch structure (approximated as previously described), with the maximum of the sine wave set to varying phases. The precise optimum value here is unknown due to the automatic phasing of the RF, but lies approximately in the expected region. There are also additional complications involved in the times of flight and the transit time factor associated with RF cavities. The effect of the sine weighting is shown in Figure 4.11 and the Gaussian mean time weighting in Figure 4.12.

Both the sine and mean time weighting show a dependence of the level of cooling on the phase of the RF. This implies that selection of the particles to be on phase is important not just in the longitudinal direction for maintaining the momentum, but also in the transverse plane for avoiding negative contributions to the cooling performance. Since MICE aims to use an RF phase of 90°, corresponding to the maximum acceleration and no drop in the cooling performance shown by this study, there is not expected to be a detrimental effect to the MICE results. It does however generate another motivation for selecting the particles which receive the maximum re-acceleration from the RF, either by this weighting algorithm or some other technique.

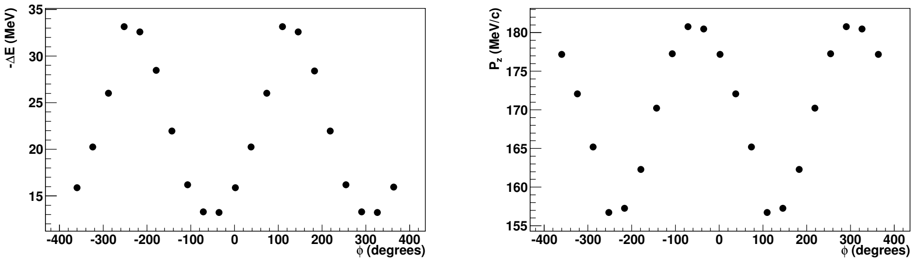

(a) Maximum probability sin(t) versus change in energy (b) Maximum probability sin(t) versus P_z

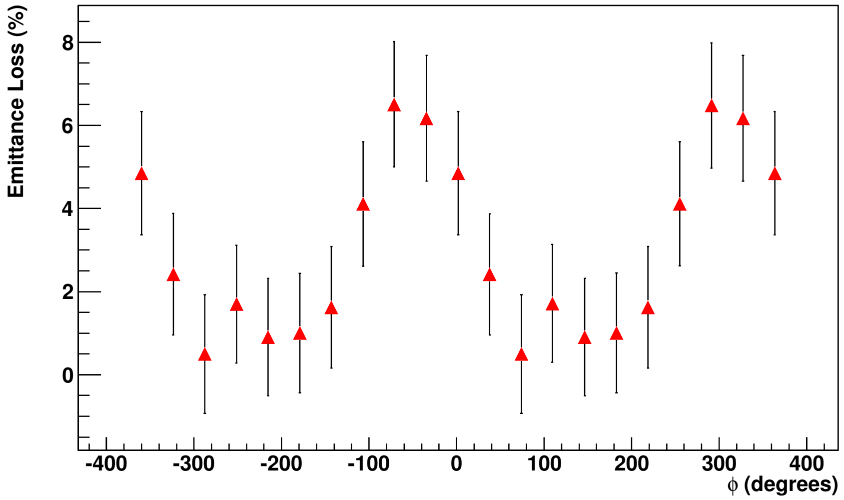

(c) Fractional emittance change versus time of maximum probability

Figure 4.11: Change in longitudinal momentum and emittance loss from weighted sets, where the weights are generated from a sine function with a phase offset. There is a symmetry in the plots due to the periodic nature of the weighting.

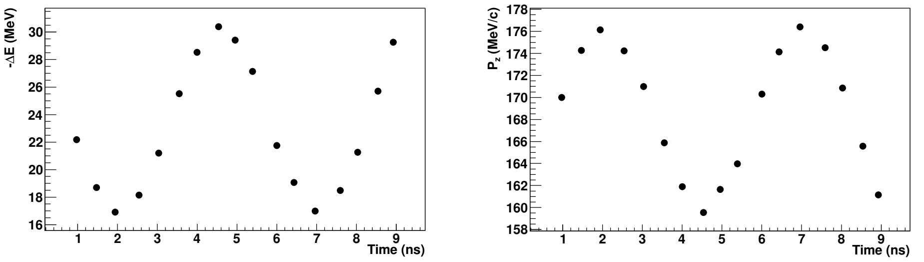

(a) Gaussian $\mu = t$, $\sigma = 0.5$ versus change in energy

(b) Gaussian $\mu = t$, $\sigma = 0.5$ versus P_z

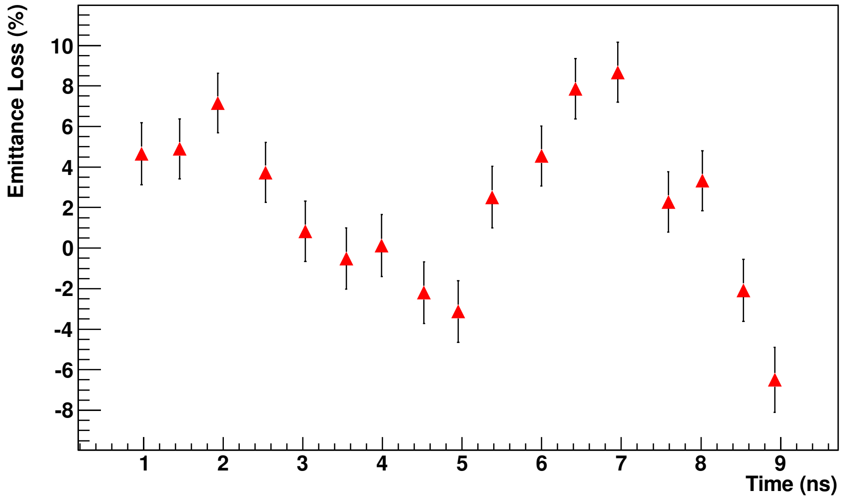

(c) Fractional emittance change versus Gaussian $\mu = t$, $\sigma = 0.5$

Figure 4.12: Change in longitudinal momentum and emittance loss from weighted sets, where the weights are generated from a Gaussian distribution with a mean time and a width of 0.5 ns.

4.7 Weighting intermediate emittance beams

The operating modes of MICE, given previously in Figure 2.2, are limited to just three emittance values; 3, 6 and 10 mm. It is desirable to know the behaviour of the cooling channel as a function of the input emittance, and so more data points are required. To achieve this, one of the existing, real data points could be re-weighted to have the properties of the intermediate points. This could be achieved by either fully calculating the covariance matrix of the new point, or scaling the beam size.

4.7.1 Beam RMS Volume Scaling

An alternative method is to scale the total volume occupied by the beam in phase space. Taking variances in each of the dimensions occupied by the beam and incrementing them, a selection of Gaussian functions can be found in each of these dimensions. The PDF of the weighting function is then defined in Equation 4.4, where G is a mono-variate Normal distribution with a mean of 0. The weight is then given in Equation 4.5, where the 1 and 2 suffixes refer to the limits of the region in which the point lies.

$$f(p) = \prod_i G(p_i, \sigma_{pi}) \tag{4.4}$$

$$w = \prod_i \left[\mathrm{erf(p_{i2})} - \mathrm{erf(p_{i1})}\right] \tag{4.5}$$

To keep cylindrical symmetry, the widths of the transverse position and momentum σ are kept the same. This method is viable as long as the β function is constant. This corresponds to a beam in which $\alpha = 0$, which is the design case for beams within the tracker solenoid, and so meets the condition.

A more involved approach, which gives more credence to the coupling of the beam coordinates but suffers from the added dimensions, would be to fully calculate the covariance matrix of the intermediate steps. With a defined ϵ, α and β the covariance matrix can be formed as previously shown. The full covariance matrix must then be diagonalised and the phase space points of each particle rotated into the diagonalised frame.

Statistical fluctuations

With the weighted sets being the same distribution of particles as each other, promulgation of statistical errors is extremely likely, causing a strong correlation between each weighted emittance point. It is also possible that weighting can exacerbate statistical fluctuations by strongly weighting a region of phase space which has varied significantly from the mean. If a particle lies in a region of phase space such that the boundaries of its Voronoi cell are limited only by the initial region (i.e. it is at the tail of the distribution), then the weight applied to that point will be limited only by the initial region size. Particle with large weights generate the following dilemma: they likely reside in the region of phase space of interest to the algorithm, but can also cause the distribution to be dominated by only a few events.

In terms of merely the phase space population of the beam, highly dense regions will not have any negative effect due to the division of the Voronoi algorithm. However, low density regions have the possibility of generating a very large weight as the PDF is integrated over the the region. It then becomes necessary to remove particles with extremely high weights, although this generates the risk of biasing the distribution away from the phase space area of interest. To mitigate the large weight effect, a cut is made at $P \simeq 0.01$ on the Poisson distribution of the weights. An example of the distribution of weights can be seen in Figure 4.13. However, this combines the particle's at the edge of the region and those with a naturally high weight together, whereas perhaps only the former should be removed. This could be achieved by removing those particles which reside in the edge regions of the Voronoi diagram, or by re-setting their weight to one.

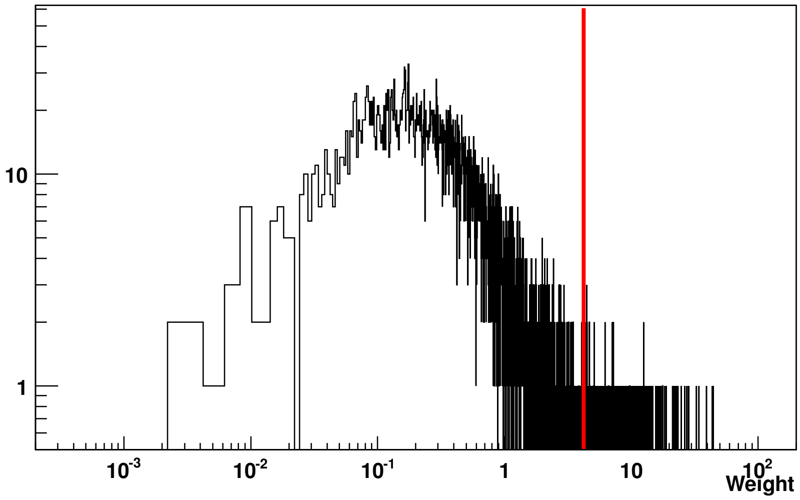

Figure 4.13: Typical distribution of weights. The standard selection for this distribution is shown by the red line in the figure.

Statistical fluctuations in the changes to the particle, which are outside of the control of the weighting algorithm, are likely to suffer the aforementioned effects, especially when considering the strong correlation between weighted sets based on the same ensemble of points.

4.7.2 Step IV Results

This weighting technique has been employed in order to measure the emittance reduction in beams which do not fit the MICE operating modes. Beginning with simulated beams corresponding to the expected measured points of MICE, weights have been calculated for intermediate emittance beams, and the weighted emittance reduction measured. The rate of change of the emittance reduction is expected to follow

$$\frac{-\Delta\epsilon_N}{\epsilon_N} = \frac{(\eta\epsilon_N - \zeta)}{\epsilon_N} \tag{4.6}$$

where η corresponds to the contribution from energy loss and ζ to the contribution from multiple scattering, as shown in Equations 4.7 and 4.8.

$$\eta = \langle\frac{dE}{dX}\rangle\frac{dX}{\beta^2 E_\mu} \tag{4.7}$$

$$\zeta = \frac{\beta_T(0.014\mathrm{GeV})^2 dX}{2\beta^3 E_\mu m_\mu X_0} \tag{4.8}$$

Here, β represents the relativistic speed of the particle and β_T the transverse β function, while X_0, $\langle\frac{dE}{dX}\rangle$, E_μ and m_μ are the radiation length of the absorber material, the mean energy loss in the material, the energy of the muon and its mass, respectively.

These values are themselves functions of constants such as the mean energy loss and radiation lengths in the absorbers, and variables which depend on the beam, such as energy and beta function. A fit can be made to any selection of $\Delta\epsilon$ points to obtain these values.

The fit, as will be shown, can be formed from weighted sets of points as well as the true, unweighted set. Once a selection of weighted sets have been calculated, it is then possible to fit this function from a single measurement.

For each of the modes of MICE the performance of Step IV was simulated, and the resultant data weighted according to methods previously described. The emittance was calculated for both the source and each weighted set at the two tracker reference planes, using Monte Carlo truth information. The diagram was drawn about the points in the first reference plane and the weights calculated according to the desired beam parameters. On a particle-by-particle basis, the weight was then re-applied to the measurement in the second tracker reference plane, and the emittance calculated (for the source set, the weight can be assumed to be 1).

Transverse β function dependence in weighted sets

Scaling the transverse position and momentum variances and weighting accordingly will provide an inflation in emittance for which the fractional change can be measured. However, this produces effects which sometimes do not match the true values to which they correspond. The residual of the fit to the true values is dependent on the ratio of the transverse variances to which the beam has been weighted. This ratio corresponds to $\sqrt{\gamma/\beta}$; where the α function is 0, this is reduced to $1/\beta$. However, simply setting the ratio of the weighted sets to match the β function of the source set (or the desired β for the lattice) may not solve the issue. If the source set itself no longer has a zero α function, or has become mismatched due to field irregularities or matter effects, the desired phase space may still not be occupied. Whilst it is possible to observe some expected effects in the β weighted sets, the variation across the full spectrum, and its high dependance on the outlying weights, make any conclusion difficult.

Available populated phase space

In the case of weighting to a Gaussian where the ratio of the variances is constant, corresponding to a constant β with $\alpha = 0$, the β dependance can be accidentally introduced by a lack of desired populated phase space. In 2D, beginning with a circle within a larger ellipse, the ratio can be kept whilst within the larger ellipse. However, when one of the axes extends beyond the available populated space within the ellipse, the area of the desired circle now is only provided by the other axis of the ellipse. This breaks the constant ratio and alters the betatron function. An example is given in Figure 4.14.

The available populated phase space is of concern to this study given the relatively low statistics. Given that the expected number of events for a precision measurement without weighting is expected to be 10^5 muons in the RF bucket, the phase space density will not be an issue for real data.

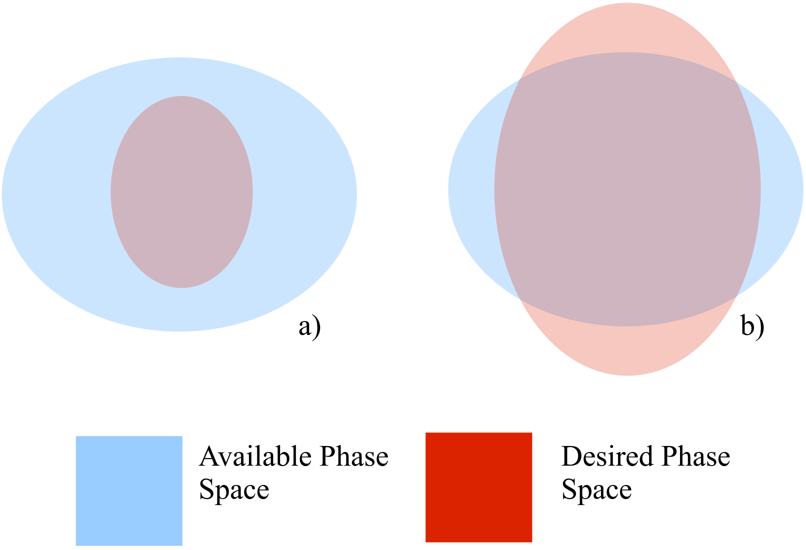

Figure 4.14: Available phase space contribution to the weighted distribution.

Step IV Results summary

Comparing the fit to simulated beams at the same momentum but differing emittance tests how well the weighting and the fit match that run mode. An example is given in Figure 4.15. To minimise the residual, the weighting function has been chosen to have a β as close as possible to that of the source set. A summary of the results of weighting for the points of Step IV is given in Table 4.2, where the χ^2/DOF and the residual δ/DOF are given between the fit based on weighted sets and seperately simulated 'true' values.

P_z (MeV/c)	ϵ_N (mm)	$\Delta\epsilon_N$ (%)	η	ζ (mm)	χ^2/DOF	δ/σ/DOF
200	3.05	3.6	7.6	15.5	1.30	1.12
200	5.32	5.4	6.2	8.0	1.22	0.85
200	8.23	6.8	5.0	9.1	6.61	2.52
240	3.10	1.9	3.3	4.9	0.28	0.12
240	5.70	2.0	4.3	7.6	0.50	-0.43
240	8.65	3.3	3.6	7.1	0.35	0.29

Table 4.2: Summary of Step IV fractional emittance losses. The corresponding values for the original set are given, the χ^2 of the fit to the original value provided in the penultimate column, and the residual between the fit and the true points in the final column.

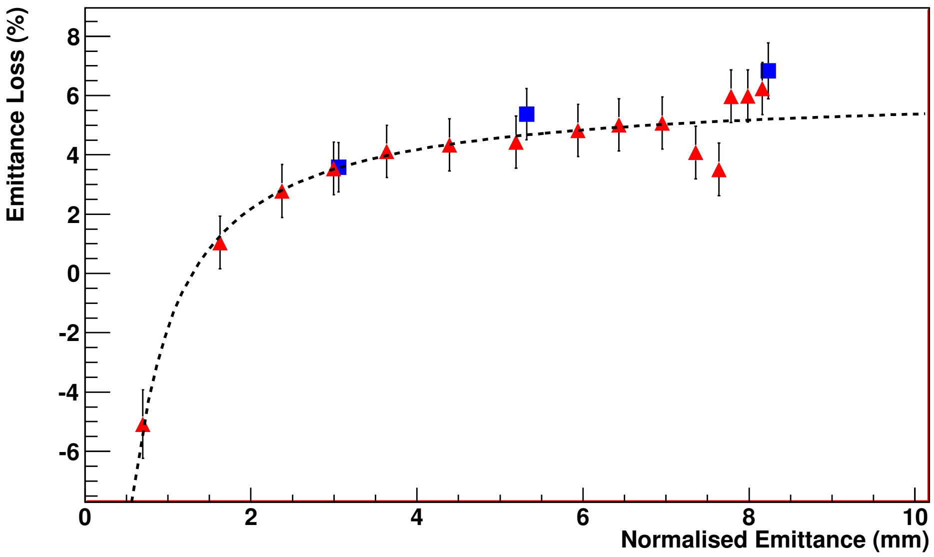

Figure 4.15: Weighted fractional emittance loss through Step IV for various beams under RMS scaling (red triangles), weighted from a single original set ($\epsilon = 6\,\mathrm{mm}$) and compared to other simulated beams in a Step 4 200 MeV/c mode (blue squares).

The suspected cause of the variations in the weighted points visible in the figures with respect to the fit line is commented on in Section 4.7.1, where strong correlations between the weighted sets reduce the effective error when they are combined. In some situations, the correlations can create the effect of causing many points to vary in the same way, due to the identical values of the points and the similar values of the weights.

The β dependence in the weighted sets provides an opportunity to study the effect of varying the β function. Using the results from the emittance weighting scheme, scaling β can be combined with the emittance scaling to provide a 2D impression of how ionisation cooling depends on the input emittance and betatron function. This is shown for the $\epsilon = 6\,\mathrm{mm}$, $P_z = 200\,\mathrm{MeV/c}$ case in Figure 4.16.

From the figure it is difficult to observe any real dependence on the β function at the reference plane. Given Equation 1.28 one would expect a lower β function to produce a larger drop in emittance, forming a diagonal gradient to this plot. As the β function is strictly a parameter of the magnetic lattice, and not the beam, and particles which corresponding to a differing β function are not, by definition, matched particles, they may not be transmitted. Additionally, the value of β at the reference plane is likely quite different from the value at the interaction point, which

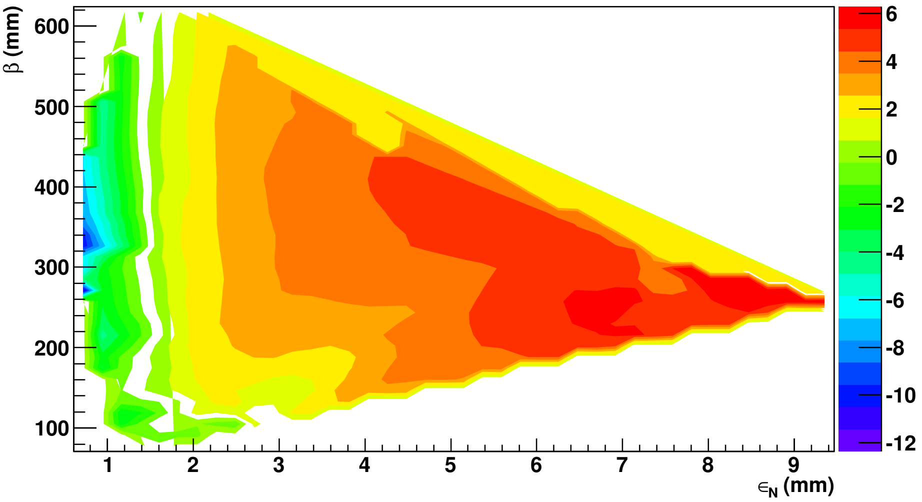

Figure 4.16: Weighted emittance loss (%) through step IV for various beams under RMS scaling, where the ratio of RMS values have been varied to produce varying β.

is the key location for ionisation cooling. In order to find this value, the β at the reference plane should be transported to the interaction point either by tracking or a transfer matrix. Information on the β function at the reference plane is therefore included for reference but it is unlikely any meaningful information can be obtained from it. Varying β in this way is also much more sensitive to weight fluctuations than the emittance scaling, which adds another region of caution.

4.8 Conclusions

By implementing a weighting technique based on approximate Voronoi diagrams, a single measurement based on a set of particles in a single MICE mode can be used to predict the effect on the beam in other modes. The change in emittance as a function of both the input emittance and β function can be predicted from a single point, and improved by the amalgamation of many points. The fitted emittance loss as a function of input emittance has been found for each of the MICE running modes.

Issues have been found with the weighting algorithm; since the loss in emittance is dependant on a number of factors, the algorithm's ability to include them all into the weighted set is not complete. Whilst scaling the emittance itself is achievable by scaling the variances of the weighting function, matching this to the β function is an added difficulty. Further issues arise if the source set does not have a zero α or its angular momentum does not cancel in the covariance matrix defined for a matched beam within a solenoid. If the beam is no longer matched and subject to non-linear effects, the available phase space becomes untenable for the target function. These issues are expected to improve as the data set size increases and the phase space becomes more populated.

In the time dimension, there are fewer problems concerning phase space and multi-parameter matching. There is no approximation when drawing the Voronoi diagram and thus no error associated with the cell sizes. By matching the PDF of the weighting function to the design phase of the RF wave, those particles which MICE is most interested in, i.e. those on the crest of the RF, are given statistical preference. Using both a sine function and low variance Gaussian functions, the relationship between the RF phase and the cooling performance was measured, suggesting that particles on the decelerating edge of the wave will see reduced cooling performance.

Statistical and correlative concerns must be taken as a caveat when considering the results, but the method and its applicability are shown to be performing well and could be used to conduct similar analyses with the real MICE data.

Chapter 5

Investigations into approximating the MICE cooling channel

5.1 Motivation

The cooling channel of a full Neutrino Factory will comprise many cooling cells such as those found in MICE. It is possible that whilst providing a proof-of-principle demonstration of ionisation cooling, MICE might also provide a prediction of the level of cooling which can be reached by a full cooling channel. This might be obtained by use of a data verified Monte Carlo or an analytical expression. Such an analysis is typically performed with a semi-analytical model of the accelerator in a matrix-based approach, rather than a full tracking simulation. It is the aim of this study to investigate whether such a matrix can be found based on the real measurements of MICE.

The typical use of transfer matrices depends on the linear behaviour of the beam though the magnetic lattice, a requirement of which is the vacuum or low mass medium through which the particles travel. The use of such a matrix in MICE would be expected to accurately model the magnetic effects, but may have difficulty in accounting for the presence of material, requiring additions to the method. In any case, the purely linear initial transfer matrix must be investigated first, and any required corrections added later.

5.2 An Ionisation Cooling Transfer Matrix

At the simplest level, the MICE channel can be formed from the magnetic lattice alone, comprised of a series of solenoids in the tracker spectrometer, the absorber focus coils and the RF coupling coils. It is possible to approximate this field using a linear transformation in the form of a transfer matrix, with the non-linear effects of the solenoidal fringe fields approximated by higher order terms in additional, larger transfer matrices.

The MICE lattice is not, however, in a vacuum. The stochastic processes from energy losses and multiple Coulomb scattering affect the momentum of the beam in a way independent of the magnetic fields, and hence the matrix approximation previously described. Energy losses have a functional prediction in the form of the Bethe-Bloch formula, and so the changes in the momentum may be predicted using the known material volumes and densities. In contrast, multiple scattering smears the possible position and momentum of the particle, albeit with a predictable distribution.

5.2.1 Forming the matrices for the cooling channel

It has already been shown that the position and angle of a particle at one point in a lattice may be transported to another point by means of a set of linear equations, combined for convenience into a matrix. To first order, this is given by:

$$\begin{pmatrix} x \\ P_x \end{pmatrix}_{out} = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} x \\ P_x \end{pmatrix}_{in} \tag{5.1}$$

Higher order effects, including the correlated terms, may be introduced by adding another differently sized matrix, as given in Equation 5.2. This is still a linear mapping and represents the magnetic part of the cooling channel. The components of these matrices could be found by any fitting technique.

$$\begin{pmatrix} x \\ P_x \end{pmatrix}_{out} = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} x \\ P_x \end{pmatrix}_{in} + \begin{pmatrix} N_{11} & N_{12} & N_{13} \\ N_{21} & N_{22} & N_{23} \end{pmatrix} \begin{pmatrix} x^2 \\ P_x^2 \\ xP_x \end{pmatrix}_{in} \tag{5.2}$$

The difference between a regular lattice and the MICE cooling channel is the large amount of material effects, both in energy losses and multiple scattering. As both of these processes are stochastic, though with a functional dependance on their mean values, the linear approach is no longer valid. Here another matrix may be introduced which contains a functional dependance on one of the elements of the particle state vector, given by:

$$\begin{pmatrix} x \\ P_x \end{pmatrix}_{out} = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} x \\ P_x \end{pmatrix}_{in} + \begin{pmatrix} 0 \\ f(P_x) \end{pmatrix} \tag{5.3}$$

The functional dependance on the transverse momentum included in the matrix would then be the initial momentum minus the mean losses from the Bethe-Bloch process given in Equation 1.30. Multiple scattering poses a problem in predicting the particle trajectory as its net effect is zero, and only introduces an additional error into the fitting algorithm. The final term is then only a correction and is not actually part of the matrix.

Accurately defining the contribution of these stochastic processes in the fitting algorithm to a degree where the iterations finally converge is an involved process, especially considering the evolution of the beam through the material. It is still necessary to implement the purely linear part of the phase space transformation, as this will contribute the effects of the magnetic lattice, and may absorb some of the matter effects into the elements of the matrix.

5.3 Finding the matrix parameters

It is conceivable that the matrix for the MICE cooling channel could be found analytically assuming the transfer matrix of a solenoid is already defined. However, the radial field components introduced by the opposing longitudinal fields and other fringe field effects complicate the simple multiplication of many solenoidal transfer matrices. Whilst analytically complicated, the channel still represents a time-independant magnetic field through which the particle motions should be well defined. Using the position and momentum measurement from the trackers, a fit could be made to the transfer map that does accurately describe the particle trajectories.

5.3.1 Neural Network

Given the nature of the MICE data, which presents a set of values *in* and *out* of the cooling channel, during which a complex set of physical processes affect the particles, an artificial neural network is a candidate for finding the mapping from the inputs to the outputs [84]. The network is formed of layers where each object in a layer is the sum of weighted inputs from the previous layer, as given in Equation

5.4. The process of calculating the weights in each layer of the network is as follows:

- **Predict Outputs** - Using the current iteration of the weights, the output of each point is calculated as the sum of the weighted inputs. Intermediate layers can be used where the outputs of each layer are the weighted sum of its inputs, and then the real outputs are the weighted sum of the final intermediate outputs. The output for a single layer is given by:

  $$O_j = \sum_i w_{ij} I_i \tag{5.4}$$

- **Calculate Partial Errors** - The error on the output as a function of each weight, $\partial F/\partial w_{ij}$, is then calculated as

  $$\left(\frac{\partial F}{\partial w_{ij}}\right) = \sum_p \left(\frac{\partial F}{\partial p_o}\right) p_i = \sum_p (p_o - t_o) p_i \tag{5.5}$$

  where p_i, p_o are the inputs and predicted outputs of the point p, respectively, t_o the true value of the output and F the error function of the distribution.

- **Update Weights** - Each weight is then updated by the partial error on that weight multiplied by the learning rate μ:

  $$w_{ij} \rightarrow w_{ij} - \left(\frac{\partial F}{\partial w_{ij}}\right) \mu \tag{5.6}$$

- **Update until convergence** - The process is repeated until the total error on the sum of the points falls below an acceptable level or a maximum number of iterations have been completed. The weights in this method are then the elements of the transfer matrix.

The learning rate μ is a sensitive parameter. If it is too small, the process may never converge to the optimum value, but if it is too large, the weights could jump too far either side of the optimum or even diverge. Likewise, the minimum acceptable error which is used to end the iterative process affects the overall efficiency. These values are largely subjective and found by experience.

A custom neural network was written in C^{++} to implement this method. The network was initially tested on known data outputs and found to converge onto the expected weights in an equivalent number of dimensions to that of the cooling channel problem. The speed of the algorithm is largely dependent on the number of points and the number of iterations which it is allowed to go through. In this case, the maximum number of iterations was set to 1000, which was invariably reached, mainly due to an overly optimistic requirement on the minimum error.

5.3.2 Mapping Complexity

To first order, the mapping of a particle state at one position in the beam line can take the form of a linear combination of the state parameters. This is true, at this level, for a purely magnetic lattice. This is included into the network as the sum of weighted inputs, where the weights are the elements of the transfer matrix. Since the linear mapping may be expanded to include higher order terms, additional matrices are also represented by the sum of the weighted inputs. Therefore, in terms of the network, the component of the predicted output O_j is the same as Equation 5.4.

Adding higher order matrices includes a larger number of additional weights, such that with each additional higher order n, the number of weights increases as $I^n O$, where I is the number of input nodes on the point and O is the number of output nodes.

Whilst a matrix approach (and by definition a linear fit) may work in a purely magnetic lattice, the MICE cooling channel, by design, breaks the linear motion of the particle by reducing the total transverse momentum without any corresponding increase in phase space (as would be the case by defocusing the beam). Multiple scattering and energy losses in the absorber, as well as momentum gains in the RF, cannot be described in a purely linear way. Their functional form may, to some degree, be included in the neural network, which need not follow the previous method of a sum of weighted inputs. The outputs may be defined in any way, although the partial derivative of the cost or error function still needs to be calculable in order to minimise the errors.

Solutions for the higher order matrices and material effects were not found in a manner in which the algorithm converged. For the higher order matrix, the difficulty lay primarily in the choice of the learning rate, as a dependence on the square of a parameter will clearly have a more dramatic effect than the parameter itself. The fit therefore began with the linear component of the particle behaviour and tested to measure the viability of this section of the algorithm.

5.3.3 Network Design

The structuring of these transfer matrices into neural networks takes the from of feed-forward networks with multi-layer perceptrons. The initial, simple case corresponding to Equation 5.1 is a single layer network with a linear activation function of the form given by Equation 5.4. This is shown schematically in Figure 5.1.

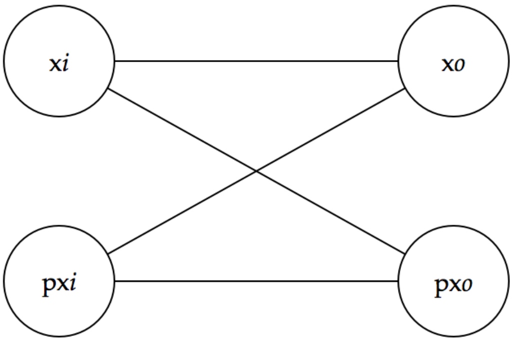

Figure 5.1: Single layer neural network to describe the first-order transfer matrix.

As the transfer matrices are expanded to include higher order terms given in Equation 5.2, a hidden layer is added to the network, as shown in Figure 5.2. Here the propagation of the hidden layer takes the form of a product, rather than a sum, to accommodate the higher order components of the inputs. This is commonly referred to as a $\Sigma\Pi$ unit, as opposed to the summation Σ unit. The outputs then become the weighted sum of the hidden and initial layers, and back-propagation is used to train the weights of the hidden layers.

Finally, the inclusion of a functional modification of the input momentum from material effects takes the form of an additional hidden layer, shown in Figure 5.3. Here the activation function is no-longer linear, taking the form prescribed by the effects of multiple scattering and energy losses.

In each of these cases the learning method is supervised as the known inputs and outputs are used to adjust the internal weights.

5.4 Results

Taking the Monte Carlo values from the tracker reference planes either side of the cooling channel, the phase space coordinates of each particle were entered into the

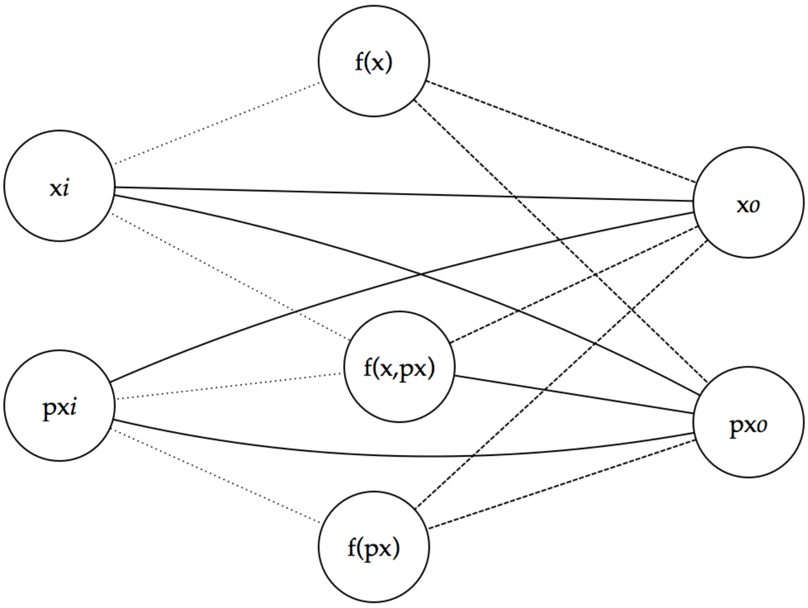

Figure 5.2: A neural network with a hidden layer to describe the higher order transfer matrix.

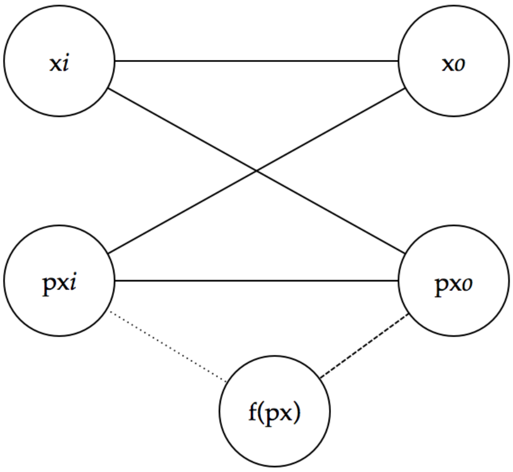

Figure 5.3: A neural network with a hidden layer to model the effects of the material interactions.

neural network as a training point. The weights which best fitted the transformation of the inputs to the outputs were then calculated using the method outlined above. The sets of points were limited to those from a common input momentum and emittance, including other time varying effects. The transfer matrices derived for the 200 MeV/c longitudinal momentum case of Step IV are given in full in Table 5.1.

Using the final matrix to map the input values to the second tracker reference plane, the residual between the matrix prediction and the true output value (equivalent to $p_o - t_o$) gives the resolution of this method in each of the phase space coordinates. This is shown in Figure 5.4 for the central points of the Step IV matrix.

ϵ	P_z				
3	200	-0.27	-1.05	0.92	-0.45
		1.49	-0.48	0.58	0.82
		-0.92	0.45	-0.27	-1.06
		-0.58	-0.82	1.49	-0.48
6	200	-0.30	-1.01	0.88	-0.49
		1.45	-0.51	0.61	0.78
		-0.88	0.50	-0.31	-1.01
		-0.60	-0.78	1.44	-0.50
10	200	-0.37	-0.89	0.77	-0.55
		1.38	-0.53	0.63	0.70
		-0.79	0.55	-0.36	-0.91
		-0.62	-0.70	1.38	-0.54

Table 5.1: Elements of the fitted Step IV transfer matrices for various input emittances with a common 200 MeV/c longitudinal momentum.

5.4.1 Longitudinal Momentum Slicing

The transfer matrix method of mapping a particle's phase space at one point in a lattice to another is dependent on the uniformity of the longitudinal momentum. This is particularly true in quadrupoles where the focusing strength is proportional to the longitudinal momentum, and less so in the case of solenoids where it is the transverse momentum that affects the focusing.

The effect of the longitudinal momentum spread on the algorithm's ability to produce an accurate transfer matrix was investigated by comparing the residual errors for a similar beam with varying momentum spreads. As shown in Figure 5.5 the quality of the single transfer matrix becomes poorer with increasing momentum spread. It would be impossible to use a single transfer matrix for the MICE cooling

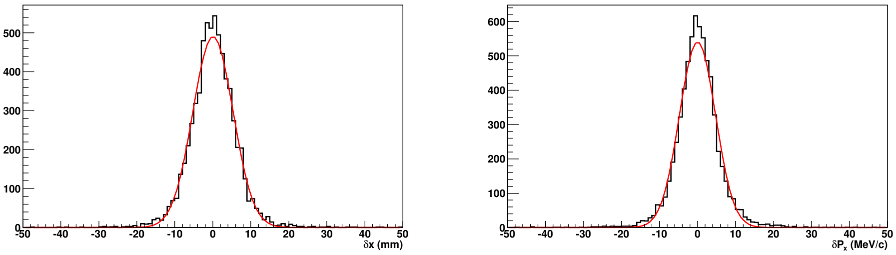

(a) δx (b) δP_x

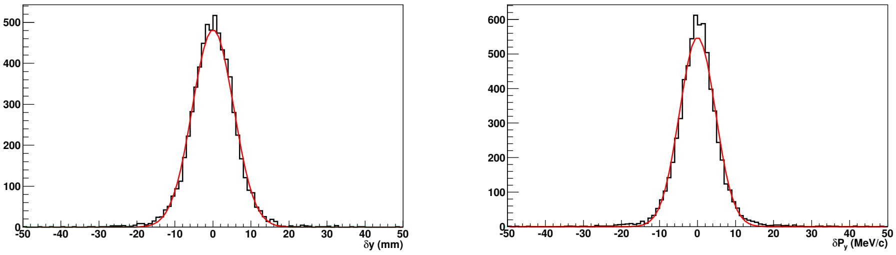

(c) δy (d) δP_y

Figure 5.4: Residuals between the matrix prediction and the actual values for position and momentum. Calculated from a 6 mm, 200 MeV/c beam passing through Step 4. The fitted σ values are $\sigma_x = 5.30\,\mathrm{mm}$, $\sigma_{P_x} = 4.75\,\mathrm{MeV/c}$, $\sigma_y = 5.44\,\mathrm{mm}$, $\sigma_{P_y} = 4.73\,\mathrm{MeV/c}$.

channel which naturally has a large longitudinal dispersion without making selection requirements on each particle's longitudinal momentum.

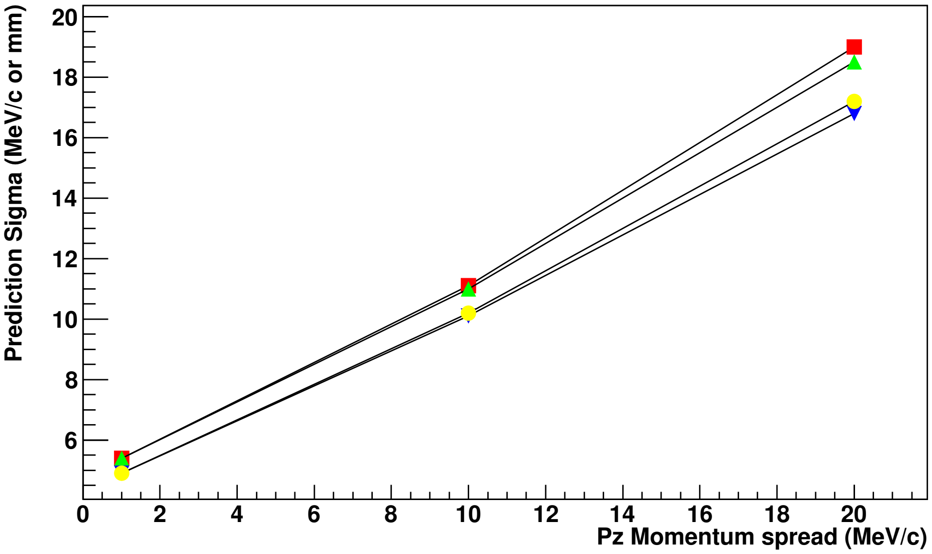

Figure 5.5: P_z spread errors. The fitted Gaussian σ in each phase space coordinate is shown for a distribution of particles with a given P_z spread. The green triangles, red squares, yellow circles and blue inverted triangles plot x, P_x, y, and P_y respectively.

To improve this effect, the beam was sliced in P_z and a separate fit performed on each slice. An example for a 200 MeV/c beam with $\sigma_{P_z} = 20$ MeV/c is shown in Figure 5.6. This led to 40 slices each 2 MeV/c wide. The elements of the transfer matrix at each slice are plotted against the mean P_z of the slice, and a sinusoidal dependance is seen. This shows that the elements of the matrix evolve with the longitudinal component of the beam momentum.

5.4.2 Results summary

Running the algorithm on simulated modes of Step IV, the resolution of each network is summarised in Table 5.2. Given the issues surrounding the accurate inclusion of matter effects as well as insuring algorithm convergence, only the results from the mapping corresponding to Equation 5.1 are given.

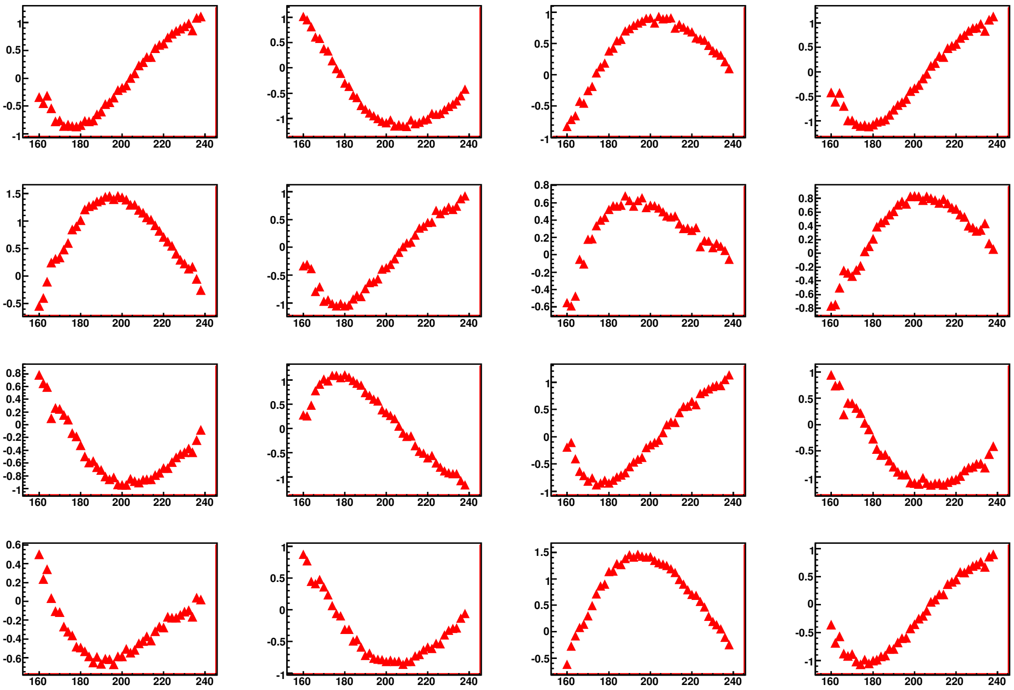

Figure 5.6: P_z dependance of the elements of a Step IV cooling channel transfer matrix. The red triangular points were found by slicing the points in 2 MeV/c bins and a separate network fit performed on each slice. The location of the graphs in the figure corresponds to the element of the matrix, where the x-axis is the mean P_z of the sub-ensemble of particles, and the y-axis is the value of the matrix element returned from the neural network.

ϵ	P_z	σ_x	σ_{P_x}	σ_y	σ_{P_y}
3	200	4.41	3.53	4.49	3.56
6	200	5.30	4.75	5.44	4.73
10	200	7.32	7.39	7.38	7.38
3	240	2.12	3.10	2.09	3.13
6	240	3.18	4.09	3.24	4.15
10	240	6.01	6.32	5.68	6.54

Table 5.2: Resolution of the transfer matrices found for each mode of Step IV, using different network approaches. Position resolutions are in mm and momentum resolutions in MeV/c.

5.5 Mapping emittance changes

Having derived a transfer matrix which can map the phase space coordinates of a particle from one point to another, it should also be possible to map the covariance matrix. The mapping of the covariance matrix is given by Equation 5.7, which allows for the emittance, which is related to determinant of this matrix, to also be mapped as given by Equation 5.8.

$$\begin{pmatrix} \sigma_x^2 & \sigma_{xx'} \\ \sigma_{xx'} & \sigma_{x'}^2 \end{pmatrix}_2 = M \begin{pmatrix} \sigma_x^2 & \sigma_{xx'} \\ \sigma_{xx'} & \sigma_{x'}^2 \end{pmatrix}_1 M^{\dagger} \tag{5.7}$$

$$\epsilon_2^2 = det\left(M\Sigma_1 M^{\dagger}\right) \tag{5.8}$$

The transfer matrix can then be used to predict the emittance change without re-measuring the beam at the end of the cooling channel. This method has a number of possible problems. For example, the error on the transfer matrix could be large enough to mitigate the prediction of the emittance change, if the ratio of the matrix resolution to the variances of the beam is high.

Using the matrices found in the previous sections, the emittance change from Equation 5.8 was calculated and compared to the actual measured change in the simulated runs. A summary of this prediction is given in Table 5.3.

Unfortunately, the quality of this prediction is too poor to serve as a indicator of the likely cooling performance in its current form. The expected emittance reduction is given by $\Delta\epsilon_T$ as calculated by the covariance matrix of the particle distribution at the second measurement plane. From this table it can be seen that the predicted values of the emittance loss by this method are too high. Errors introduced by multiple scattering and energy straggling outweigh the convergence of the fit onto values which accurately represent the motion due to the magnetic field

ϵ	P_z	ϵ_N	$\Delta\epsilon_T$	$\Delta\epsilon_M$
3	200	3.04	3.1	7.4
6	200	5.27	5.1	9.3
10	200	8.05	6.5	15.7
3	240	3.11	1.3	5.0
6	240	5.84	2.3	7.6
10	240	8.53	2.8	13.1

Table 5.3: Prediction of emittance change (%) based on the transfer matrices, with the true fractional change in emittance, $\Delta\epsilon_T$, and the mapped change, $\Delta\epsilon_M$, where the error on $\Delta\epsilon_T \simeq \pm 0.8\%$.

and mean energy losses in the absorber, which was expected when using the purely magnetic style of the matrix to account for all the physical processes in the cooling channel. There is a notable correlation between the quality of this emittance prediction and the errors from the fits given in Table 5.2, themselves corresponding to the longitudinal and transverse components of the beam. As the momentum increases, and the scattering decreases, the quality of the prediction improves; as the emittance increases, and the energy losses in the transverse plane increases, the fit becomes poorer. Both of these effects are attributed to the absence of energy losses and multiple scattering in this version of the network. Although this method alone produces inaccurate results, the nature of the errors verify the suspected issues and indicate that the inclusion of stochastic processes would improve the fit performance.

By stepping the particle through the absorber and calculating the effect of the material at each step, a more effective expression of the electromagnetic processes could be found. Whilst more complicated than a single step approach, it would still be simpler than a Geant based tracking algorithm, which would calculate these steps for every particle, rather than assuming the same values for the entire beam.

There is an additional feature linked to the evolution of the beam envelope through the cooling channel. As the β function changes the density of phase space is moved from the momentum to the position planes. With the β function minimised in the absorber the momenta are maximised; it is here the losses to the momenta from interactions are applied, and it is these values which should be calculated. However, as the beam passes through additional magnetic fields prior to the measurement plane the envelope evolves and the momentum phase space, including the losses, are transferred to the position plane. For this reason a correction due to the stochastic processes made at the measurement planes would have to apply to both the position and momentum values. Considering purely an adjustment to the momentum of the mapped phase space coordinates applied to the absolute values of momentum, the effect on the emittance loss can be seen in Figure 5.7. This is equivalent to using the found transfer matrix to map each particle, and then applying a correction such as that given by Equation 5.3.

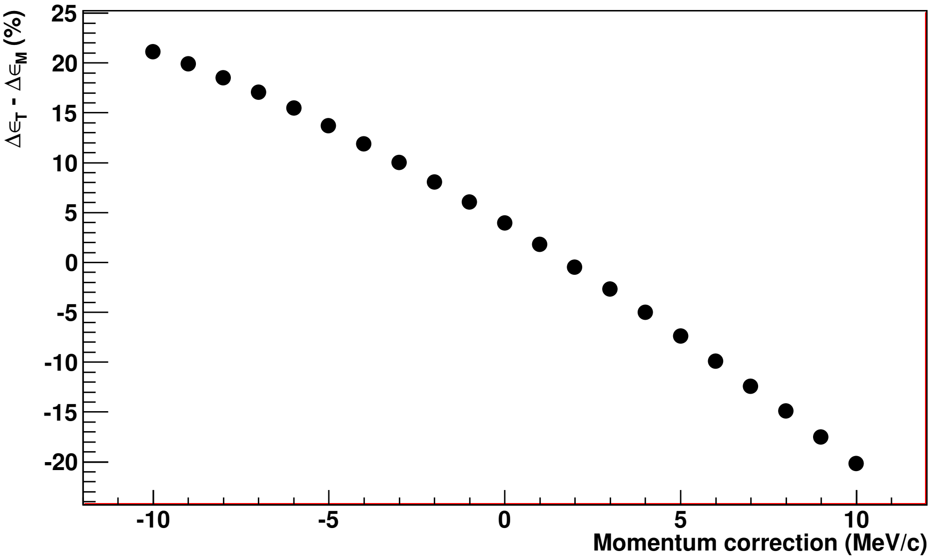

Figure 5.7: Effect of varying the momentum correction applied to each point after it has been mapped using a matrix found by the neural network.

Inflating the momentum distribution at the second reference plane in this way with knowledge of the true value is clearly a biased measurement. However it shows that a functional correction to the matrix prediction exists and if it could be found in an un-biased way by minimising the residual between the prediction and the true value in the second reference plane, would enable a prediction of emittance reduction through the MICE cooling channel without the need for a second measurement.

5.6 Conclusions

A method of finding the elements of a transfer matrix through Step IV of MICE has been developed using a neural network approach, taking the measured inputs and outputs of the cooling channel and finding the weights which cause each output to be a sum of weighted inputs. The optimum weights, which correspond to the matrix elements, were found using the network by minimising the error function of the ensemble of particles. For beams with very small longitudinal momentum spread the result was a matrix which could reasonably predict the position of the particle after the cooling channel, although with a resolution less than that of the trackers, and that required for an accurate measurement of emittance. As the longitudinal momentum spread increased, the resolution of the prediction, and hence the reliability of the matrix, became poorer, and a correlation between the two has been shown.

By slicing the beam into longitudinal momentum components and applying the algorithm to the sub-sets of particles, a functional dependance of the transfer matrix elements was found, and this could be used to apply the matrix method to beams with a high momentum spread.

A clear danger in this data fitting campaign is that the results, whilst reasonably accurate, may only be a fit to the data or limited to the current experimental setup. The periodic optics of the cooling channel go through only one repetition in Step VI, and it is difficult to extrapolate the particle oscillations from earlier phases to that of Step VI. Even with a fit to Step VI data, it cannot be assumed that a two cell cooling channel will have a transfer matrix equivalent to the multiplication of two Step VI matrices, or that quantities such as emittance will be predicted by them, without data from a cooling lattice larger than that of MICE.

The fitting of the linear, magnetic components of the phase space mapping is a required first step towards a complete model of the cooling channel and its effects. Additional work is needed to try to accommodate the material effects which are the key component of the ionisation cooling channel, and to find the relationship, if any exists, between multiple cells of the MICE cooling channel. Inclusion of the analytical interpretations of energy losses and multiple scattering may help, and could be expanded to account for the changes in the momenta as the beam passes through the material.

Chapter 6

Conclusions

Current experiments in the field of neutrino physics are providing ever more accurate measurements of the parameters of neutrino oscillation, and closing in on the limits of previously unknown parameters. In the event that θ_{13} still lies beyond the statistical limit of modern experiments, or to measure the values of the CP phases and mass hierarchies through many interaction channels, a Neutrino Factory is the ideal candidate facility. In order to build such a muon accelerator many fields of research must be opened up and technical challenges addressed. One such challenge lies in how to accelerate a beam of muons which has been produced with a large transverse emittance from interactions in the target. Ionisation cooling under investigation at the Muon Ionisation Cooling Experiment should provide a solution by reducing the transverse momentum of a muon beam in a time scale to fit the muon lifetime.

The muon beamline at MICE has been constructed and preliminary measurements of its characteristics have been made. In order to make a precision measurement of the emittance of the MICE beam, scintillating fibre trackers have been constructed to be placed inside superconducting solenoids. The readout electronics for the tracker were developed as part of this research and have been used in a cosmic ray test stand. To achieve the maximum light yield (and hence particle detection efficiency) in each plane of the trackers whilst minimising noise, a calibration procedure was developed and documented in this thesis. By varying the bias voltage applied to groups of the photon detection devices (VLPCs) and measuring the noise and light yield, where these are obtained by sampling the ADC in each channel when an external light has and has not been applied, the bias dependence of the noise and efficiency was found for each channel. As the channels are grouped an averaging algorithm was used to find the optimal biases based on defined parameters.

In addition to the optimisation of the analogue signal, timing and zero suppression features have been shown to operate and a calibration procedure prepared, although this has yet to be implemented due to availability of the detector when not taking cosmic ray test data.

Reconstruction algorithms were also implemented to find the optimum combination of overlapping fibres to produce a space point, and this also resulted in a measurement of internal misalignments. A Kalman filter has been written to fit the points to a particle trajectory based on a known uniform solenoidal field. Whilst the results of this Kalman filter indicated the tracker would meet the position and momentum resolution requirements, assumptions of field uniformity are optimistic. Once the spectrometer solenoid is completed and its field mapped, a tracking algorithm which takes into account local field inhomogeneities will need to be developed.

Based on cosmic ray data, the track position resolution, light yield and plane efficiency was measured and found to be consistent with expectations and design requirements. The MICE trackers are therefore expected to meet the criteria for making a precision measurement of emittance into, and out of, the MICE cooling channel.

With the phase space coordinates of each particle measured before and after the cooling channel the emittance reduction of the beam can be measured from the limited number of input points. As the channel contains RF cavities and the beam structure is significantly larger than the period of the RF wave, time-acceleration correlations appear in the beam. To expand on the limited number of input points, and to correct for the RF phasing, a weighting algorithm has been developed using approximate Voronoi diagrams. The viability of the method has been investigated, and then applied to simulated beams at operating modes of Steps IV and V.

The additional weighted points have been used to fit the expected theoretical prediction of cooling to simulated data, with good agreement between the prediction resulting from weighted sets and separate simulated points. Using a sinusoidal function to weight particles according to the likely level of their acceleration in the RF field, the effective performance of sustainable cooling was investigated in Step V. The weighting method has proven successful at extrapolating the cooling performance to varying input emittance, and characterising the performance of the re-accelerating channel.

As transfer matrices are a common form of transporting a particle's position and momentum between points in an accelerator lattice, an attempt was made to fit such a matrix based on the measurements in and out of the MICE cooling channel. A neural network type linear regression fit was implemented and found to produce a varying quality matrix, with the residual between the true value and the matrix prediction dependent on the longitudinal momentum and transverse emittance. The failings in the fit lay in the approximation, or lack thereof, of the stochastic processes of energy losses and multiple coulomb scattering. Using the matrices found by the fit to extrapolate the emittance through Step IV gave varying, but primarily quite poor results. As expected, further work into including the material effects in the fit would be required to improve this calculation.

Using particle physics detectors to measure the parameters of a muon beam on a single particle basis allows for an extremely accurate measurement not just of the primary effects of ionisation cooling, but the subtle behaviour in the lower statistics regions of the beam not normally accessible by macroscopic measurements. Some of these analyses have been investigated during the course of this thesis, and have hopefully highlighted the sensitivities in these regions which must be of concern in constructing a high flux muon beam for use in a collider or neutrino source.

The verification and characterisation of ionisation cooling at MICE will form just one component in any future muon accelerator. It is, however, still a key factor, and when combined with other technological advancements will allow for the muon to be used in high precision measurements at the energy frontier, and of physics beyond the standard model.

Appendix A

Additional Results

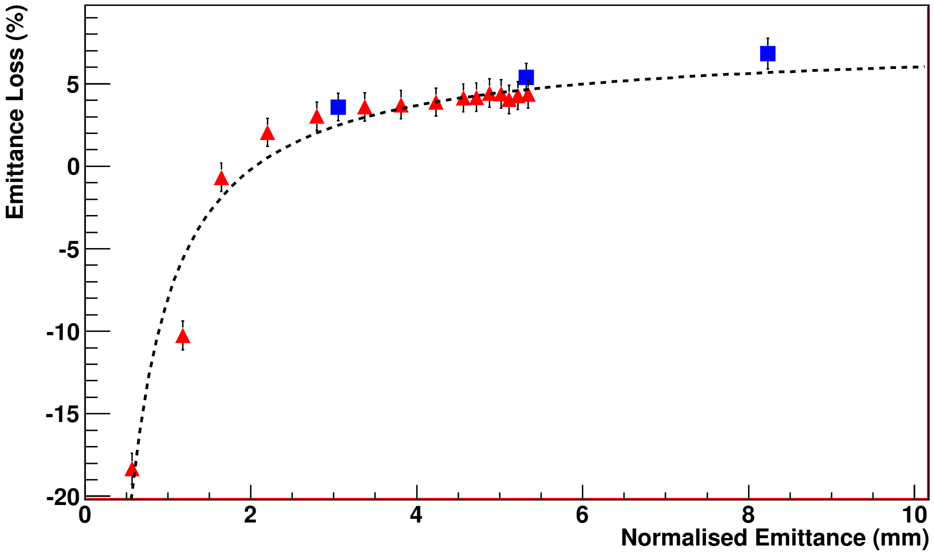

Figure A.1: Weighted fractional emittance loss through Step IV for various beams under RMS scaling (red triangles), weighted from a single original set ($\epsilon = 3\,\mathrm{mm}$) and compared to other simulated beams in a Step IV 200 MeV/c mode (blue squares).

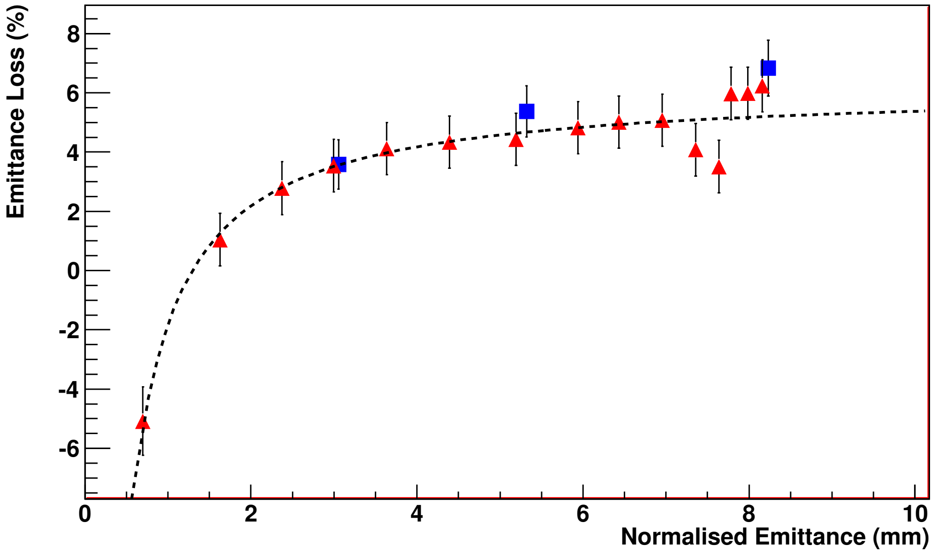

Figure A.2: Weighted fractional emittance loss through Step IV for various beams under RMS scaling (red triangles), weighted from a single original set ($\epsilon = 6\,\mathrm{mm}$) and compared to other simulated beams in a Step 4 200 MeV/c mode (blue squares).

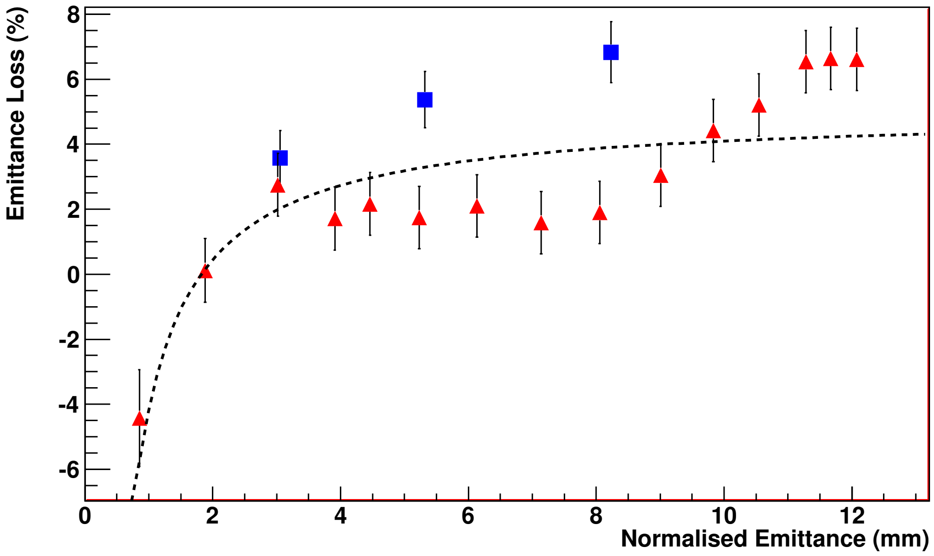

Figure A.3: Weighted fractional emittance loss through Step IV for various beams under RMS scaling (red triangles), weighted from a single original set ($\epsilon = 10\,\mathrm{mm}$) and compared to other simulated beams in a Step IV 200 MeV/c mode (blue squares).

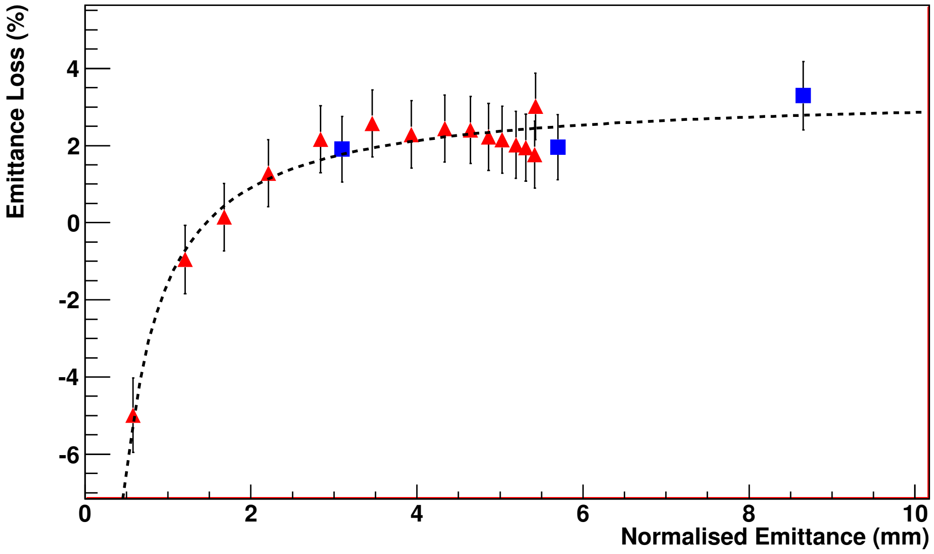

Figure A.4: Weighted fractional emittance loss through Step IV for various beams under RMS scaling (red triangles), weighted from a single original set ($\epsilon = 3\,\mathrm{mm}$) and compared to other simulated beams in a Step IV 240 MeV/c mode (blue squares).

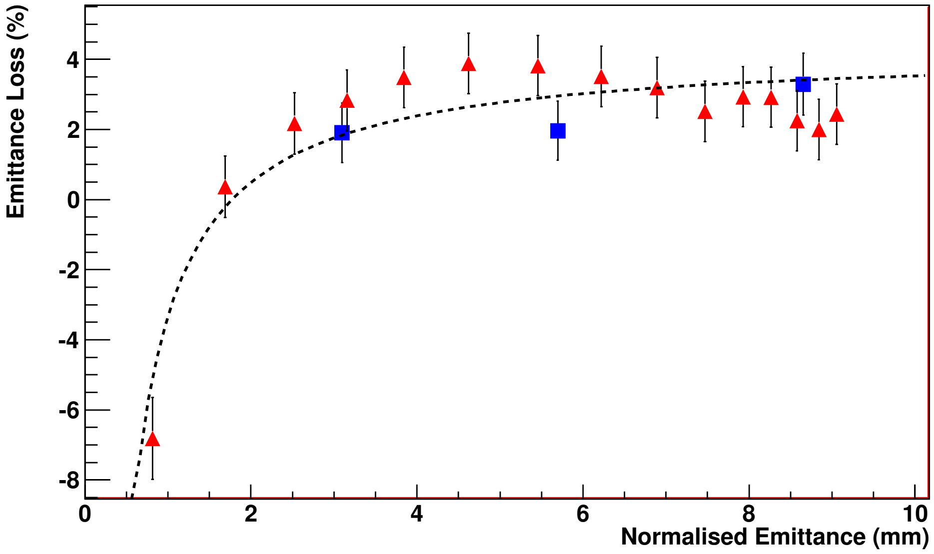

Figure A.5: Weighted fractional emittance loss through Step IV for various beams under RMS scaling (red triangles), weighted from a single original set ($\epsilon = 6\,\text{mm}$) and compared to other simulated beams in a Step IV 240 MeV/c mode (blue squares).

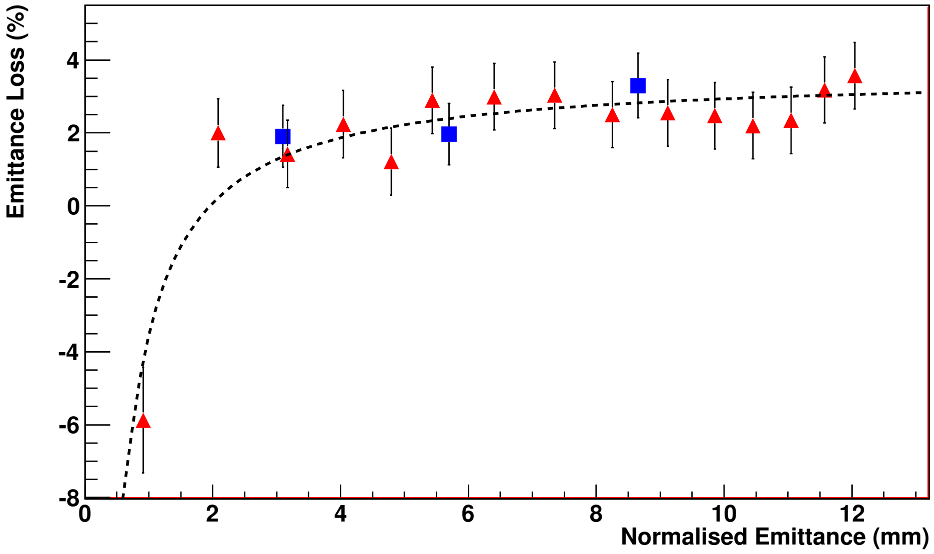

Figure A.6: Weighted fractional emittance loss through Step IV for various beams under RMS scaling (red triangles), weighted from a single original set ($\epsilon = 10\,\mathrm{mm}$) and compared to other simulated beams in a Step IV 240 MeV/c mode (blue squares).

Bibliography

[1] C. S. Wu et al., Experimental Test of Parity Conservation in Beta Decay, Physical Review, 105:4:1413-1415 (1957)

[2] C.L. Cowan, Jr., F. Reines, F.B. Harrison, H.W. Kruse and A.D. McGuire, Detection of the Free Neutrino: A Confirmation, Science 124:103 (1956)

[3] M. Goldhaber, L. Grodzin and A. W. Sunyar, Phys. Rev. 109:1015 (1958)

[4] E. J. Konopinski and H. M. Mahmoud, The universal Fermi interaction, Phys. Rev. 92:1045-1049 (1953)

[5] S. L. Glashow, Partial Symmetries of Weak Interactions, Nucl.Phys 22:579-588 (1961)

[6] S. Weinberg, A model of leptons, Phys. Rev. Lett. 19:1264-1266 (1967)

[7] A. Salam, Renormalization of Gauge theories, Phys. Rev. 127:331-334 (1962)

[8] G. Arnison et al, Physics Letters122B:103 (1983)

[9] G. Arnison et al, Physics Letters126B:398 (1983)

[10] K. Nakamura et al. (Particle Data Group), JP G 37, 075021 (2010)

[11] Z Maki, M. Nakagawa and S. Sakata, Remarks on the unified model of elementary particles, Progress of Theoretical Physics, 28:5:870-880 (1962)

[12] C. Guinti and C. W. Kim, Fundamentals of neutrino physics and astrophysics, Oxford University Press (2007)

[13] D. G. Michael et al., The magnetized steel and scintillator calorimeters of the MINOS experiment, Nucl. Instrum. Meth. A 596:190-228 (2008)

[14] P. Adamson et al., Improved search for muon-neutrino to electron neutrino oscillations in MINOS, Phys. Rev. Lett. 107.181802 (2011)

[15] T2K Collaboration, The T2K experiment, Nucl. Instrum. Meth. A 659:1:106-135 (2011)

[16] K. Abe et al., Indication of Electron Neutrino Appearance from an Accelerator-Produced Off-Axis Muon Neutrino Beam, Phys. Rev. Lett. 107 (2011)

[17] R. Acquafredda et al., The OPERA experiment in the CERN to Gran Sasso neutrino beam, JINST 4 (2009)

[18] A. Agafonova et al., Observation of a first ν_τ candidate event in the OPERA experiment in the CNGS beam, Physics Letters B 691:3:138-145 (2010)

[19] A. Tonazzo, The double Chooz detector, IPRD06 Proceedings (2006)

[20] Y. Abe et al, Indication for the disappearance of reactor electron antineutrinos in the Double Chooz experiment, arXiv:1112.6353v1 [hep-ex] (2011)

[21] The Daya Bay Collaboration, Daya Bay Proposal - A Precision Measurement of the Neutrino Mixing Angle θ_{13} Using Reactor Antineutrinos at Daya Bay, arXiv:0701029 [hep-ex] (2006)

[22] RENO Collaboration, RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang, arXiv:1003.1391v1 [hep-ex]

[23] G. L. Fogli et al., Evidence of $\theta_{13} > 0$ from global neutrino data analysis, arXiv:1106.6028 (2011)

[24] The ISS physics working group, Physics at a future Neutrino Factory and super beam facility, International scoping study of a future Neutrino Factory and super-beam facility (2007)

[25] K. Abe et al., Detectors and flux instrumentation for future neutrino facilities, International scoping study of a future Neutrino Factory and super-beam facility (2007)

[26] J. S. Berg et al., Accelerator design concept for future neutrino facilities, International scoping study of a future Neutrino Factory and super-beam facility (2007)

[27] K. T. McDonald et al., The MERIT high-power target experiment at the CERN PS, in proceedings of First International Particle Accelerator Conference, Kyoto Japan, May 23–28 2010, pg. 3527.

[28] A. Letchford et al., Status of the RAL front end test stand, IPAC10 Proceedings (2010)

[29] J.R.J. Bennett et al., Lifetime and strength tests of tantalum and tungsten under thermal shock for a Neutrino Factory target, Nucl. Instrum. Meth. A 646 (2011) 1-6.

[30] J.J.Back, Energy deposition studies for the Neutrino Factory target station, JINST 6 (2011) P06002.

[31] A. Moretti et al., Effects of High Solenoidal Magnetic Fields on RF Accelerating Cavities, PRSTAB 8, 072001, (2005)

[32] www.fnal.gov/projects/muon_collider/cool/cool.html

[33] R. Barlow et al., EMMA - The world's first non-scaling FFAG. Nucl. Instrum. Meth. A 624:1:1-19 (2010)

[34] N. Bliss et al., Technical description and status of the EMMA non-scaling FFAG, Proc. of Cyclotrons 2007, Giardini Naxos, Italy (2007)

[35] S. L. Smith, First commissioning results from the non-scaling FFAG accelerator, EMMA, Proc. of Cyclotrons 2010, Lanzhou, China (2010)

[36] G. I. Budker, An effective method of damping particle oscillations in proton and anti-proton storage rings, Atomic energy 22:5:438-440 (1967)

[37] D. Mohl, G. Pertrucci, L. Thorndahl, S. van de Meer, Physics and technique of stochastic cooling, Physics Reports, Vol 58 Iss 2 (1980)

[38] G. Ising, Prinzip einer Methode zur Herstellung von Kanalstrahlen hoher Voltzahl, Arkiv for matematik o. fysik 18:1-4 (1924)

[39] R. Wideroe, Uber ein neues Prinzip zur Herstellung hoher Spannungen, Arch. f. Electrot 21:387-406 (1928)

[40] S. Y. Lee, Accelerator Physics, World scientific publishing (2004)

[41] E. Wilson, Introduction to particle accelerators, Oxford University Press (2001)

[42] A. W Chao and M. Tigner, Handbook of accelerator physics and engineering, World scientific publishing (1999)

[43] G. Penn, Beam Envelope Equations in a Solenoidal Field, MuCool Note 71

[44] The MICE collaboration, The Beam Line of the Muon Ionization Cooling Experiment at ISIS: Construction and Performance, to be submitted to JINST (2012)

[45] C. Rogers et al., Study of the Proton Absorber in the MICE Beamline, MICE Note 294 (2010)

[46] J. H. Cobb, A compendium of MICE β functions, MICE Note 83 (2004)

[47] T. Carlise and J. Cobb, Ionization Cooling in MICE Step IV, Proceedings to the International Particle Accelerator Conference (2011)

[48] M. A. Green et al, Focussing solenoids for the MICE cooling channel, MICE Note 63

[49] C. Rogers, Effect of Reduced RF Gradient on MICE Cooling Performance, MICE Note 265

[50] D. Li et al, Normal conducting RF cavity for MICE, MICE Note 314

[51] Y. Karadzhov et al, TOF detectors time calibration, MICE Note 251 (2009)

[52] M. Adinolfi et al., The KLOE electromagnetic calorimeter, Nucl. Instrum. Meth. A 482 (2002)

[53] D. Lietti, M. Bari, D. Bolognini, P. Chimenti, G. Giannini, et. al., The prototype of the MICE Electron-Muon Ranger: Design, construction and test, Nucl. Instrum. Meth. A 604:314318 (2009)

[54] S. Agostinelli et al., Geant4 - A simulation toolkit, Nucl. Instrum. Meth A 506:250-303 (2003)

[55] J. Allison et al., Geant4 developments and applications, IEEE Transactions on Nuclear Science 53:1:270-278 (2006)

[56] M. Ellis, MICE Software Design and Physics Performance, Proceedings of the 10th ICATPP Conference on Astroparticle, Particle and Space Physics, Detectors and Medical Physics Applications 763-770 (2008)

[57] C. Tunnel, The MICE Analysis User Software, Proceedings of IPAC 2011 (2011)

[58] The ALICE data acquisition software framework DATE V5, 15th International conference on computing in high energy and nuclear physics (2006)

[59] Experimental Physics and Industrial Control System, www.aps.anl.gov/epics/

[60] A. J. Dobbs, Particle Rate and Host Accelerator Beam Loss on the MICE Experiment, PhD Thesis, Imperial College London (2011)

[61] M. A. Rayner. The development of a novel technique for characterizing the MICE muon beam and demonstrating its suitability for a muon cooling measurement. PhD thesis, The University of Oxford (2011)

[62] D. Adey, First measurements of emittance using particle physics detectors at MICE, Proceedings of the 11th International workshop on a future neutrino factory (2011)

[63] M. Ellis et al, The design, construction and performance of the MICE scintillating fibre trackers, Nucl. Instrum. Meth. A 659:136-153 (2011)

[64] B. Wang et al, The Design and Construction of the MICE Spectrometer Solenoids, IEEE Transactions on Applied Superconductivity 19:3

[65] DØ Collaboration, V. M. Abazov et al., The Upgraded DØ Detector, Nucl. Instrum. Meth. A 565:463-537 (2006)

[66] A. Bross, E. Flattum, D. Lincoln, S. Gruenendahl, J. Warchol, M. Wayne, and P. Padley, Characterization and performance of visible light photon counters (VLPCs) for the upgraded DØ detector at the Fermilab Tevatron, Nucl. Instrum. Meth. A 477:172-178 (2002)

[67] S. Takeuchi, J. Kim, Y. Yamamoto, and H. Hogue, Development of a high quantum-efficiency single-photon counting system, Appl. Phys. Lett. 74, 1063 (1999)

[68] S. L. Chuang, Physics of photonic devices, Wiley (2009)

[69] P. Rubinov, AFEII design specifications. Fermilab TM-2488-PPD

[70] J. Estrada, C. Garcia, B Hoeneison, P. Rubinov, MCM II and the Trip Chip, DØ Note 4009 (2002)

[71] Mil-Std-1553 Designer's guide, www.ddc.web.com

[72] S. Rapisarda et al., VLSB Specifications. Fermilab TM 2489-PPD.

[73] R. E. Kalman, A new approach to linear filtering and prediction problems, Transactions of the ASMEJournal of Basic Engineering, 82 (Series D): 35-45 (1960)

[74] H. W. Sorenson, Least-squares estimation: from Gauss to Kalman, IEEE Sptectrum 7:63-68 (1970)

[75] A. Cervera-Villanueva et al, Kalman filter tracking and vertexing in a silicon detector for neutrino physics, Nucl. Intrum. Meth. A:486:639-662 (2002)

[76] M. Ellis et al, Simulated performance of the MICE trackers, MICE Note 90

[77] A. Cervera-Villanueva, J. J. Gomez-Cadenas, and J. A. Hernando, RecPack: A reconstruction toolkit, Nucl. Instrum. Meth. A 534:180-183 (2004)

[78] G. Voronoi, Nouvelles applications des parametres continus a la theorie des formes quadratiques, Journal fur die Reine und Angewandte Mathematik, 133:97-178 (1909)

[79] B. Delaunay, Sur la sphre vide. A la mmoire de Georges Vorono, Bulletin de l'Acadmie des Sciences de l'URSS, Classe des sciences mathmatiques et na 6:793800 (1934)

[80] C.T. Rogers, Statistical weighting of the MICE Beam, Proceedings of the 11th European Particle Accelerator Conference (2008)

[81] S. Fortune, A sweepline algorithm for Voronoi diagrams, Algorithmica 2:153-174 (1987)

[82] J. H. Cobb, Statistical errors on emittance measurements, MICE Note 268 (2009)

[83] J. H. Cobb, Statistical errors on emittance and optical functions, MICE Note 341 (2011)

[84] C. Peterson, Neural networks in high energy physics, Proceedings of the conference on Computing in high energy physics (1991)