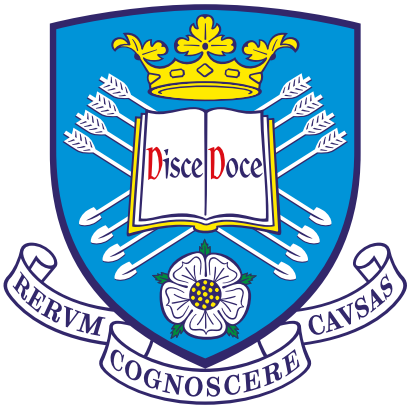

Studies and Developments within the Muon Ionisation Cooling Experiment

Edward Overton

Supervisor: Dr. Chris Booth

A Thesis Submitted for the Degree of Doctor Of Philosophy in Physics

Department of Physics and Astronomy
University of Sheffield, September 2014

Abstract

The study of neutrino physics has developed significantly during the past decades. A neutrino factory promises unmatched precision in studies of the CP violating phase in neutrino oscillation. In addition the facility paves the path for other unique technologies, such as a muon collider. In order to construct such a facility a number of experiments must successfully demonstrate viability. The Muon Ionisation Cooling Experiment (MICE) is one of these.

To demonstrate successful cooling MICE contains a cooling cell sandwiched between instrumentation. A single particle measurement technique is adopted, where the emittance of each particle is measured using tracking Spectrometer Solenoids. The single particles are generated parasitically within the ISIS Synchrotron at the Rutherford Appleton Laboratory.

This thesis has three topics, the first topic focuses on mechanical diagnostics for the MICE target, which ultimately allow an indirect friction measurement. The developed diagnostics are then applied to the quality assurance procedure, to ensure reliable targets are installed on ISIS.

The next topic is the Tracker, which forms a central component of the spectrometer system. In order for the tracker to operate effectively, a system to synchronise the front end electronics gate to particle arrival has been built and tested. In addition a LED calibration system has been designed and installed.

The final topic in this thesis contributes to the alignment of the MICE beamline, using a beam based alignment technique. The technique was capable of studying a number of quadrupoles and the TOF1 detectors alignment, using data collected by the TOF system.

Acknowledgements

I would like to thank my supervisors Chris Booth and Paul Hodgson for the insight, direction and inspiration given, and especially for their comments while writing this thesis. In addition Paul Smith and John McMillan have been invaluable in developing my practical laboratory skills, electronics in particular.

Many thanks go to the MICE collaborators which I have had the fortune of working with, your friendliness, encouragement and criticism is graciously accepted. A special thanks go to Adam Dobbs and Craig MacWaters for a cheerful face even on the rainiest of days.

Finally, this thesis would not have been possible without the support of my family and friends.

Cheers!

Declaration

Chapter one is a brief summary of neutrino oscillation and the neutrino factory concept, taken entirely from available publications and other sources on relevant topics.

Chapter two is an overview of ionisation cooling in the context of MICE, the material for which originates from publications, internal documents to MICE and discussions with collaborators.

Chapter three focuses on the MICE target mechanism, where section 3.1 describes the design of the mechanism and again stems from internal documents and discussions. Sections 3.2 and onwards are mostly my own work, the exception being the "width" and "accelerating" diagnostics which were suggested by Paul Hodgson.

Chapter four focuses on the preparation of the MICE fibre tracker. Section 4.1 describes the tracker from internal documents and discussions. Section 4.2 focuses on synchronisation of MICE to ISIS bunches, a study that was suggested by Daniel Kaplan, however all the following work is entirely my own. Section 4.3 addresses the need for a calibration system, the requirements for which were developed from discussion with the tracker group and the design was inspired from discussions with John Macmillan. Again the resulting design and construction is entirely my own work.

Chapter five originates from discussions with Chris Rogers, who planned the initial data taking runs. The following analysis of this data is my own work.

Contents

List of Tables

This page is intentionally left blank.

List of Figures

This page is intentionally left blank.

Chapter 1

Introduction

1.1 Neutrino Physics

1.1.1 History

The concept of a neutrino was hypothesised by W. Pauli[1] to explain the missing momentum and continuous energy spectrum observed from beta decay. Pauli postulated that the neutrino should be light, electrically neutral, spin $\frac{1}{2}$ particle and was an unseen component of the decay. Pauli described this as a "very terrible thing", because the postulated particle seemed undetectable.

To detect such an impossible particle Cowan and Reines first suggested placing a detector close to a nuclear explosion, which was never realised. In 1956 the pair used anti-neutrinos from the Savannah River nuclear reactor in the interaction:

$$\overline{\nu}_{\mathrm{e}} + \mathrm{p} \rightarrow \mathrm{e}^{+} + \mathrm{n} \tag{1.1}$$

and made the first direct detection of an anti-neutrino[2]. In 1955 Davis[3] had failed to observe the interaction:

$$\overline{\nu}_{\mathrm{e}} + {}^{37}\mathrm{Cl} \rightarrow \mathrm{e}^{-} + {}^{37}\mathrm{A} \tag{1.2}$$

which identified the anti-neutrino as a distinct antiparticle of the neutrino.

In 1960 Pontecorvo questioned if the neutrinos emitted from beta decay were the same particle as those in pion decay. This was investigated in the first neutrino beam experiments[4], where neutrinos originating from pion and kaon decay were studied. Muons from the decays were absorbed by substantial shielding and neutrino

interactions were observed in a spark chamber, from the interactions:

$$\nu_\mu + \mathrm{n} \rightarrow \mu^- + \mathrm{p} \qquad \overline{\nu}_\mu + \mathrm{p} \rightarrow \mu^+ + \mathrm{n} \tag{1.3}$$

$$\nu_\mathrm{e} + \mathrm{n} \rightarrow \mathrm{e}^- + \mathrm{p} \qquad \overline{\nu}_\mathrm{e} + \mathrm{p} \rightarrow \mathrm{e}^+ + \mathrm{n} \tag{1.4}$$

The experiment noted an absence of electron showers and demonstrated that neutrinos were different particles ($\nu_\mu \neq \nu_\mathrm{e}$). Following the discovery of the τ lepton in 1975, the corresponding neutrino (ν_τ) was directly observed by the DONuT collaboration[5] in 2000. Finally studies of the Z^0 resonance peak width at LEP measured the number of light neutrino generations as $\mathrm{N}\nu = 2.92 \pm 0.05$ [6], indicating that all the light neutrinos (where $2m_\nu < m_\mathrm{Z}$) are accounted for.

Parity Violation in Weak Interactions

One of the many interesting discoveries from the early experiments on beta decay is parity violation. Under a parity transformation all the spatial axes are inverted about the origin. Parity is conserved if the probability of a process occurring is unchanged by the transformation. This was assumed to be the case for the weak interaction, however Lee and Yang were unable to find any supporting evidence[7].

In 1957 Wu[8] prepared a sample of ^{60}Co, which had the nuclear spin polarised using an external magnetic field. The counting rate of electrons emitted both parallel and anti-parallel to the polarisation of the atoms was measured. The results showed a clear asymmetry between the rates and emission of electrons preferentially occurred anti-parallel to the nuclear spin; careful analysis indicated that parity is maximally violated. One year later the Goldhaber[9] experiment showed that neutrinos have spin anti-parallel to their momentum (left handed), while anti-neutrinos have spin parallel to momentum (right handed).

Solar Neutrino Puzzle

From the developed understanding of beta decay and the development of the solar model, it was possible to both predict and measure the solar neutrino flux at the earth. The Homestake experiment[10], took place in a gold mine and used 380 m^3 of $\mathrm{C_2Cl_4}$, where Chlorine is a target for ν_e. The number of detected interactions was averaged over the lifetime of the experiment and yielded 1/3 of the expected rate.

To provide an independent measurement of the neutrino flux, the Kamiokande-II[11] detector was re-purposed to observe Cherenkov radiation from recoil electrons

that were scattered by neutrinos from the sun. Cherenkov radiation provided a real-time measurement technique that included directional information. The results clearly correlated the source of neutrinos as solar, however the rate was 1/2 of that predicted by the solar model.

A hypothesis which explained these results was that the ν_{e} neutrinos were oscillating into ν_{μ} and ν_{τ} flavours, which were not being (fully) detected. This was confirmed by the Sudbury Neutrino Observatory (SNO) which used a 1 kt D_2O detector and could differentiate between three interactions:

$$\nu_{\mathrm{e}} + \mathrm{d} \rightarrow \mathrm{p} + \mathrm{p} + \mathrm{e}^{-} \qquad \text{(Charged Current)} \tag{1.5}$$

$$\nu_{\mathrm{x}} + \mathrm{d} \rightarrow \nu_{\mathrm{x}} + \mathrm{p} + \mathrm{n} \qquad \text{(Neutral Current)} \tag{1.6}$$

$$\nu_{\mathrm{x}} + \mathrm{e}^{-} \rightarrow \nu_{\mathrm{x}} + \mathrm{e}^{-} \qquad \text{(Elastic Scattering)} \tag{1.7}$$

The Charged Current interaction was only attainable by ν_{e}, the Neutral Current interaction was equally achievable by all flavours, while Elastic Scattering had a reduced sensitivity to ν_{μ}, ν_{τ}. The rate of neutral current interactions agreed with the solar model, yet the charged current rate was again one third of the expectation, providing direct evidence of neutrino oscillation[12].

1.1.2 Neutrino Oscillations and Mixing

The theory of neutrino oscillation developed by Pontecorvo, Maki, Nakagawa, Sakata is founded on the principle that the flavour eigenstates of the weak interaction, ν_{α} (where $\alpha = \mathrm{e}, \mu, \tau$) are not mass eigenstates, ν_{k} (where $k = 1, 2, 3$). The flavour eigenstates are then considered a mixture of the mass eigenstates[13]:

$$|\nu_{\alpha}\rangle = \sum_{k} U^{*}_{\alpha k} |\nu_{k}\rangle \tag{1.8}$$

where U is the PMNS matrix:

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \tag{1.9}$$

and $c_{ab} = \cos\theta_{ab}$, $s_{ab} = \sin\theta_{ab}$. The parameters θ_{12}, θ_{13} and θ_{23} are mixing angles, while δ is a CP violating phase. Note that if the neutrino is Majorana (self antiparticle) then an additional two CP violating phases are introduced by means of

a diagonal matrix.

The mass eigenstates are the eigenstates of the Hamiltonian, the time dependence of which can be found by solving the Schrödinger equation. In the case of vacuum oscillation, plane waves solutions are adopted:

$$|\nu_k(t)\rangle = e^{-iE_k t}|\nu_k\rangle \tag{1.10}$$

where $E_k = \sqrt{\mathbf{p}^2 + m_k^2}$ is the energy of the eigenstate, with mass m_k and momentum $\mathbf{p}$. In order to measure the parameters of U an established technique is to study the probability of flavour change. The probability (P) of a neutrino state to change from an initial flavour(α) to final flavour (β) as a function of time is:

$$P_{\alpha\to\beta}(t) = |\langle\nu_\beta|\nu_\alpha(t)\rangle|^2 \tag{1.11}$$

$$= |\sum_k U^*_{\alpha k} U_{\beta k} e^{-iE_k t}|^2 \tag{1.12}$$

$$= \sum_{k,n} U^*_{\alpha k} U_{\beta k} e^{-iE_k t} U_{\alpha n} U^*_{\beta n} e^{+iE_n t} \tag{1.13}$$

The following relativistic approximations are then used:

1. $t = L$, where L is the distance travelled
2. $E_k - E_n \approx \dfrac{\Delta m^2_{kn}}{2E}$, where $\Delta m^2_{kn} = m^2_k - m^2_n$ and $E = |\mathbf{p}|$

to enable equation 1.13 to be expressed in terms of distance and energy, as follows

$$P_{\alpha\to\beta}(L,E) \approx \sum_{k,n} U^*_{\alpha k} U_{\beta k} U_{\alpha n} U^*_{\beta n} \exp\left(-i\frac{\Delta m^2_{kn} L}{2E}\right) \tag{1.14}$$

In many cases a two neutrino mixing is considered, where $U = R(\Theta)$, the two dimensional rotation matrix. This yields the two dimensional oscillation probability equation:

$$P_{\alpha\to\beta,\alpha\neq\beta}(L,E) = \sin^2(2\Theta)\sin^2\left(1.27\frac{\Delta m^2[\mathrm{eV}^2]L[\mathrm{km}]}{E[\mathrm{GeV}]}\right) \tag{1.15}$$

1.2 Neutrino Oscillation Experiments

The study of neutrino mixing utilises a known neutrino source, followed by a distant detector, at which the neutrino flux is measured. Where possible experiments prefer to

use a detector close to the source to measure the neutrino flux before the oscillation has occurred. A key component of the experiment is the baseline and energy, commonly referred to as L/E, which determines how sensitive the experiment is to oscillation. In reality a number of complimentary experiments study neutrino oscillation from a variety of sources. Some of these experiments are briefly described below.

1.2.1 Atmospheric

Cosmic rays interacting in the upper atmosphere generate a shower of particles, a substantial fraction of which are pions. The pions decay into muons, which then decay into electrons. A number of neutrinos are emitted during these decays, which can then be studied by atmospheric neutrino experiments.

One of the key experiments to study atmospheric neutrinos was the Super-Kamiokande detector, a 50 kt water Cherenkov, instrumented by in excess of 10,000 photomultiplier tubes. By using the direction of the observed particles, the detector was able to estimate the zenith angle of the neutrino and calculate the baseline of each event. With the energy of the interaction also available the experiment observed ν_μ disappearance at $L/E = 500$ km/GeV [14].

1.2.2 Reactor

Another substantial source of neutrinos are power reactors, which generate approximately $2{\times}10^{20}$ $\overline{\nu}_\mathrm{e}$/s/GW, with energies around 3 MeV [13]. The low energy restricts the studies to $\overline{\nu}_\mathrm{e}$ disappearance, however it also significantly reduces the necessary baseline distance required, ultimately increasing the neutrino flux.

The current generation of reactor experiments include: Daya Bay[15], Double-CHOOZ[16] and RENO[17]. Each of the experiments utilises a near detector to reduce systematic error on the neutrino flux from the reactor. In all experiments a deficit of $\overline{\nu}_\mathrm{e}$ has been observed at the far reactor, consistent with neutrino oscillation.

1.2.3 Accelerator

Neutrinos generated in accelerator sources nominally use a high energy (several GeV) proton beam which collides with a fixed target. The interaction of the beam generates secondary particles, including pions and kaons, which are commonly focused by magnetic horns. The in flight decay of the secondary (and tertiary) particles then yields a directional neutrino beam.

The tertiary decay process generates a neutrino beam with a large energy spectrum, over a broad range of angles. The complexity of the interactions requires a near detector to precisely characterise the neutrino beam and ultimately reduce systematic uncertainty in the results. A number of accelerator experiments are in operation, including:

T2K

The Tokai to Kamioka experiment in Japan uses a 30 GeV proton beam incident on a graphite target to generate neutrinos (mainly ν_μ) destined for the Super-Kamiokande far detector. The near and far detectors are placed at 280 m and 295 km respectively. Both detectors are offset by 2.5° from the neutrino beam axis, in order to substantially reduce the energy spread of the observed neutrinos.

In 2013 the T2K experiment delivered results that demonstrated ν_e appearance to $> 7\sigma$ [18], showing $\sin^2\theta_{13}$ to be non-zero.

MINOS (MINOS+)

The Main Injector Neutrino Oscillation Search[19] utilises the NuMi beam-line at Fermi National Accelerator Laboratory, Chicago to send neutrinos to the Soudan mine, Minnesota. The NuMi beam-line accepts a 0.4 MW 120 GeV proton beam from the Main Injector, which is to be upgraded to 0.7 MW for future experiments.

The MINOS detector is on axis and as a result has a broad neutrino energy spectrum, in the range of 1 to 30 GeV. The far detector is a magnetised 5.4 kt of alternating steel plate and scintillator strips, located 735 km from the target. The near detector is functionally equivalent to the far detector, but significantly lighter at 0.98 kt and located 1 km from the target.

The MINOS experiment has now been superseded by MINOS+, which has begun operating in the upgraded NuMi beamline.

NOνA

The Numi Off-axis ν_e Appearance[20] experiment shares the NuMi beamline at Fermi National Accelerator Laboratory and features a 14 kt far detector located 810 km from Fermi National Accelerator Laboratory, with a 0.3 kt near detector. Both detectors are 0.8° off axis, producing an neutrino energy spectrum with a peak at 2 GeV.

The detectors are constructed from rows of tubes of extruded plastic, with enhanced reflectivity and filled with liquid scintillator. The direction of each row is perpendicular to both the neutrino beam and the previous row. The main objective of the experiment is to provide a high precision measurement of $\sin(\theta_{13})$, Δm^2_{23} and studies of δ.

OPERA

The Oscillation Project with Emulsion-tRacking Apparatus uses neutrinos produced by the CERN Neutrinos to Gran Sasso (CNGS) beamline, which provides a 0.51 MW, 400 GeV proton beam. The far detector utilises 1.2 kt array of emulsion film interleaved with 1 mm thickness lead, formed into bricks. In addition a number of scintillators and tracking elements between the bricks provide real time reconstruction and energy measurements.

The OPERA detector has observed four ν_τ appearances from a ν_μ beam, which demonstrates ν_τ appearance with 4.2σ significance [21].

1.2.4 Experimental Results and Outlook

From the various studies of neutrino oscillation the parameters have been measured, shown in table 1.1. Experiments presently in operation will continue to improve the accuracy to which we know the mixing angles.

Both Δm^2_{21} and $|\Delta m^2|$ have been determined, however there is insufficient data to determine the mass hierarchy (Normal: $m_1 < m_2 < m_3$, or Inverted: $m_3 < m_1 < m_2$). The result is expected in the near future from the Noνa and T2K experiments.

Finally the relatively large $\sin^2\theta_{13}$ may allow an early measurement of δusing the current generation of experiments. However this parameter is likely to be determined accurately by the next generation of long baseline experiments.

1.3 Next Generation Experiments

The next generation of long baseline neutrino experiments are tasked with the accurate determination of δ, for which a number of experiments have been proposed:

T2HK: Tokai to Hyper-Kamiokande, operates over the same baseline as T2K, however a new 990 kt (560 kt fiducial) water Cherenkov detector is to be con-

Parameter	best-fit (1σ)	units
Δm^2_{21}	$7.54^{+0.26}_{-0.22}$	$10^{-5}\,\mathrm{eV}^2$
$\lvert\Delta m^2\rvert$	$2.43^{+0.06}_{-0.10}$	$10^{-3}\,\mathrm{eV}^2$
$\sin^2\theta_{12}$	$0.307^{+0.018}_{-0.016}$	
$\sin^2\theta_{23}$	$0.386^{+0.024}_{-0.021}$	
$\sin^2\theta_{13}$	0.0241 ± 0.0025	

Table 1.1: Global best fits of neutrino oscillation, from [6]. The definition of Δm^2 used is: $\Delta m^2 = m^2_3 - (m^2_1 + m^2_2)/2$.

structed. The accelerator complex at Tokai would also be upgraded to deliver a 0.75 MW proton beam.

LBNO was the proposed very Long Baseline Neutrino Oscillation experiment which could use a 100 kt liquid argon time projection chamber as a far detector. A 0.8 MW proton beam would produce a neutrino beam, that could operate over a 2.3 Mm baseline.

LBNE is the Long Baseline Neutrino Experiment, using a 2.3 MW proton beam to deliver neutrinos over a 1.3 Mm baseline to a 34 kt liquid argon detector.

ESS could utilise the European Spallation Source proton linac to deliver a 2 GeV, 5 MW proton beam. The far detector would be a water Cherenkov with $\approx$ 500 kt fiducial mass over a 360 km baseline.

Daeδlus is a novel design, where a fixed 300 kt Gadalinium-doped water Cherenkov detector is used. The neutrino source is from a stopped pion decay chain, producing ν_e, ν_μ, $\overline{\nu}_\mu$, with very little $\overline{\nu}_\mathrm{e}$. The pions are produced from an intense cyclotron source, which can be moved between different baselines: 1.5 km, 8 km and 20 km.

NF the Neutrino Factory would utilise decays of 10 GeV μ^-, μ^+ in storage rings, generating $10^{21}/10^7\,\mathrm{s}^{-1}$ muon decays to produce high purity neutrino beams. The neutrinos would be studied over 1.5 Mm to 2.5 Mm baseline, using a 100 kt magnetised iron detector.

nuStorm neutrinos from STORed Muons, would produce neutrino beams from an un-accelerated muon storage ring.

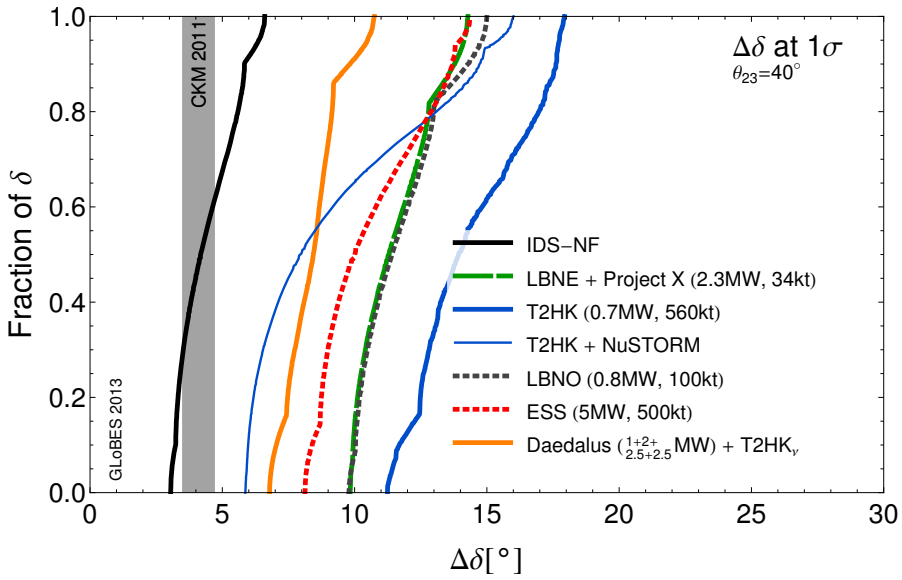

Figure 1.1: 1σ resolution of next generation facilities to measure the CP violating phase, δ [22]. Note that the fraction of δ is a fraction of π, and is displayed on the vertical axis.

Figure 1.1 shows a comparison of the next generation facilities' ability to measure the CP violating phase δ, now that $\sin^2\theta_{13}$ has been measured. The most precise of these is the Neutrino Factory, which could determine δ with a precision of a few degrees.

None of the proposed next generation neutrino facilities is trivial to design and construct. The neutrino factory, with the need to efficiently deliver a 10 GeV muon beam is no exception. The remaining pages of this chapter will give an overview of such a facility, and how this technology can apply to the energy frontier in particle colliders.

1.4 Muon Accelerators

1.4.1 Neutrino Factory

Production of a accelerated muon beam is a task which requires many steps, illustrated by figure 1.2. Due to the short lifetime of the muons the first part of the facility is dedicated to efficient muon generation. The following steps focus on efficiently transporting the muons to the decay rings, and due to the limited muon lifetime must be incredibly rapid. The decay rings aim the muons towards the destination detector,

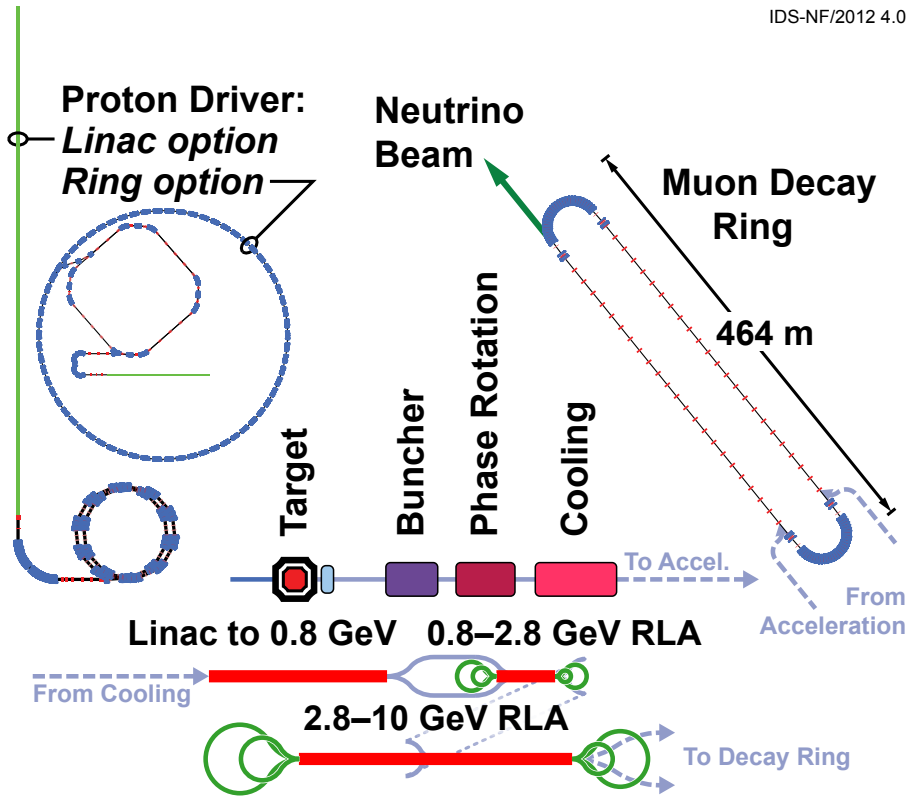

Figure 1.2: Schematic sketch of a neutrino factory. [22]

and maximise on-target decays by featuring long straight sections.

The following subsections describe each of the functional components of a neutrino factory, as found in [23].

Proton Driver

The first element in the neutrino factory is the proton driver, which must generate short (1 ns to 3 ns) proton bunches, at 8 GeV with a beam power of 4 MW. The initial acceleration technique depends drastically upon available infrastructure at the target laboratory. A neutrino factory constructed at either CERN or FNAL would use superconducting linacs to achieve the final energy, while a factory at RAL would use a combination of a linac and Rapid Cycling Synchrotron to deliver the final energy.

The final requirement of the proton driver is to deliver short bunches, and this is achieved using a compressor ring. In essence this step shortens the spread in time by increasing the energy spread of the beam.

Target

The target in the neutrino factory must be capable of accepting the large beam power, which produces significant thermal shock and requires substantial cooling. A liquid Mercury jet is used for the target, which was demonstrated by the MERIT[24] experiment at CERN.

Interactions within the target produce a range of particles, including π^+,π^-. In order to simultaneously capture both charges a 20 T solenoidal field is used, which is reduced to 1.5 T over a distance of 15 m. During this time a significant fraction of the π^+ (π^-) decay to μ^+ (μ^-).

Longitudinal Drift

Following the target is a long (40.8 m) drift section, which maintains the 1.5 T solenoidal field. As particles traverse this section a time-energy correlation develops, which can be exploited by the phase rotation step, further downstream.

Adiabatic Bunching

The next task is to group the muons into periodic bunches, which can then be efficiently manipulated by Radio Frequency(RF) cavities. This task is completed by using a number of RF cavities that provide a varying amount of acceleration to muons depending on the muon position and cavity phase. The varying acceleration causes the muons to bunch.

In the neutrino factory this is accomplished using a number of different frequency RF cavities, where frequency begins at 358.92 MHz and decreases to 235.4 MHz. Note that since the μ^- and μ^+ have opposite charge, they are simultaneously bunched into alternating bunches of μ^- and μ^+.

Phase Rotation

The phase rotation exploits the time-energy correlation developed in the longitudinal drift step. Here a number of de-phased cavities apply a decelerating field to the first bunch, while an accelerating gradient is applied to the last. This substantially decreases the bunch to bunch energy spread, and therefore reduces the overall energy spread of the beam.

Cooling

The muons in the beamline are a secondary product of the initial protons, and as such have a large transverse phase space. The cooling step of the process rapidly reduces the phase space of the beam, in order to maximise the fraction of particles which are accepted by the RF acceleration system. The Muon Ionisation Cooling Experiment (MICE) was proposed to test the cooling necessary for a neutrino factory, and is presented in chapter 2 . Note that the neutrino factory cooling design has diverged from that constructed in MICE, however MICE still presents a viable facility to test different cooling materials.

At the end of the cooling step, approximately 0.066 good μ^- are captured per 8 GeV proton on target, for a neutrino factory.

Acceleration

Acceleration of the muon beams to the final energy is handled in a number of linacs, each of which combine RF cavities and a solenoidal field to focus the beam. The first linac accelerates the beam to 0.8 GeV in a single pass; due to the μ^-, μ^+ phasing both species are conveniently accelerated simultaneously.

Next, a recirculating "dog-bone" linac accelerates the muons to 2.8 GeV in 4.5 passes. As the beam exits the linac a dipole field bends the particles into a return arc. Due to different beam momenta after each pass a different return arc for each energy is required. The energy range each recirculating linac can accept is restricted by the accumulation of error between the ideal phase of the cavity for acceleration, and the phase when the bunch arrives.

Finally a second recirculating linac performs the final acceleration to 10 GeV, using a similar technique to the first.

Decay Rings

The final component in the neutrino factory complex is the decay rings, in which the muons circulate and decay, producing a neutrino beam. The decay ring is shaped as a racetrack, containing two straight sections. Only decays occurring in a single one of these straights will produce a beam towards the target detector, with approximately 35.5% of muon decays on target. The efficiency can be improved by adjusting the geometry of the racetrack, to a maximum value $< 50\%$, at extra cost. In order to target a far detector 2 Mm from the facility the decay ring enters the ground at $9\,^\circ$.

1.4.2 Muon Colliders

In high energy particle collider experiments, e^- e^+ machines provide the cleanest events, since the colliding particles are fundamental. This enables the exact interaction energy to be known and tunable (to exploit resonances).

One key constraint in the design of e^- e^+ colliders is the emission of synchrotron radiation, which increases with relativistic γ^4, but decreases with bend radius. Due to the light mass of the electron this restricts the minimum bend radius and maximum energy which can be achieved. To mitigate this effect, a linear collider can be built, such as the International Linear Collider(ILC), presently in the design phase.

Since muons have a much higher mass, γ is significantly smaller at the TeV energy range and they thus emit negligible synchrotron radiation. This allows the use of compact circular muon colliders[25]. The development of some of the key technologies for a neutrino factory are shared with that of such a collider, in particular the proton driver, target and initial cooling.

The complete construction of a muon collider requires a significantly more advanced cooling scheme and a larger acceleration complex. In addition muon decays in flight will provide a substantial source of electrons, which will appear as background in particle detectors instrumenting the interaction point.

1.5 Outlook

The initial postulation of the neutrino appeared an impossible particle to detect, however through the construction of ever bigger detectors, reactors and accelerators the neutrino has been studied with increasing accuracy. These studies have revealed the neutrino oscillation phenomenon and begun the process of measuring the PMNS matrix.

A neutrino factory produces a clean neutrino beam, however a number of key technologies must be developed and tested. As part of these R&D efforts the MICE experiment was conceived to demonstrate and study ionisation cooling. Beyond the neutrino factory, muon storage rings present an interesting prospect for collider studies.

The following chapter in this thesis describes the concept of ionisation cooling in the context of the MICE experiment.

This page is intentionally left blank.

Chapter 2

The Muon Ionisation Cooling Experiment

The Muon Ionisation Cooling Experiment (MICE) aims to demonstrate that ionisation cooling is a feasible technology for reducing the phase space of a muon beam. The development of such a technology is crucial for a future Neutrino Factory or Muon Collider. In order to achieve this goal MICE must construct, operate and characterise a cooling channel.

MICE is hosted at the Rutherford Appleton Laboratory in the UK, where a dedicated beamline has been constructed to deliver particles to MICE from the ISIS Proton Synchrotron. The experiment is constructed in stages, which allows each part of the experiment to be tested before moving forward to the next.

This chapter will briefly introduce beam emittance and the principle of ionisation cooling, followed by the design of the MICE cooling channel, beamline and detectors.

2.1 Emittance

Within a beamline each particle can be described in a six dimensional space. The trace space is defined as (x, x', y, y', t, E) and can be divided into three two-dimensional spaces. There are two transverse spaces (x, x') and (y, y'), plus a longitudinal space (t, E), where $x' = dx/dz$, $y' = dy/dz$ and z is defined along the beamline.

In an ideal beamline, with only conservative processes, transverse particle motion can be described using Hills differential equations, the solution to which is known

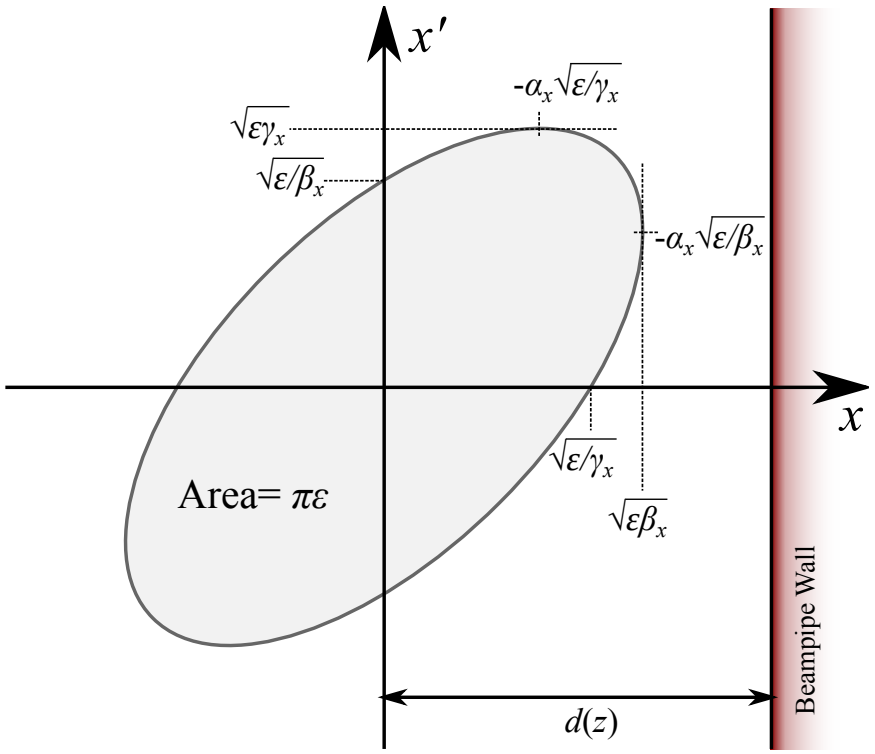

Figure 2.1: Trace space ellipse, labelled with Twiss parameters. An example transverse constraint is also present, in the form of the beam pipe wall.

as the Courant-Snyder invariant[26]:

$$\gamma_x x^2 + 2\alpha_x x x' + \beta_x x'^2 = \epsilon \tag{2.1}$$

where α_x, β_x and γ_x functions are known as the Twiss parameters and are a function of the position along the beamline (z). For clarity the Twiss parameters will always be denoted with a subscript, however relativistic parameters will remain unaccompanied. The equation describes a ellipse in trace space, shown in figure 2.1 with an area $\pi\epsilon$ and is referred to as the geometrical emittance. Since the area of the ellipse is explicitly declared, the equation is over-constrained and the Twiss parameters are related by:

$$\beta_x \gamma_x - \alpha_x^2 = 1 \tag{2.2}$$

Liouville's theorem states that the volume in phase space is constant relative to time, providing the particles obey the canonical equations of motion. This is generally true for accelerators, where large fields bend and focus the beam. However, in reality there

also many non-conservative forces, such as space charge and multiple scattering, which act to increase the emittance of a particle beam. Beam cooling is a term given to techniques which can reduce the emittance of a beam.

2.1.1 Acceptance

In complement to emittance, the acceptance is defined as the maximum emittance particle which can be transported by a beamline. An example of this in the transverse direction is the restriction imposed by the beam pipe:

$$A_x = \left(\frac{d^2(z)}{\beta_x(z)} \right)_{\text{min}} \tag{2.3}$$

where $d(z)$ is the width of the beam pipe. A similar restriction can also be present in the longitudinal dimension where particles must synchronously arrive at a RF cavity to be correctly accelerated. The transport and acceleration of large emittance beams is an involved task, requiring strong magnetic fields and large apertures which tend to significantly increase the cost of a beamline.

2.1.2 Normalisation

By definition the geometrical emittance of a beam reduces as it is accelerated and the (x', y') components are reduced as a result. When considering a beam with varying energy this is often inconvenient and hence an alternative definition is used which does not change. The normalised emittance, ϵ_{N} is defined as[27]:

$$\epsilon_{\text{N}} = \gamma\beta\epsilon \tag{2.4}$$

where β and γ refer to the relativistic quantities.

2.1.3 RMS Emittance

To extend the emittance concept to an entire beam a modified definition must be used since a beam is a statistical object, combining many particles with different emittances and has no defining edge. The RMS emittance, ϵ_{RMS} is defined as the area in trace space which contains 1σ of the beam. This quantity is in essence a measure of how well collimated a beam is and can be determined using[27]:

$$\epsilon_x^{\text{RMS}} = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle x x' \rangle^2} \tag{2.5}$$

The measurement can also be thought of in terms of a covariance matrix, which is defined in two dimensions as:

$$\Sigma_x = \begin{pmatrix} \sigma_{xx} & \sigma_{xx'} \\ \sigma_{xx'} & \sigma_{x'x'} \end{pmatrix} \tag{2.6}$$

From this matrix, the emittance can be calculated using the following relation:

$$\epsilon^{\mathrm{RMS}} = \sqrt[m]{\det \Sigma} \tag{2.7}$$

where m is the number of dimensions.

2.1.4 Solenoid Dynamics

Solenoid magnets feature heavily in transport of high emittance beams and will be used in MICE to measure the emittance of each particle. As particles traverse the fringe field of the solenoid, transverse components of the field lines exert an azimuthal force on the beam and cause the beam to gain kinetic angular momentum, which confines the beam. Upon exiting the solenoid, transverse components of the field lines are reversed, and an equal amount of angular momentum is removed.

Since the beam has angular momentum, the transverse dimensions are interconnected and can no longer be treated independently. The covariance matrix for a matched Gaussian beam in a solenoid is given by[28], and has been be translated to trace space co-ordinates:

$$\Sigma = \frac{mc\epsilon_N}{P_Z} \begin{pmatrix} \beta_\perp & & & \\ -\alpha_\perp & \gamma_\perp & & \\ 0 & \beta_\perp\kappa - \mathcal{L} & \beta_\perp & \\ -\beta_\perp\kappa - \mathcal{L} & 0 & -\alpha_\perp & \gamma_\perp \end{pmatrix} \tag{2.8}$$

where κ is a focusing term and $\mathcal{L}$ is related to the canonical angular momentum of the beam. Note that now the twiss parameters α, β and γ are accompanied by a $\perp$ subscript, indicating the beam is cylindrically symmetric. Using the approach in equation 2.7 the emittance can be determined from the matrix:

$$\epsilon^{\mathrm{RMS}} = \sqrt[4]{\det \Sigma} = \frac{\epsilon_N}{\gamma\beta} \tag{2.9}$$

2.1.5 Measuring Emittance

The measurement of emittance is a crucial diagnostic for accelerators and can be determined from the covariance matrix for the beam, which requires knowledge of the phase space of the beam and cannot be determined solely from a position measurement. A number of methods exist to determine the covariance matrix, some of these include:

- Multiple profile monitors, where knowledge of the beam transport between each monitor enables the measurements to be combined and populate the covariance matrix. A variant of this was used as an early diagnostic in MICE [29].
- Variable strength quadrupole followed by a beam profile monitor, allowing the transport to the monitor to be varied to populate the covariance matrix. This technique was used in chapter 5.
- A "pepper pot" followed by profile monitor. The pepper pot restricts transmission of the beam to a well defined x-y pattern, allowing the profile monitor to determine both the position spread and divergence of the beam.

In MICE each particle will be tracked through a spectrometer solenoid, allowing the entire 6D phase space to be determined for each particle. The covariance matrix can then be found from the single particle measurements. This technique was adopted in order to make a precision (0.1% absolute error) emittance measurement, that is statistically limited.

2.2 Ionisation Cooling

The cooling of a muon beam requires a rapid technique in order to minimise decay of particles during the process, making existing techniques unsuitable for muons. To rapidly cool muons a different approach is required, namely ionisation cooling[30].

The ionisation cooling process is conducted in two stages. First the beam is passed through material (absorber), which induces energy loss by ionisation and reduces both longitudinal and transverse momentum (process 1 in figure 2.2). Secondly the beam is accelerated to restore the lost longitudinal momentum(process 3 in figure 2.2), allowing the process to be repeated.

As the transverse momentum is reduced, the transverse phase space of the beam is also reduced and the beam is cooled. Within the absorber multiple scattering will

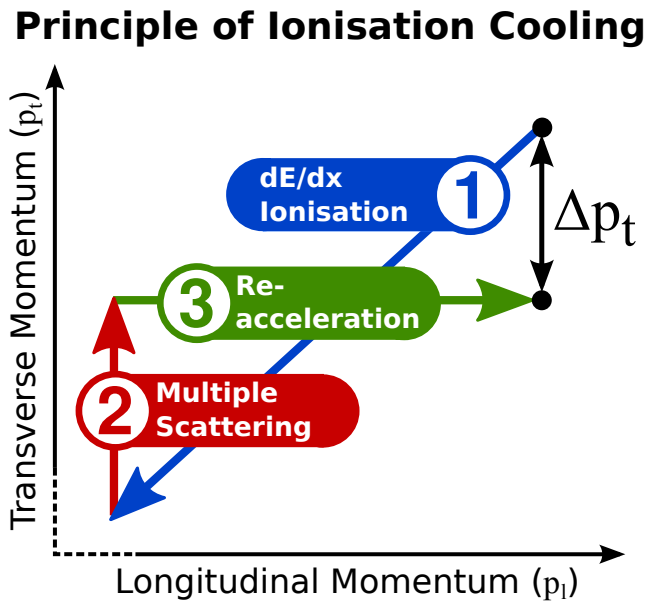

Figure 2.2: Change in transverse and longitudinal momentum during the ionisation cooling process. The ionisation and multiple scattering processes occur simultaneously inside the absorber.

also occur and serve to increase the phase space and heat the beam(process 2 in figure 2.2). The normalised transverse emittance reduction in the absorber can be estimated by the following equation[31]:

$$\frac{d\epsilon_N}{dz} \approx -\frac{\epsilon_N}{\beta^2 E_\mu}\left\langle\frac{dE}{dz}\right\rangle + \frac{\beta_\perp (0.014\text{ GeV})^2}{2\beta^3 E_\mu m_\mu X_0} \tag{2.10}$$

where β is the particle velocity, $\frac{dE}{dz}$ is the rate of energy loss, $\beta_\perp$ is the cylindrically symmetric twiss parameter beta and X_0 is the radiation length in the absorber. The first term refers to the cooling effect, while the second accounts for the heating effect. Equilibrium emittance occurs when the two effects are in balance:

$$\epsilon_{eq} \approx \frac{\beta_\perp (0.014\text{ GeV})^2}{2\beta m_\mu X_0}\left\langle\frac{dE}{dz}\right\rangle^{-1} \tag{2.11}$$

For successful cooling of the beam two key conditions must be respected, first an absorbing material must be chosen such that $X_0\langle\frac{dE}{dz}\rangle$ is maximised. Secondly, the beam must be strongly focused in order to minimise $\beta_\perp$ within the absorber. Design studies commonly feature solenoid magnets in order to simultaneously focus in both transverse dimensions. These magnets give the beam kinetic angular momentum,

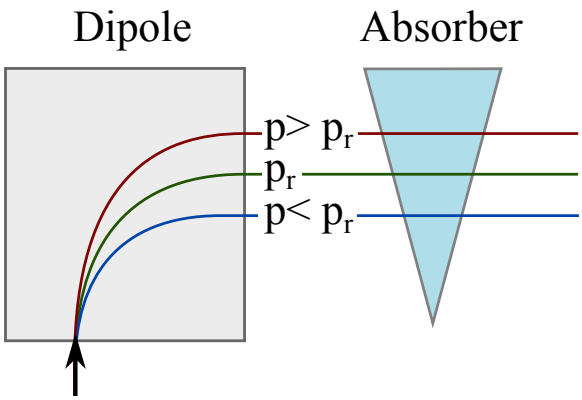

Figure 2.3: Emittance exchange using a wedge shaped absorber. Low momentum particles have a smaller bend radius and pass through less absorber, resulting in a reduced energy loss.

which is also reduced by the absorber. This loss can accumulate over a long cooling channel, reducing the solenoids' ability to confine and focus the beam. The effect can be mitigated by operating the channel in flip mode, where the field direction alternates in each absorber.

2.2.1 Longitudinal Phase Space

During the ionisation cooling process the longitudinal phase space of the beam tends to increase due to energy straggling within the absorber. Furthermore the energy of the cooling channel is below the minimum ionizing point, such that particles with decreased energy have a larger loss in the absorber. This effect can accumulate over time and cause particles to exceed the maximum longitudinal acceptance of the beamline and be lost. To prevent excessive longitudinal growth a technique for cooling longitudinal emittance can also be considered.

To cool the longitudinal phase space, an exchange process can be used to translate longitudinal emittance to a coolable transverse emittance. One possible implementation of the emittance exchange process operates by using different bend radii inside a dipole magnet to separate low from high momentum particles by position. A wedge shaped absorber is then placed such that the higher momentum particles must travel further through the medium than the low, causing an increased energy loss, as shown in figure 2.3. As the beam now has an increased spread in position the transverse phase space has been increased by the exchange.

Material	X_0 (g cm^{-2})	dE/dz (MeV g^{-1} cm^2)	ϵ_{eq} (mm)
LH_2	63.04	4.103	1.6
LiH	79.62	1.897	2.6
C	42.70	1.742	5.3
Al	24.01	1.477	10.3
Cu	12.86	1.403	22.3

Table 2.1: Potential materials for MICE Step IV cooling study. ϵ_{eq} calculated from using equation 2.11, with $P_z = 207$ MeV/c and $\beta_\perp = 42$ cm [34]. Note that the cooling equation does not consider the effect of safety windows surrounding the LH_2.

2.3 MICE Cooling Channel

The MICE cooling channel is a prototype cooling cell realised from the Neutrino Factory Feasibility Study II[32]. The cell is expected to reduce the emittance of the beam by 10%[33], using a total of three absorbers and a pair of accelerating cavities. Inside the cooling channel, focusing is achieved using solenoidal fields produced by superconducting coils. The cell can be run in either solenoid mode, where each magnet has the same polarity, or flip mode with alternating polarity. In flip mode the cell can be repeated to further reduce the emittance, until equilibrium is reached.

2.3.1 Absorber and Focus Coil (AFC) Module

The focus coil is designed to strongly focus the beam at the centre of the absorber in order to minimise the beta function of the beam. This is achieved using a pair of superconducting coils cooled using a pair of cryocoolers. The coils can be operated in one of two modes; flip mode where the coils operate with opposing polarity and solenoid mode where the coils have the same polarity.

MICE aims to test a number of different absorber materials, shown in table 2.1. The Liquid Hydrogen (LH_2) absorber is expected to give the greatest cooling effect and is retained inside a 21 litre aluminium container at the centre of the focus coil (see figure 2.4). To mitigate the risk of LH_2 leak, a pair of safety windows enable a separate vacuum to be formed around the hydrogen container. Since the additional material in the beam will cause extra heating, the windows are curved in order to minimise the thickness. The study of solid absorbers can also be conducted, by removing the aluminium container and safety windows.

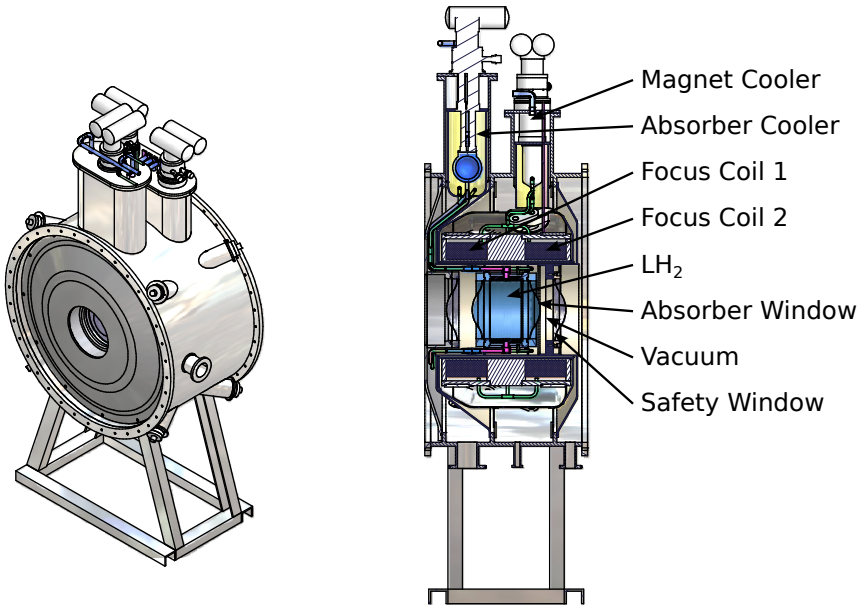

Figure 2.4: MICE absorber focus coil, configured for LH_2 operation.

2.3.2 RF Cavity and Coupling Coil (RFCC) Module

Re-acceleration of the muon beam after the absorber is performed using four 201.25 MHz normal-conducting radio frequency cavities. Transverse acceptance is maximised by using large aperture cavities and a 2.5 T magnetic field. Thin beryllium windows are then used to close the cavities. The cavities are designed to operate at up to 17 MV m^{-1} [35], and accept 4.6 MW of peak input power. In order to minimise power supply and dissipation requirements, the cavities will be operated in 1 ms pulses.

Surrounding the cavities is a large superconducting coupling coil, which provides the solenoidal field. The external magnetic field can focus dark current within the cavity, which causes damage to the internal surface and induces breakdown at a lower gradient[36]. The MICE cavities are electropolished to reduce the dark current and minimise risk of damage from this effect. The testing of the cavities for MICE is under way at the MuCool Test Area (MTA) facility at Fermi National Accelerator Laboratory.

The power for the cavities will be derived from a total of four amplifiers, each driving a pair of cavities[37]. This will restrict the maximum gradient to 8 MV m^{-1} in MICE. A low level RF control system will regulate the power of the output and

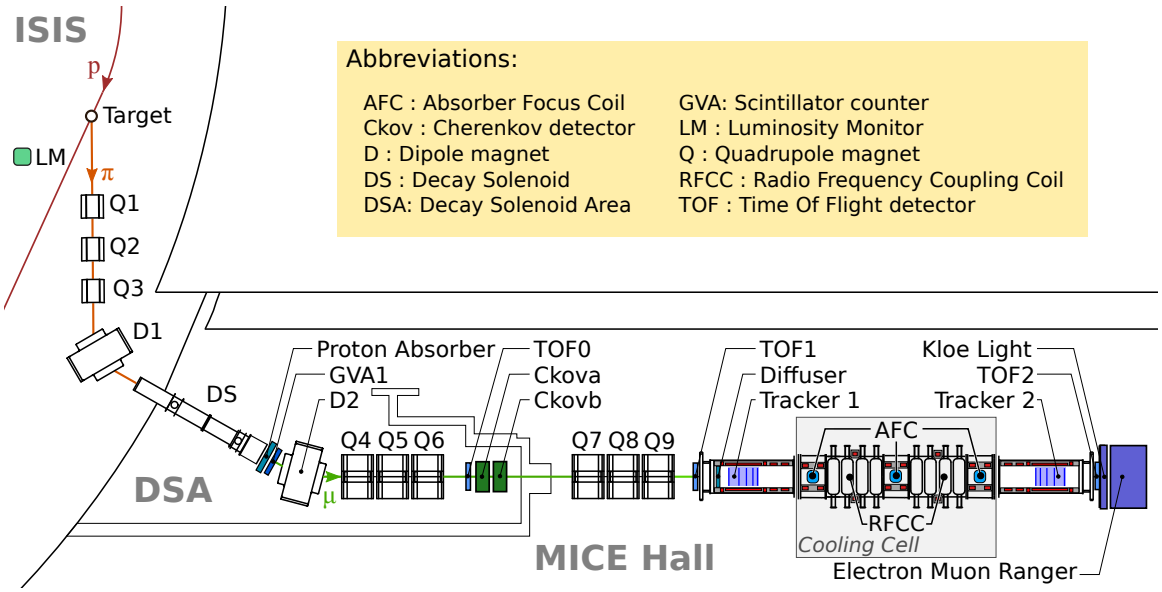

Figure 2.5: Drawing of the complete MICE beamline. Adapted from [39]

prevent reflections exceeding the peak capacity of coaxial lines transferring the RF to the cavities.

Particles for MICE are derived from an external source which is not correlated with the MICE accelerating RF and therefore the phase of the RF will be different for each particle. To estimate the energy gained in the cavities, the phase for each particle will be measured with 5° precision[38].

2.4 Muon Beamline

2.4.1 Particle Production Target

The particles used in MICE are derived from the ISIS Synchrotron, which accelerates protons to a maximum energy of 800 MeV. The acceleration cycle takes 10 ms, at the end of which particles are extracted for use in the ISIS target stations.

When required, the MICE target mechanism inserts a small titanium tip to parasitically sample the protons during the ISIS acceleration cycle. To minimise losses within the accelerator, the mechanism only intersects the beam during the final 3 ms of the ISIS spill, when the pion production cross section is the highest.

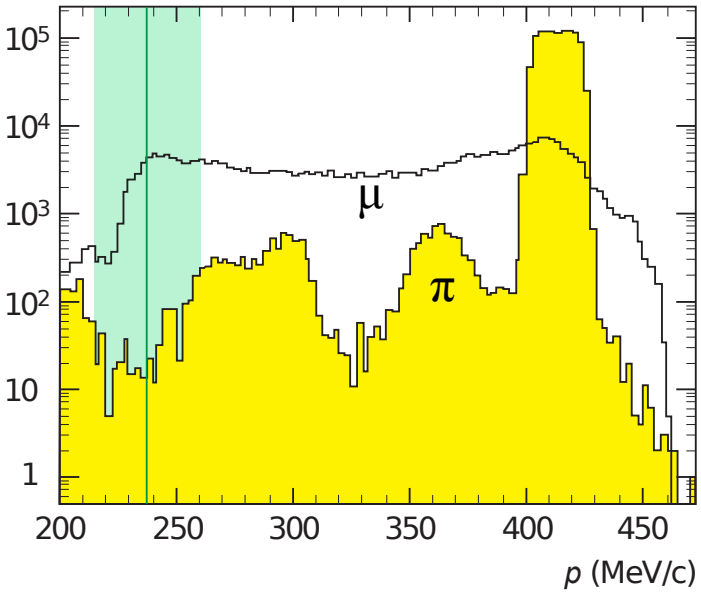

Figure 2.6: Estimated spectrum of muons and pions arriving at D2, for a 400 MeV/c selection at D1. The green band shows the accepted momentum range for D2. [40]

2.4.2 Beam Transport

A fraction of the particles from the target interaction are collected using a quadrupole triplet (Q1-3) and transported to a dipole magnet (D1). The dipole enables an initial momentum selection and directs the beam towards the MICE hall. Next the beam enters a long 5 Tesla superconducting decay solenoid (DS), which serves to increase the path length of the particles and increase the proportion of muons in the beam. After this a number of variable thickness plastic sheets can be inserted into the beam, which absorb the vast majority of protons which could saturate downstream detectors. The beam next enters a second dipole magnet(D2) to deflect particles into the MICE hall and provide a second point of momentum selection. Finally a pair of quadrupole triplets(Q4-6, Q7-9) guide the beam into the upstream spectrometer solenoid, which makes the upstream emittance measurement. Upon entering the solenoid a number of variable thickness irises can be deployed, which diffuse the beam and enable a number of different input emittances to be studied.

2.4.3 Beam Selection

By carefully selecting the momentum of the dipoles different beams can be generated. To produce a high purity muon beam one can use backward decaying muons by selecting a low momentum in D2. For calibration a much higher rate beam containing electrons, muons and pions can also be achieved by setting the momentum of D1 equal to that of D2. Figure 2.6 shows the estimated spectrum of pions and muons at D2, the green line indicating the selected momentum at D2 to maximise muon beam purity.

2.5 Detectors

A variety of particle detectors are installed along the MICE beamline, each particle being identified before and following the cooling channel. The momentum is measured before and after cooling using a pair of high precision fibre trackers. Finally a calorimetry system is used to perform particle identification after the second momentum measurement. In the final run of MICE the TOF system will be combined with a phase measurement of the RF to determine the energy restored in the cavities.

2.5.1 Cherenkov

To aid in separation of pions and muons prior to the cooling channel, a pair of high density aerogel cherenkov detectors are used [41]. The upstream detector provides discrimination for the 240 MeV/c beam setting with a refractive index of $n = 1.07$, while the other detector is suitable for 200 MeV/c beams with $n = 1.12$.

Inside each detector, the aerogel is imaged by four large photomultiplier tubes, whose signals are next digitised using a high speed ADC, the CAEN V1731. Once digitised analysis software can sum the collected charge from all four photomultipliers in order to determine the particle species.

2.5.2 Time of Flight Detectors

The Time of Flight (TOF) detectors are used in all MICE stages and have proven to be the most versatile detectors in the beam thus far, allowing both a time and position measurement.

Each TOF detector is constructed from a number of fast scintillator bars orientated in both horizontal and vertical directions as shown in figure 2.7. Each bar is optically coupled to a pair of high speed photomultipliers (rise time $<$ 1 ns), mounted

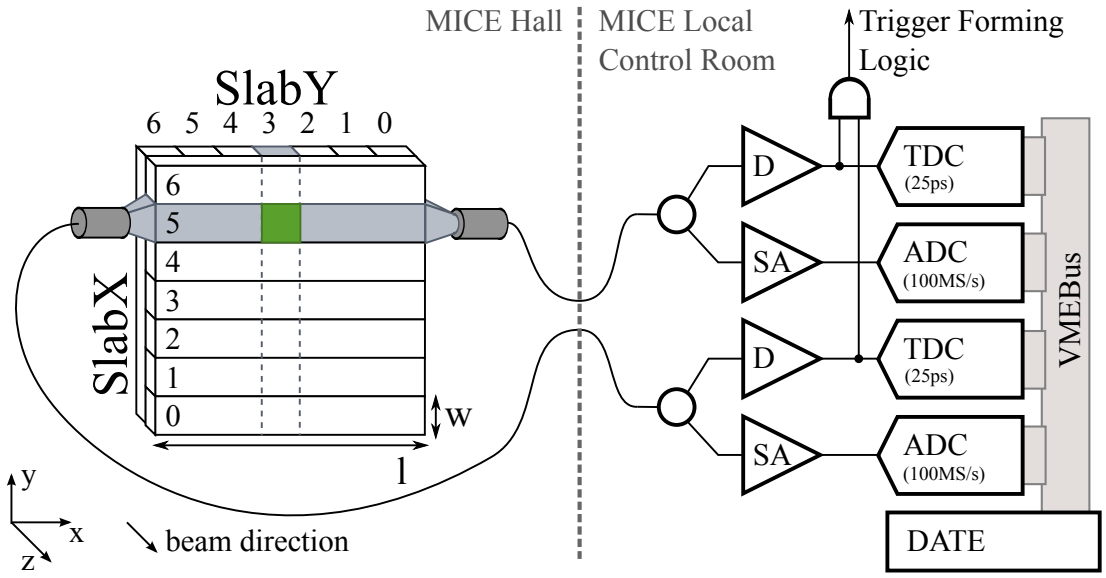

Figure 2.7: Sketch of the TOF1 detector, showing a single instrumented scintillator bar for clarity. The names of each slab as read in software are also labelled.

at each end of the bar. To allow the photomultipliers to operate in the stray magnetic field produced by the beamline, each photomultiplier is shielded externally.

Signals from each photomultiplier are then sent to the MICE Local Control Room (MLCR) and split into a discriminator(D) and a shaping amplifier(SA). The discriminator provides a digital signal for the Time to Digital Converter(TDC) and can be used in coincidence to provide a trigger source to the DAQ trigger logic. A shaping amplifier elongates the signal, such that a lower bandwidth ADC can record the pulse height, which is required for a time walk correction[42].

The position of the particle can be determined using two methods. The first relies on using the horizontal and vertical slabs to form pixels, where each pixel has a size the width (w) of the bar. This measurement can be improved using the differential time measured by opposing PMTs. In order to use this technique the effective speed of light in the scintillator and differential cable delays must be known[46].

In addition to position measurements, the time of flight can also be used in conjunction with a momentum measurement to perform particle identification. Figure 2.8 shows the time of flight for a pion beam, where the electron, muon and pion peaks can be clearly identified. The muon beam has a much higher dispersion, leading to a larger variation in time of flight.

Parameter	TOF0	TOF1	TOF2
Slab Length, l (mm)	400	420	480
Slab Width, w (mm)	40	60	60
Number of Slabs	10	7	8
Scintillator	BC-420	BC-404	BC-404
TOF Resolution (ps)	52.2 ± 0.9	59.5 ± 0.7	52.7 ± 1.1
Coarse Position Resolution (mm)	11.5	17.3	17.3
Fine Position Resolution (mm)	9.8	11.4	-

Table 2.2: Characteristics of the MICE TOF system [43, 44, 45]. The coarse position resolution is from the pixel based measurement, while the fine resolution is from the differential time measurement.

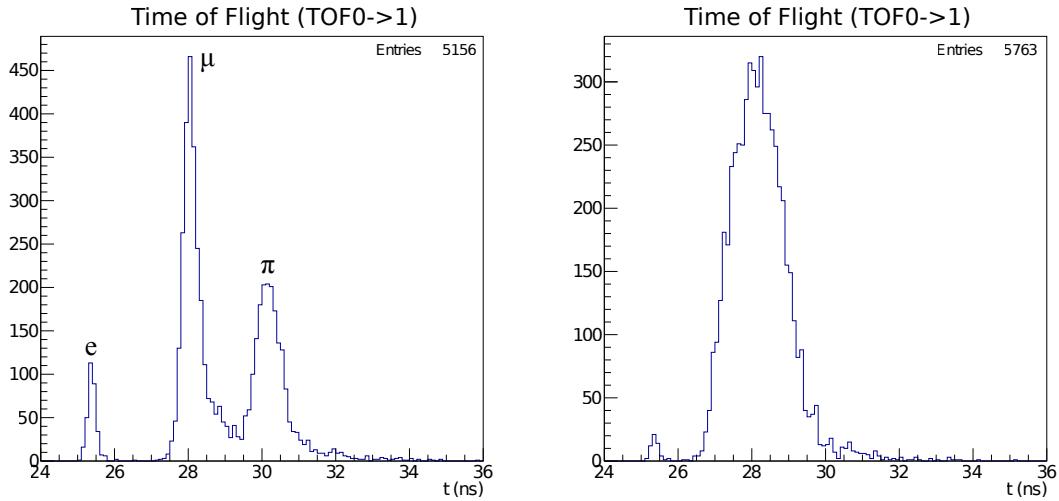

Figure 2.8: Left: Time of flight spectrum for a 272 MeV/c pion beam, showing electron, muon and pion peaks. Right: Time of flight spectrum for a 237 MeV/c muon beam, showing a broad muon peak.

Parameter	RMS Resolution
x (mm)	0.54
y (mm)	0.44
t (ns)	<1
P_x (MeV/c)	2.1
P_y (MeV/c)	1.5
P_z (MeV/c)	4.5

Table 2.3: Estimated performance of the MICE Scintillating fibre trackers [47].

2.5.3 Tracking Spectrometer

The primary detector for measuring beam position and momentum is a scintillating fibre tracker, mounted inside a superconducting solenoid magnet. This causes traversing particles to travel in a helix, enabling a measurement of momentum in addition to position.

Each tracker contains a number of measurement stations, positioned along the tracker. The station contains three planes of parallel fibres orientated at 120°, hits from each plane can be combined to produce a spacepoint. Using pattern recognition software the spacepoints are combined into a track, which is processed by a Kalman filter to determine the precise position and momentum of the particle. The expected resolution of the tracker is outlined in table 2.3.

The momentum information can be combined with the time of flight from the TOF0 and TOF1 detectors to provide clean upstream particle identification. To determine the emittance change a pair of almost identical trackers are placed immediately before and after the cooling channel.

2.5.4 KLOE Light

The KLOE-Light (KL) detector is a smaller variation of the KLOE[48] electromagnetic calorimeter and forms a component of the downstream particle identification. The active area is constructed from extruded lead foils, inlaid with bicron BCF-12 scintillating fibre as shown in figure 2.9, with a ratio of ~1:1. Scintillation light is then guided from each slab into a total of six photomultiplier tubes (three each side of the slab). Mu-metal magnetic shields are fitted to each photomultiplier to mitigate against large stray magnetic fields from the beamline. Output from each photomultiplier is shaped and sampled by a 14-bit flash ADC.

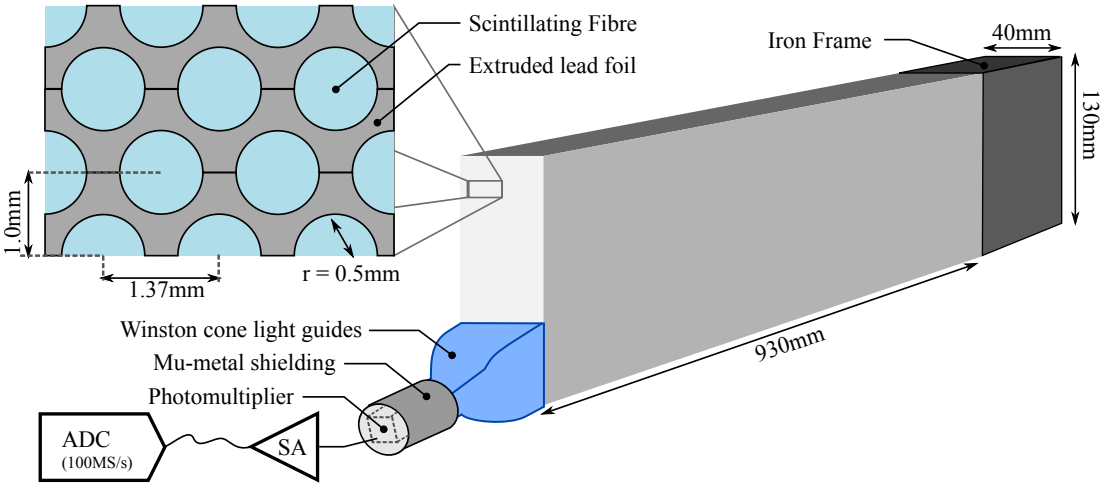

Figure 2.9: Drawing of a single slab from the KL detector, a total of seven are used to instrument the muon beam.

The detector occupies a volume of 93 cm × 93 cm × 4 cm, and initiates a shower for electrons. Alone the detector has been used to estimate the level of pion contamination within the MICE muon beams to be around 1% [49].

2.5.5 Electron Muon Ranger

The Electron Muon Ranger[50] is a totally active tracking calorimeter and is the final detector in the MICE beamline. The active area is constructed from triangular extruded scintillator bars 1.1 m in length; 59 bars are next combined to form a plane 99 cm wide and 1.65 cm deep. A total of 48 planes constitute the detector; the planes have alternating horizontal and vertical orientation to allow x and y tracking of particles.

Light from each scintillator bar is extracted using a wavelength shifting fibre. One end of the fibre is connected to a single channel of a multi anode photomultiplier, which is read-out using a MAROC ASIC[51]. The MAROC performs discrimination of the signals and records time of hit and time over threshold for each channel[52], which can be used to determine the path of the particle. All detector hits during the ISIS spill are buffered within the detector in combination with external trigger times. This configuration provides zero dead time and enables the reconstruction of an entire event, including decay electrons.

In addition, the energy deposited in each layer is measured by connecting the remaining 59 ends of the wavelength shifting fibre to a single anode photomultiplier.

The signal is digitised using a 500 MS/s flash ADC, upon the receipt of a trigger from the MICE DAQ. The analysis software will use information from the tracking, hit timing and energy deposited per layer in order to separate muons from decay electrons.

2.5.6 Luminosity Monitor

The luminosity monitor[53] is the only detector installed outside of the MICE beamline and serves to provide a relative measurement of the number of particles produced by the target, independent from the beamline settings. This enables normalisation for detectors and an aid in verification of simulation codes. The detector is installed on the inner side of the ISIS ring, at an identical angle to the MICE beamline.

The active detector is built from four pieces of scintillator, each imaged by a single photomultiplier. The scintillators are placed in pairs and separated by a 15 cm deep polyethylene block. The polyethylene stops protons below 500 MeV and pions below 150 MeV. To determine the luminosity the signals from each photomultiplier are discriminated and coincident pulses between each photomultiplier pair are counted by a simple scaler. The number of counts recorded at the end of the spill is proportional to the MICE beam intensity.

2.6 DAQ, Computing and Control

2.6.1 Control and Monitoring

Successful operation of the experiment requires the control of a large number of systems, ranging from magnet control to detector settings. The control and monitoring of these systems is handled by the Experimental Physics and Industrial Control System (EPICS)[54], which enables the sharing of Process Variables (PVs) over a network. The beamline components are connected to a number of Input Output Controllers (IOCs) which translate PVs to settings on the equipment and sensor values back to PVs.

During running a large number of PVs are alarmed, which enables rapid detection of tripped off equipment degrading data runs. In addition the PVs are also logged to disk, which allows long term monitoring of systems.

A run control system is also under development to ensure that the beamline and detector parameters are loaded correctly. The system uses a database of settings

as a template to verify the current settings against and inhibit DAQ operation until satisfied.

2.6.2 Configuration Database

A large number of different configurations of MICE will be run to configure the beamline, detectors and study ionisation cooling. To allow meaningful simulation and reconstruction it is essential to have a clear record of which calibration should be applied to each detector, the field strengths of the magnet and geometry of the beamline. In MICE this information is stored in the Configuration Database (CDB). During running settings are committed to the database by the run control system and calibrations are uploaded by detector groups. A public read-only interface to the database is provided to allow run settings to be retrieved.

2.6.3 Data Acquisition

The detector DAQ architecture of MICE is to use a number of VME crates, each connected to a computer for readout. At the beginning of an ISIS spill, a start-of-spill signal is distributed to each VME crate and signals particle data is imminent. Next an electronic gate opens, that enables particle triggers to be created from hits in either GVA1 (a scintillation counter upstream of D2), TOF0 or TOF1. Each particle trigger instructs the front end electronics to digitise and store the appropriate data. Once the ISIS spill has concluded, a physics event trigger initiates the transfer of the buffered data into computer memory. Finally an end-of-spill trigger signals the conclusion of spill data collection.

The software employed by MICE to combine the readout from various computers is DATE[55], developed by the ALICE collaboration. Each computer connected to a VME crate is a Local Data Concentrator (LDC) and sends the collected data over a local network to a Global Data Concentrator where it is written to disk. At the end of each shift the data is uploaded to the GRID and CASTOR tape for permanent storage.

In addition to detector data, several of the more complex subsystems such as the target and superconducting magnets use independent DAQ systems to allow monitoring and diagnostics.

2.7 Analysis and Simulation

2.7.1 MICE User Analysis Software

Simulation, Reconstruction and Analysis of the MICE cooling channel and detectors is accomplished using the MICE User Analysis Software (MAUS)[56]. The software accepts data collected by the DAQ and loads appropriate calibration and geometry information from the MICE CDB.

Monte Carlo simulation of the cooling channel is also performed in MAUS, which in turn uses the Geant4 tracking code and physics processes. MAUS loads the necessary geometry and field maps into Geant4. In addition it creates a realistic beam starting from the second dipole magnet (D2). The output from Geant4 is then translated into realistic looking particle hits by using digitisation routines to imitate each detector.

2.7.2 G4Beamline

The entire beamline has also been modelled in G4Beamline[57], which again uses Geant4 tracking codes. G4beamline allows the rapid construction of beamline models by placing magnets and materials in the beam path. The simulation input files are referred to as decks. G4Beamline is used to generate the input beam for MAUS.

2.8 Staging and Timescale

MICE is a staged experiment, allowing the integration of necessary hardware to be completed over several years, as the hardware is completed. This enables preliminary measurements and testing of equipment prior to the final measurement. The stages can be summarised as follows:

Step I (complete 2010): Install and commission the target mechanism and upstream beamline (up to Q9). Install the Cherenkov, Time of Flight and KLOE-Light detectors along with necessary subsystems.

Step I.1 (complete 2013): Install and commission the Electron Muon Ranger.

Step II-III (merged into Step IV): Install each spectrometer and tracker detector. Substantial equipment delays caused the schedule to converge at Step IV.

Step IV (est. Q2 2015): Install and commission both tracker modules and a single absorber. This stage is divided into a number of sub-steps to be carried out:

.0 Align the detectors and test global tracking with no field and an empty absorber focus coil.

.1 Study detector resolutions and beam properties, using magnetic fields and an empty absorber focus coil.

.2 Study emittance change and transmission of particles in a liquid H_2 absorber, with magnetic fields.

.3 Repeat previous study for the solid absorbers, such as lithium hydride.

.4 Insert a solid wedge absorber and study 6D emittance exchange.

Step V [1] (possible) Install a second AFC module and the first RFCC module.

Step VI [1] Study the complete MICE cooling channel, containing all three AFC modules and pair of RFCC modules.

2.9 Conclusion

This chapter has outlined the concept of beam emittance, facilitating a description of the ionisation cooling process. This has enabled a description of the cooling cell design used in the second feasibility study of a neutrino factory, which forms the foundation upon which the MICE experiment is built. The MICE cooling cell has been described, followed briefly by a summary of the beamline components, detectors and software required for the experiment to operate.

It is the aim of this thesis to investigate three aspects of the experiment, required for a successful demonstration of ionisation cooling at the Rutherford Appleton Laboratory:

- Chapter 3 develops indirect diagnostic techniques for the MICE target mechanism, which was originally developed at the University of Sheffield. The diagnostic techniques fulfil a key role in the quality assurance procedure and ongoing

[1] These steps of the experiment have been abolished since the writing of this section. A replacement for step V is known as the demonstration of ionisation cooling, which uses two absorbers mounted each side of a pair of RF cavities. Focusing is provided by the AFC modules and the RFCC module is removed.

monitoring of the mechanism, by monitoring the mechanical performance of the mechanism. The ultimate goal is to identify early signs of a poor or failing target before incurring significant downtime for MICE.

- Chapter 4 explores the additional hardware needed to successfully integrate the MICE Scintillating Fibre trackers in Step IV of the experiment. This includes how to synchronise the charge integration gates on the front end electronics to the arrival of single particles in the beamline. Following this an LED calibration system is designed, which is capable of operating in the high magnetic field of the spectrometer solenoid magnets.
- Chapter 5 investigates alignment of the Step I MICE beamline, using a Quadrupole magnet scan and the TOF detectors.

This page is intentionally left blank.

Chapter 3

Diagnostics of the MICE Target Mechanism

The MICE target mechanism is responsible for the generation of pions which are collected by the MICE beamline. The target itself is designed to operate parasitically on the ISIS Synchrotron at Rutherford Appleton Laboratory. Parasitic operation dictates that the target must intersect the beam towards the end of a 10 ms acceleration cycle, but be outside the beams radius at injection 10 ms later.

These requirements place great demand on the target mechanism, requiring an acceleration in excess of 80 g. In addition the high vacuum environment prevents the use of lubrication and the radiation field restricts materials choices. Finally the limited access of the synchrotron vault requires that the target mechanism operates reliably for several months without direct intervention or inspection.

A key component of a reliable target mechanism is the use of indirect measurements to verify and monitor the mechanism prior to and during operation in ISIS. The ultimate goal is to observe early signs of failure and replace the target before incurring downtime on the MICE experiment. This chapter is devoted to the understanding of how the mechanism operates and what observables can be measured in order to indirectly monitor the mechanical performance of the mechanism.

3.1 Target Design

The MICE target mechanism contains a number of systems each contributing to its operation; the components fitted to the mechanism are shown in figure 3.1. This section provides an overview of the key subsystems and construction of the target

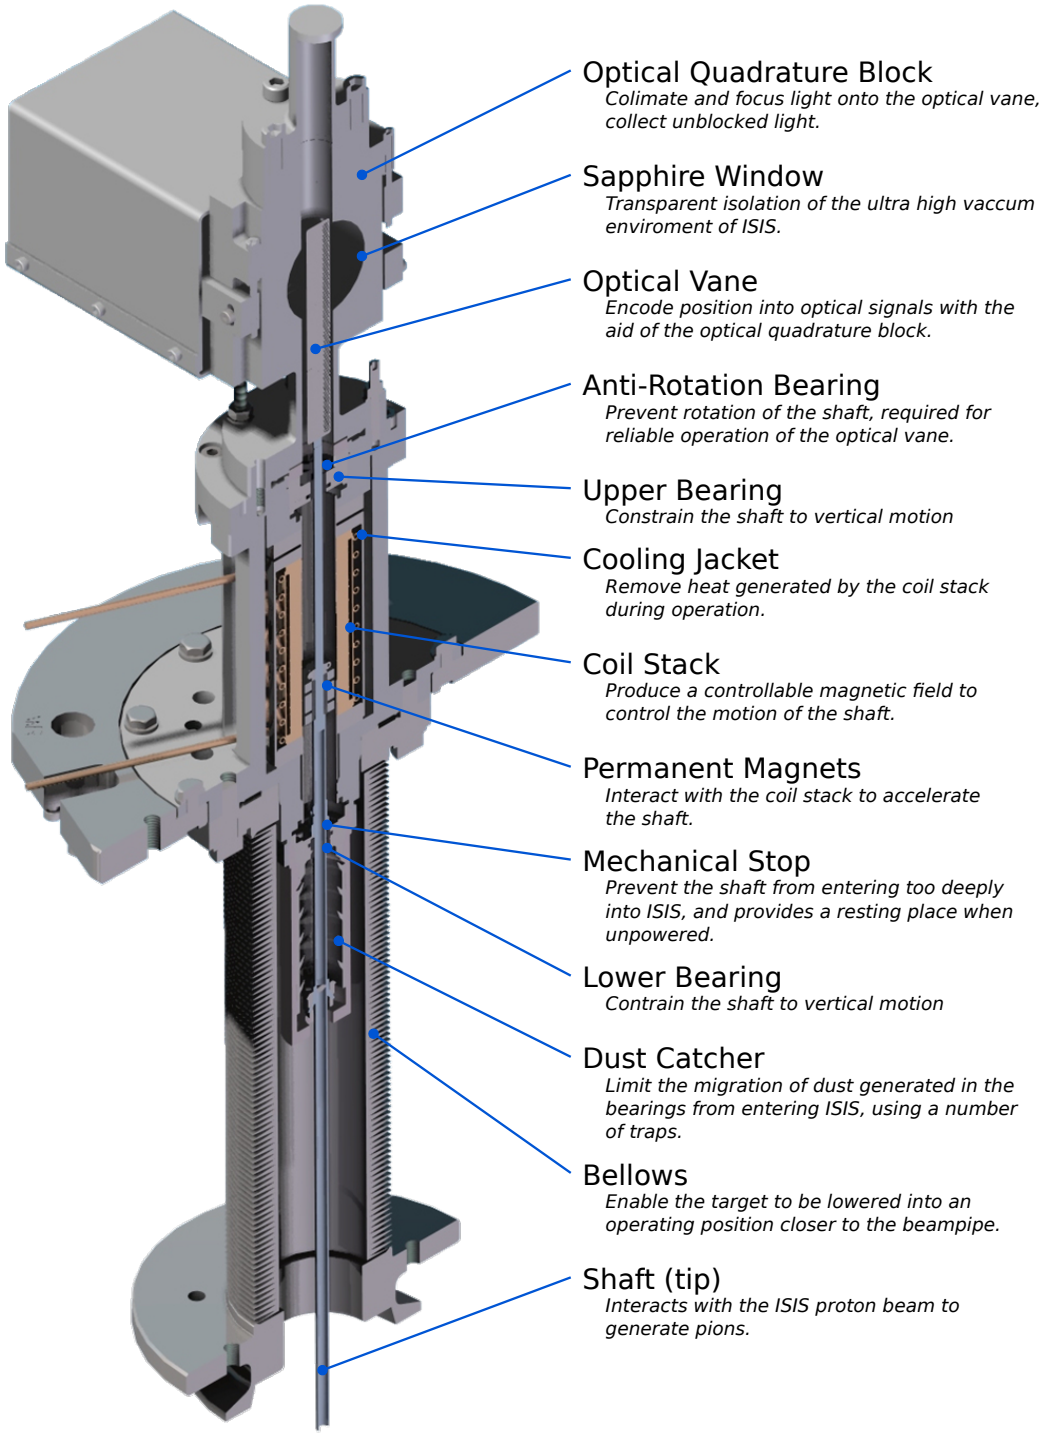

Figure 3.1: Labelled drawing of the MICE target mechanism.

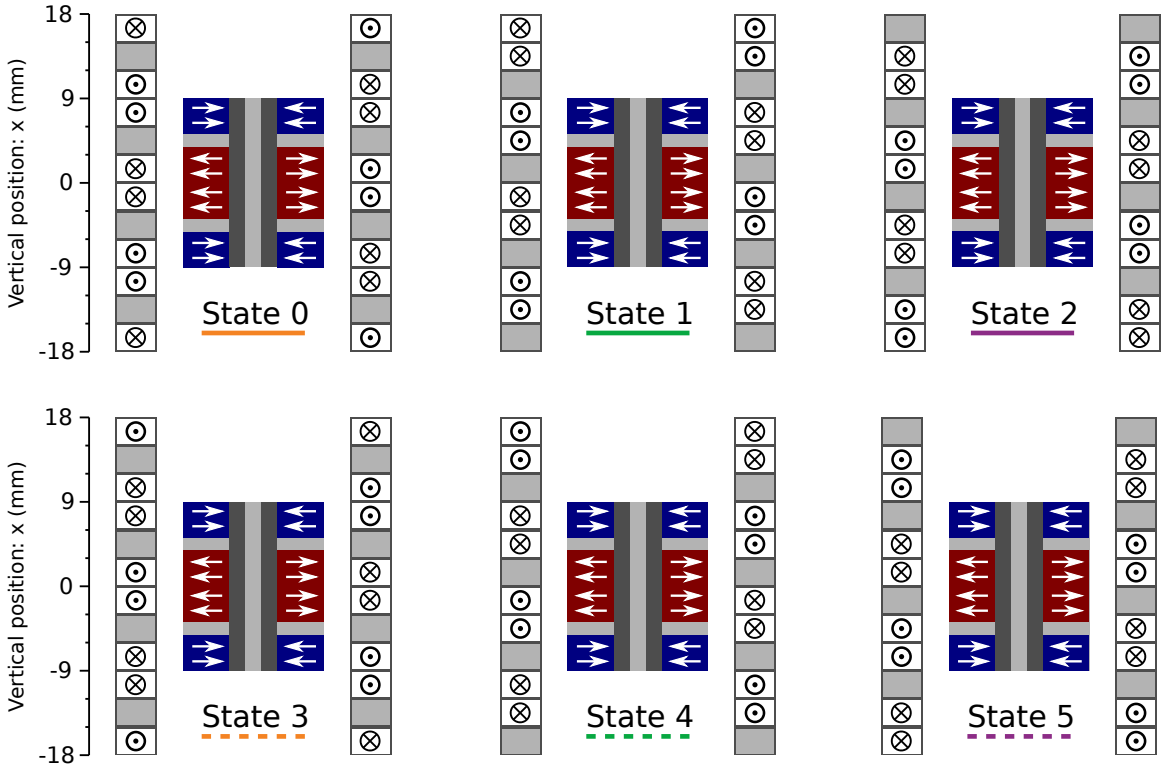

Figure 3.2: Side view of the magnet and coil assembly, showing the powered coils in the six states necessary for two phase operation. The direction of current flow is indicated by cross and dots on the diagram, which indicate current flow into and out of the page respectively, while the unpowered coil is shown as a grey block.

mechanism. A more complete description can be found in [58].

3.1.1 Motor Design

The target is driven by the interaction between a permanent magnet assembly mounted directly on the central shaft of the mechanism and a number of coils fixed to the target body, shown in figure 3.2. The permanent magnets are radially polarised, with the middle section having outbound field lines, while the top and bottom both have inbound field lines.

Surrounding the magnets is a coil assembly containing a total of 24 coils, each of which is connected to one of three phases. Within each phase the coils are connected serially and one end of the phase terminates at a common start point, while the remaining end is driven. This drive method requires at least two phases to be active at one time and provides six different drive states, also shown in figure 3.2. The direction

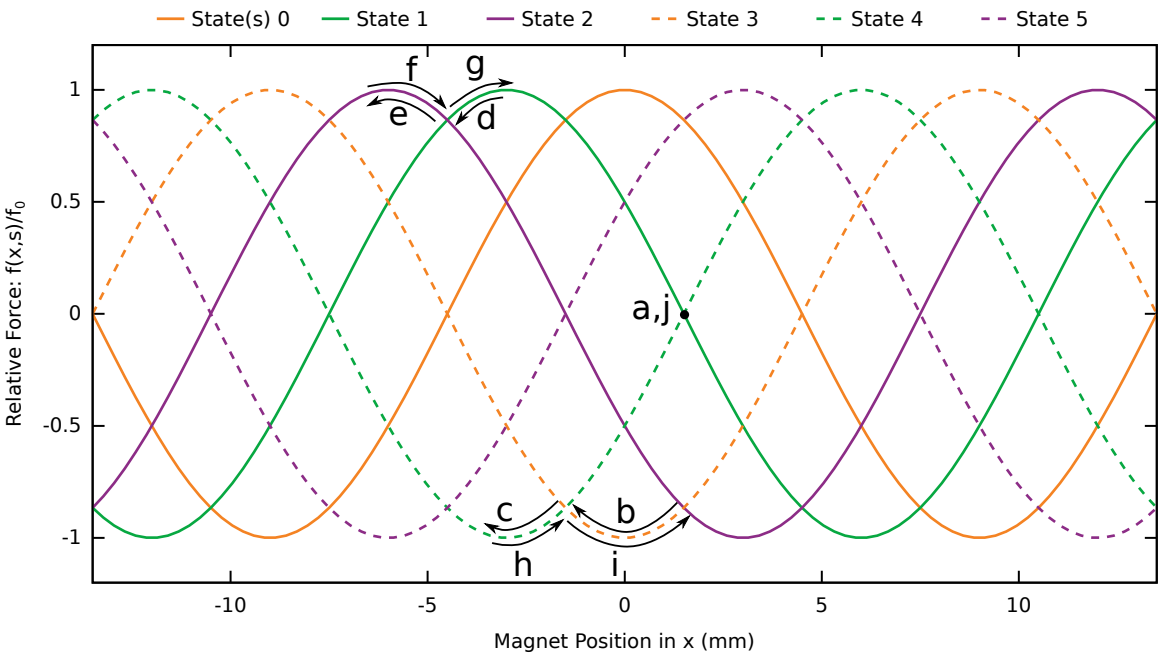

Figure 3.3: The force delivered as a function of magnet position for each state. The states required for a short actuation have been identified (a-j).

of current flow in each drive state is indicated by cross and dots on the diagram, which indicate current flow into and out of the page respectively, while the unpowered coil is shown as a grey block. The pattern of active coils repeats every 18 mm, which allows an arbitrary vertical point to illustrate the entire operation. When the magnets are positioned at 0 mm, drive state zero would produce an upwards force.

Studies of the magnetic interaction of the magnet and coil assemblies from a finite element model[59] indicate that the resultant force on the magnet assembly can be represented by a trigonometric function:

$$f(x,s) = f_0 \cos\left(2\pi \frac{x[\text{mm}] + 3s}{18}\right) \tag{3.1}$$

where x is the magnet position, s is a state number in the range 0 to 5 and f_0 is the absolute maximum force. This function is plotted in figure 3.3, for each of the six states.

The accelerating force can then be controlled by selecting the appropriate state for the current goal. An example of an actuation, where the target starts from a levitating state, moves downwards, then upwards, then stops can be described as follows (illustrated on figure 3.3):

(a) Starting at $x = 0$, the target would be levitated by using state 1 or 2.

(b) To begin downwards movement, the magnets would be accelerated downwards, by selecting state 3.

(c) After 1.5 mm of downwards travel, the state would be switched to state 4, in order to maintain maximum force.

(d) To slow the target, an upwards force would be delivered by selecting state 1.

(e) To maintain the upwards force while the target is slowed the powered state would switch to state 2, at -4.5 mm.

(f) The magnets would then stop, and begin to move upwards.

(g) To maintain maximum upwards acceleration, state 1 be powered when the -4.5 mm point is crossed.

(h) To slow the target to a stop, a downwards force would then be delivered by state 4.

(i) The downwards force would be maintained by switching to state 3 after crossing -1.5 mm.

(j) Upon detecting that the target velocity is zero, the system can return to a levitation state, again switching to either state 1 or 2.

Permanent Magnets

In order to maximise the acceleration, the shaft mass must be kept to a minimum while maximising the field generated by the permanent magnets. The magnet assembly is formed from a number of radially magnetised neodymium-iron-boron (NdFeB) segments, which are glued to a central mild steel core and weigh a total of 25 g. A fully assembled shaft has a mass of ≈ 51 g.

Coil Stack

There have been two key design iterations relating to the coil assembly. The initial design, which includes stators T1 and T2, had each coil individually wound with thirty six turns on a 18.1 mm diameter former. Each coil was 2.85 mm thick and was set in epoxy before being individually stacked on the stator. Once stacked the

Characteristic	T1,T2	S1,S3
Mechanical		
Bore Diameter (mm)	18.1	16.2
Outer Diameter (mm)	30	25
Coil Thickness (mm)	2.85	2.85
Turns per Coil	36	50
Coil Pitch (mm)	3	3
Electrical		
2 Phase Stator Inductance (μH)	502	851
2 Phase Stator Resistance (Ω)	3.42	5.43
Wire Area (mm^2)	0.24	0.19
Single Coil Inductance (μH)	38	56

Table 3.1: Comparison of key mechanical design features from T1 & T2 to S1 & S3

coils were surrounded in a water cooling jacket, which extracts heat generated by operation of the coils.

The process of stacking the coils on the stator introduced the possibility of an offset on each of the coils from the shaft's axis, which in turn leads to off-axis forces on the shaft, degrading performance. To correct this the stator body was redesigned to enable the coils to be wound directly onto the stator. In addition this revision also reduced the inner coil radius and allowed an increased number of turns in each coil in order to further improve performance. A single stator (S1) has been built to this specification, with a second (S3) nearing completion. The key design parameters are highlighted in table 3.1.

During levitation only a fraction of the available force is required to combat gravity, so the target is driven in a reduced current mode, which minimises the energy deposited in the coils. This is achieved using a technique called pulse width modulation, by which the motor switches between an un-powered state and a fully powered state at a rate of 2 kHz. The duty cycle of the powered state is limited to $\approx 3.5\%$, which substantially reduces the operational temperature.

3.1.2 Position Monitoring

In order to apply the correct acceleration to the target it is necessary to track the position of the shaft during actuation. For this task an optical quadrature system has been employed. At the top of the shaft is an optical vane containing a total

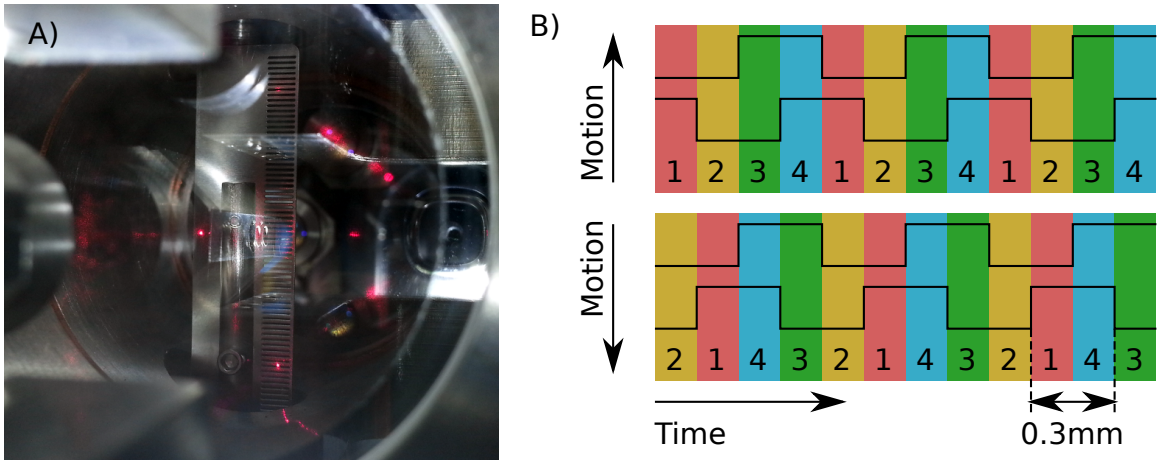

Figure 3.4: A) Image of an aligned quadrature system. B) example of decoding quadrature signals to determine distance and direction information. Each state of the two input signals is denoted a state, the ordering of which reverses when the direction of travel is reversed.

of 157, 300 μm wide apertures. A pair of light sources are tightly focused on the optical vane and set with a 90 ° phase offset, shown in figure 3.4. As the shaft moves the transmission of light through the vane pulses, with each rising edge separated by 600 μm. The offset of the second channel enables the direction of travel to also be determined from the ordering of the rising/falling edges. The second channel also increases the resolution, resulting in a quadrature step size of 150 μm. For reliable readout the tolerance of the centre of each slit in the vane was kept to within 7.5 μm, while the absolute tracking error along the vane was also kept below 7.5 μm. The vane is attached directly on the shaft and adds less than 1 g in total weight.

To identify the exact position of the target a unique reference point is provided using a third signal which is only transmitted at a single point in the vane, and is referred to as the index mark. The index marker is crossed twice per actuation, which enables miscounts to be identified and corrected. Due to the four distinct quadrature states a missed pulse results in a miscount in multiples of four. Miscounts of one can be common during operation and indicate that the index mark is close to a quadrature transition. These miscounts are safely ignored by the control system to deliver more consistent performance. In practice the quadrature system is incredibly robust and can operate for millions of actuations without observing a substantial (> 1) miscount.

Readout of the optical vane is accomplished using an optical block, which collimates and focuses light from a single mode optical fibre onto a small spot on the vane.

Light transmitted through the vane is then collected and coupled into a multi-mode optical fibre, which can be read out remotely. The entire optical assembly is mounted outside the vacuum environment and light is transmitted to the vane using sapphire windows, rated for both the ISIS radiation field and ultra high vacuum.

In order to avoid placing sensitive electronics components within the ISIS vault where they will be irradiated, all lasers and photo-diodes are located in the MICE Local Control Room. Light into the vault is sent via three single mode optical fibres, while the returned light uses three multi mode fibres. The optical signals are immune against electrical interference and ground loops over the considerable cable runs. Three 2.5 mW 635 nm commercial diode lasers, which nominally operate at 1.0 mW, are used to supply the single-mode fibre. The quadrature signals on the multi-mode fibres are received using broad spectrum photo-diodes, and are passed to transimpedence amplifiers, programmable gain amplifiers and comparators to generate digital signals. A quadrature decoder circuit incorporated within the control electronics then processes the three digital signals and determines the absolute position.

The absolute position of the target is the distance measured from the centre of the ISIS beam pipe to the target tip and is referred to as the Beam Centre Distance (BCD). During levitation outside of the beam the target nominally sits around 64 mm BCD, while the maximum insertion possible is around 20 mm BCD.

3.1.3 Bearings

Motion of the shaft is constrained to the vertical dimension using a pair of bearings mounted each side of the coil stack. The shaft has two flat edges above the magnets, which the upper bearing meets in order to also prevent rotation of the shaft.

The bearing assembly is mounted within the ISIS vacuum, which prohibits the use of lubrication and some materials due to out-gassing. Furthermore the vault also imposes activation and radiation damage constraints. Finally substantial dust release, caused by bearing wear, is also prohibited.

The two bearings are the only mechanical interface between the moving shaft and stationary body of the target. Over time the bearings wear and are the component with the shortest life, and therefore limit the time a target can be installed in ISIS. A substantial amount of development has gone into the choice of materials for the bearings. Initially a hard Diamond Like Carbon (DLC) layer was used to coat both the bearings and shaft. After a number of tests this technique was deemed unsuitable

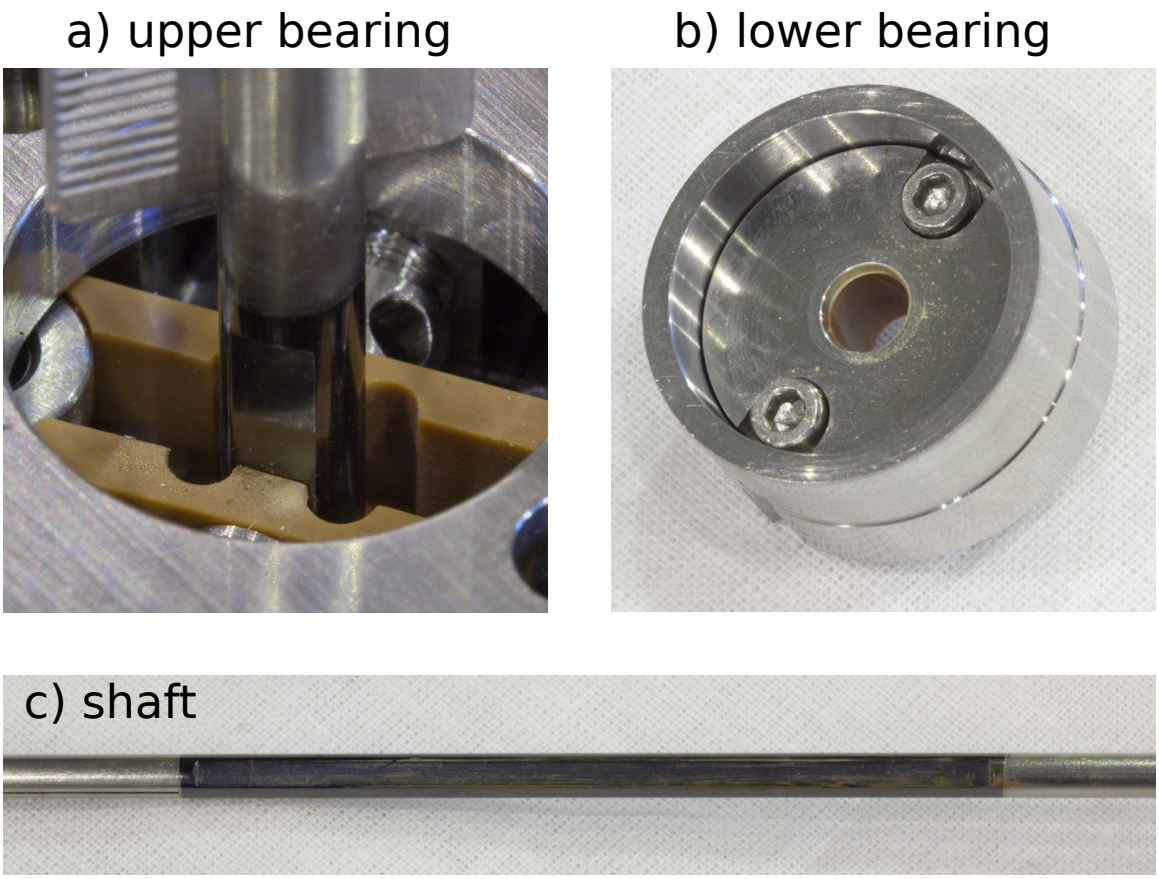

Figure 3.5: Upper bearing, lower bearing and shaft imaged after three million actuations. Small amounts of VESPEL dust can be observed as a orange powder.

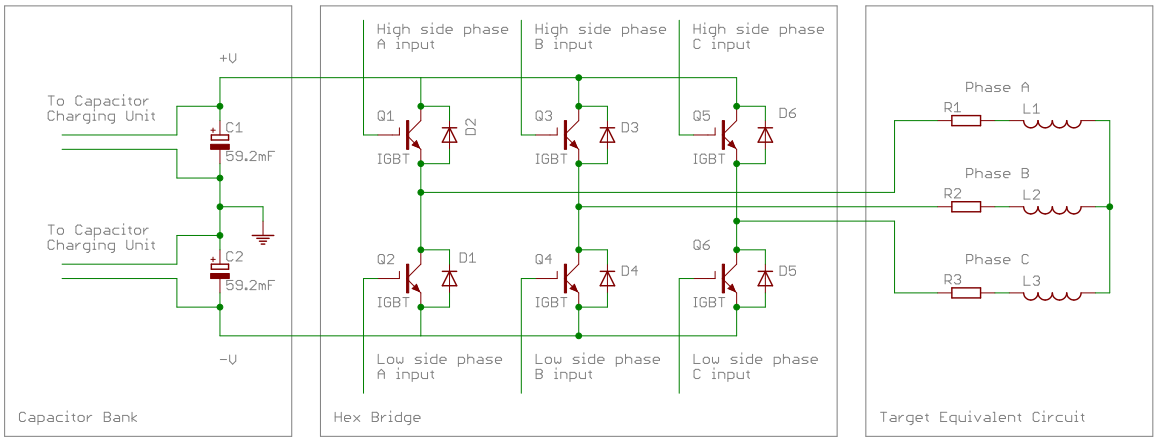

Figure 3.6: Schematic showing the core hex bridge circuit, capacitor bank and target stator.

due to variability between bearing coatings, which ultimately made the mechanism unreliable.

The coated bearings were then replaced with a hard polymer named VESPEL[60]. Unlike the coating it is machined to fit the shaft and not susceptible to the same variations that made the target unreliable. During operation the VESPEL surface wears slightly, generating a very small amount of dust particles. The amount of dust produced by the mechanism after three million actuations is shown in figure 3.5, and was deemed small enough for operation by ISIS. As a further level of protection a dust trap is mounted below the lower bearing to catch any material attempting to migrate into the synchrotron.

3.1.4 Power Electronics

During actuation the target mechanism requires up to 6 kW of instantaneous power, at up to ± 150 V. This energy is delivered from a pair of 59.2 mF capacitor banks, each of which is charged from a Capacitor Charging Unit (CCU). The limited current draw when levitating the target reduces the average power consumption to < 60 W and the CCUs can deliver up to 600 W.

In the event of a fault in either the control electronics or the chiller, the temperature of the mechanism can rapidly increase. To protect the stator a number of interlocked thermocouples initiate shut down of the charging units and discharge the capacitor banks in the event the temperature exceeds 90 °C.

Power from the capacitor banks is directed to the target stator using a number

of insulated gate bipolar transistors which are configured in a hex-bridge shown in figure 3.6. A hex-bridge is commonly used for motor drive circuits and enables each terminal to be connected to either the positive voltage, negative voltage or left in a high impedance state.

To minimise cable losses the power electronics are located as close as possible to the synchrotron, approximately 25 m from the stator. The hex bridge circuit is controlled from a remote control system, located around 70 m from the power electronics in the MLCR. The interconnections between the major subsystems are displayed in figure 3.7. A separate control signal is sent to engage each transistor in the hex bridge, requiring a total of 6 links in total. Each link is a single optical fibre, chosen for negligible susceptibility to noise and a high reliability.

3.1.5 Target Frame

The target body is mounted on a moving frame which is used to lower the target into the synchrotron during operation. After data taking the frame is then raised to remove the target from ISIS and prevent possible issues while left unattended. A gate valve can also be engaged, which completely separates the target system from the ISIS vacuum and is normally in an open state.

A series of interlocks only allow the frame to be lowered once the target is actively levitated. In the event of a fault the frame automatically raises and cannot be lowered until the fault is rectified.

3.1.6 Control Electronics and Operation

Decoding of the optical quadrature signals and driving of the hex bridge circuit is conducted by a custom built control system; the target controller. At the heart of the target controller is a Field Programmable Gate Array (FPGA) which is mounted on a printed circuit board named the USBDaq[61]. The board contains a USB connection for communication with a computer and a large number of Input/Ouput connections. A pair of daughter cards were specifically constructed to convert the inputs from the quadrature system and various status and trigger signals to digital signals which are handled by the FPGA. Outputs from the FPGA are also buffered and converted into various forms, including drive signals for the hex bridge and interlocks with the ISIS Beam Protection System.

The FPGA circuit is configured into a state machine which allows the target to be moved between four key states:

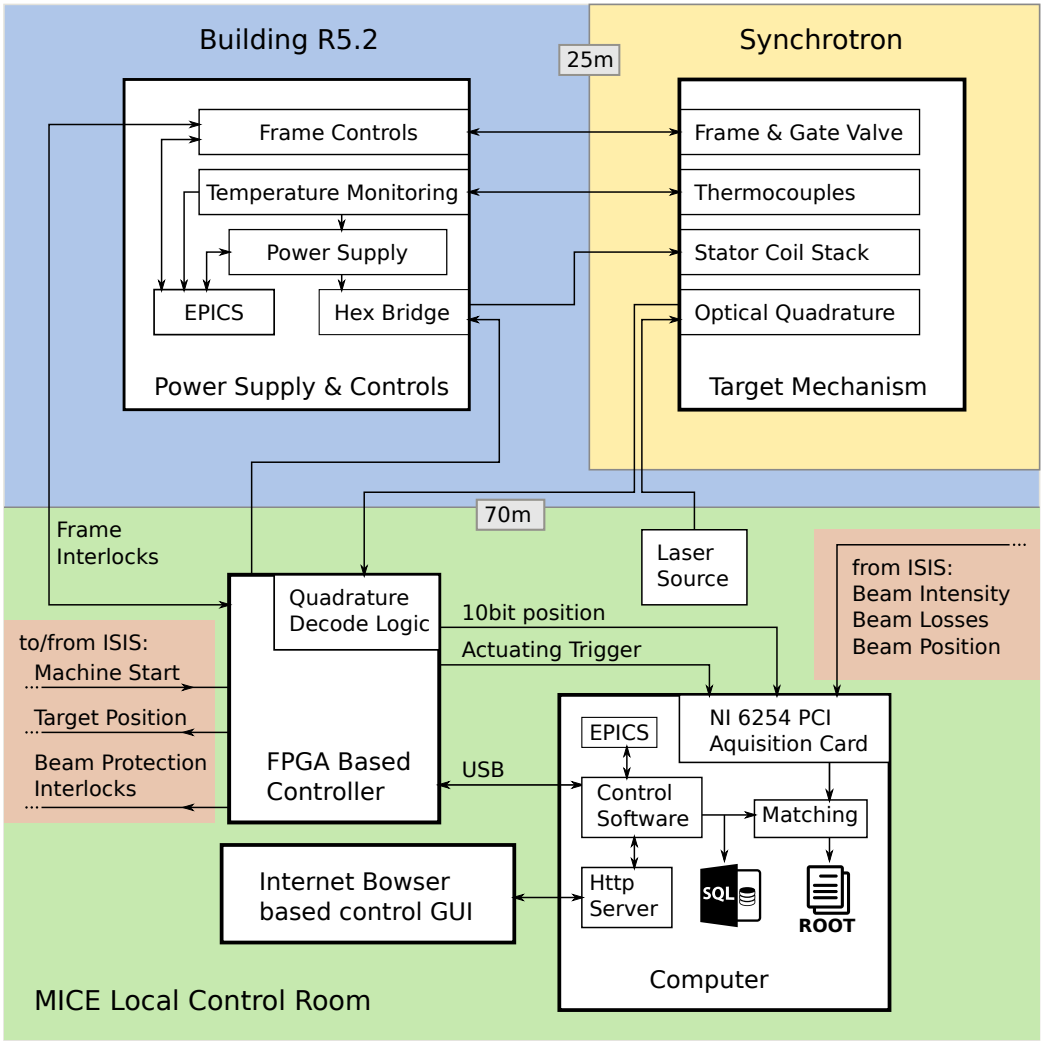

Figure 3.7: Major components of the MICE target system, showing interconnects between the subsystems and ISIS.

Park: Target is in an unpowered state, normally resting on the mechanical stop, with a position around 15 mm BCD. This state is disabled when the target frame is lowered, however un-recoverable errors will result in the target entering this state.

Hold: Target is levitating above the beam, with a position around 64.1 mm BCD. Once in this state the frame can be lowered into ISIS.

Actuate Enable: Target is levitating, but a Machine Start trigger will begin an actuation after a specified delay.

Actuating: The target is currently undergoing high acceleration in order to dip into the ISIS beam.

Once the actuating state is entered a number of sequenced steps are executed, shown in figure 3.8. The first step levitates the target for a predefined time, before the actuation start point is reached and the first quadrant is entered. While in the first quadrant the mechanism applies a large downwards force. Once the target reaches Switch Point 1 the acceleration is reversed which brings the target to a minimum position and passes to the third quadrant. The third quadrant maintains a upward acceleration until the second switch point is reached, where the acceleration is again reversed. Once inside quadrant four a downward acceleration is maintained until the target begins moving downwards, indicating the maximum position has been reached. The "kickback" step applies a upward acceleration to reduce the residual downwards motion. Finally the system enters a "capture" step, where the motor is driven identically to the levitate state. In this state any residual energy from the actuation is seen as an oscillation in position, which is passively damped. The target system remains in this step for a considerable time to ensure the oscillation is sufficiently decayed before initiating the next actuation.

The precise timing and depth of the target is adjusted by changing the delay to the start point and the position of the first switch point. During the actuation a number of parameters from the state changes are logged, some of which are interlocked to protect ISIS in the event of a rogue actuation. Time is measured by the controller's 40 MHz crystal oscillator, making the resulting granularity negligible.

The USB link allows control over the mechanism, by changing between three of the four states. The link also enables the readout of information collected during actuations and configuration of the actuation parameters and advanced settings.

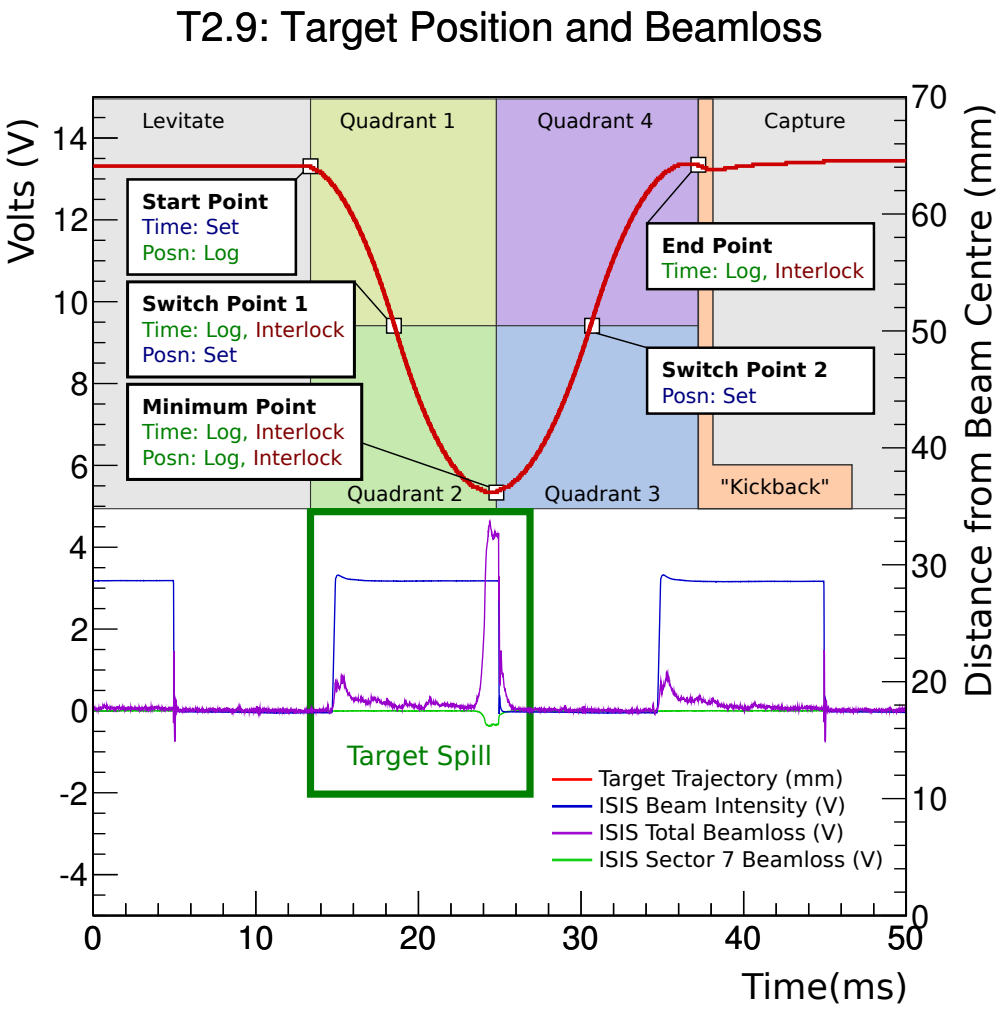

Figure 3.8: Annotated target trajectory, showing the key actuation states and the transitions between. The ISIS beam intensity and losses are shown below.

Control software present on a local computer sets and receives information over the USB link and enables control from a user friendly HTML/Javascript interface.

3.1.7 Data Acquisition

During actuation data is collected by both the control system and a waveform digitiser to monitor the performance of the mechanism.

Several key points on the actuation trajectory are logged in the controller data (see figure 3.8) and these are augmented with further information, such as the laser intensity of each quadrature channel. This data is then readout by the controller's USB link and written immediately to a MySQL database.

In order to also record the exact target trajectory and beam-loss levels within the synchrotron, a National Instruments (NI) 6254 PCI card is used. The position of the target is directly output from the quadrature decoder over a 10-bit parallel connection to the card and sampled at 200 kS/s. Signals from ISIS are read out simultaneously by the card using a number of ADC channels, each of which samples at 100 kS/s. The recorded signals are as follows:

- Beam Intensity
- Total Ring Beamloss
- Beamloss in Sector 7 (target's location)
- Beamloss in Sector 8 (immediately downstream from the target)
- Horizontal/Vertical beam position in sector 7.

The digitized waveform data is matched in real time with data from the target controller and appended to a ROOT file. Each hour a new ROOT file is written, in order to keep the file size manageable.

3.1.8 Testing Programme

In order to test and verify the target mechanism for use on ISIS, a testing platform was constructed in building R78. A complete mock-up including a section of beam pipe, ion pump, target frame and control systems has been built. This system is identical in as many ways as possible and also operates as drop in replacements for the systems on ISIS.

The platform reproduces the hard vacuum present on ISIS and has allowed an extensive testing and development program. Each test is denoted by the stator identifier and bearing number, for example S1.2 is the test of stator S1, bearing set 2. A new unused set of bearings is fitted for each test. In addition a standard set of test conditions are also used, whereby the target actuates at a rate of 0.83 Hz and dips over the full range of the shaft, to a minimum position at around 20 mm BCD. The test also includes a one hour daily stop, which allows the target to cool and closely mimics the non-continuous running of the production target on ISIS.

As part of the Quality Assurance (QA) procedure the target is checked on a daily basis to spot early signs of failure and arrange a swap over. Production targets are also tested for a single day in R78 and a decision is made on their suitability for ISIS based on the single day performance. To test the QA procedures this process is completed for all targets.

Once the test is complete the target is carefully inspected for evidence of dust and a photograph taken of each of the bearings, shaft and dust catcher for future reference. After switching to the DLC/VESPEL combination each test has shown a small amount of dust due to wear, but not sufficient quantities to be of concern.

3.2 Target Numerical Model

Understanding and interpreting the performance of the MICE target is an involved task, due to the number of physical, electronic and control effects present in the output data. A model which allows a range of effects to be included, adjusted and ignored is a useful tool in the study of the mechanism.

The goal of this model is to mimic the behaviour of the real mechanism, inclusive of effects from the control system. The mimic data can then be converted into a waveform which can be used to test analysis code used on the real data.

3.2.1 Model Design

The use of a numerical code was chosen, since this provides the greatest flexibility and would enable modelling of the control system logic. At the heart of the model is a fourth order Runge-Kutta code, which determines how the position and velocity change given a calculated acceleration. The acceleration (a) of the mechanism can be described as a function of position (x), velocity (v) and time (t) using the following

equation:

$$a(x,v,t) = a_0 C(x,t)\,\exp\left(\frac{-t}{t_{RC}}\right) - Ev - F(x)\,\mathrm{sgn}(v) - g \tag{3.2}$$

The first term accounts for the driving force from the coils; a_0 is the absolute maximum driving acceleration, $C(x,t)$ is how efficiently the acceleration is applied by the electronics where the sign indicates the direction of acceleration, t_{RC} is the discharge time of the capacitor bank. The second term is the effect of eddy currents induced in the target stator by the permanent magnets and is characterised by E, the deceleration per unit velocity. Finally $F(x)$ accounts for friction effects, while g is the acceleration due to gravity.

The value of $C(x,t)$ describes both the direction and fraction of total available acceleration which is driving the motor. It is evaluated by summing the product of the fraction of coil current in every possible state from the hex bridge with the fractional force each of those states delivers, shown below:

$$C(x,t) = \sum_{i=0}^{5} I_i \sin\left(2\pi \frac{x[\mathrm{mm}] + 3i}{18}\right) \tag{3.3}$$

where I_i is the fraction of total current in each phase of the motor during this iteration of the model. This value is determined by the last time the phase was switched and the coil stack time constant (t_{S}). Note this is caused by the inductance of the coil stack and is significantly different from the capacitor decay time(t_{RC}). In order to calculate the switching time of each phase, a simplified mimic of the control logic has been implemented within the model. There is no normalisation in the equation because only a single state is active at a single time; the summation allows the code to account for residual current in unpowered states. During each iteration of the target model the following steps are taken:

1. The true position of the target is converted to a quadrature number, with a pitch of 150 μm in a similar way to how the quadrature decoder circuit operates.

2. A state machine monitors the quadrature position and mimics the state changes present in the real controller. Depending on the current state a direction for acceleration is determined. Records are also kept at each state transition for a "true" measurement.

3. The quadrature position is used to determine the optimum phase of the motor to be used and switch the active phase as necessary.

4. The values of I_i for the current simulation step are then calculated, based on when the phase was activated.

5. Finally $a(x, v, t)$ is evaluated using the above equations, and is then passed to the Runge-Kutta code.

3.2.2 Frictional Model

In the real target, the friction induced by the dry bearings and shaft depends on a large number of factors including: specific materials, surface roughness, velocity of travel and temperature. Furthermore it is fair to assume that the perpendicular load on the bearings is also not constant as the shaft moves. Under non-vacuum conditions the coefficient of friction between a similar VESPEL poly-amide and DLC was observed to vary between 0.55 and 0.80[62] as a function of velocity. Stiction effects are of no consequence to this study since data is only collected when the system is moving rapidly.

Without detailed measurements of the friction using the exact materials and conditions used in the target system, it is an involved task to estimate what variations of friction should be considered. Under ideal operating conditions the transverse forces on the system are zero, however transverse forces can be created by a misalignment between the coil and magnet centres. The magnetic model of the target system indicates that the largest transverse force is due to the torque, which varies from 27 mN m to 32 mN m for an offset of 0.5 mm and scales linearly[59] with offset.

The model adopted by the numerical code is based on coulomb friction, with a normal force which can vary by several percent and is accomplished by re-evaluating the friction after every 2 mm of travel. The local friction value is selected from a Gaussian random number generator, with a mean of the average friction and width 10% of the average friction. Note that the friction accounts for less than 15% of the overall acceleration.

3.2.3 Model Performance

A single actuation generated by the model is shown in figure 3.9. The position follows a trajectory similar to that of an actual target, demonstrating that the model mimics the steps required to actuate the target. The velocity of the model also shows the expected shape and includes the overshoot at the end of the actuation, derived from

the limited position resolution of the control system. Finally the acceleration of the mechanism shows several modelled effects:

1. The acceleration follows the downward (quadrant 1), upward (quadrants 2 & 3), downward (quadrant 4), upward (kickback) pattern desired to actuate the target.

2. The peak section of a sinusoidal pattern can be seen overlaid on the block structure, which is a result of the driving force changing as a function of position, shown in figure 3.2. With increased velocity the period of the sinusoid decreases, while the switching frequency increases.

3. Friction can be seen as a sudden change in the acceleration at the minimum point, where the direction of motion changes.

4. Eddy currents can be seen as a decrease in acceleration at higher velocities.

5. The slow decay of the capacitor bank can be seen in a decrease of maximum acceleration over time.

6. Finally the effect of the pulse width modulation can be observed as a number of spikes in the acceleration at the end of the actuation. The coil rise time prevents the acceleration reaching the maximum value and results in a sawtooth shaped waveform.

The final step in understanding the model performance was to determine the optimum step size, for use in testing the analysis code. Due to the complicated nature of the model, the result is a trajectory which is difficult to predict analytically. Adopting a simplified model which is easier to predict does then not test the detailed effects which are of interest. To select the optimum step size the model was run with identical starting conditions, while varying the step sizes to look for convergence. The convergence of the model was determined by the point at which the target enters the capture stage of the actuation. Figure 3.10 shows that model converges with variations below 5 μm and 1 μs when a step size below 100 ns is used, and a value of 50 ns was selected.

3.3 Analysis Techniques

Data collected by the target system comes in two complementary forms. The first is collected during the actuation directly by the control system (illustrated in figure 3.8).

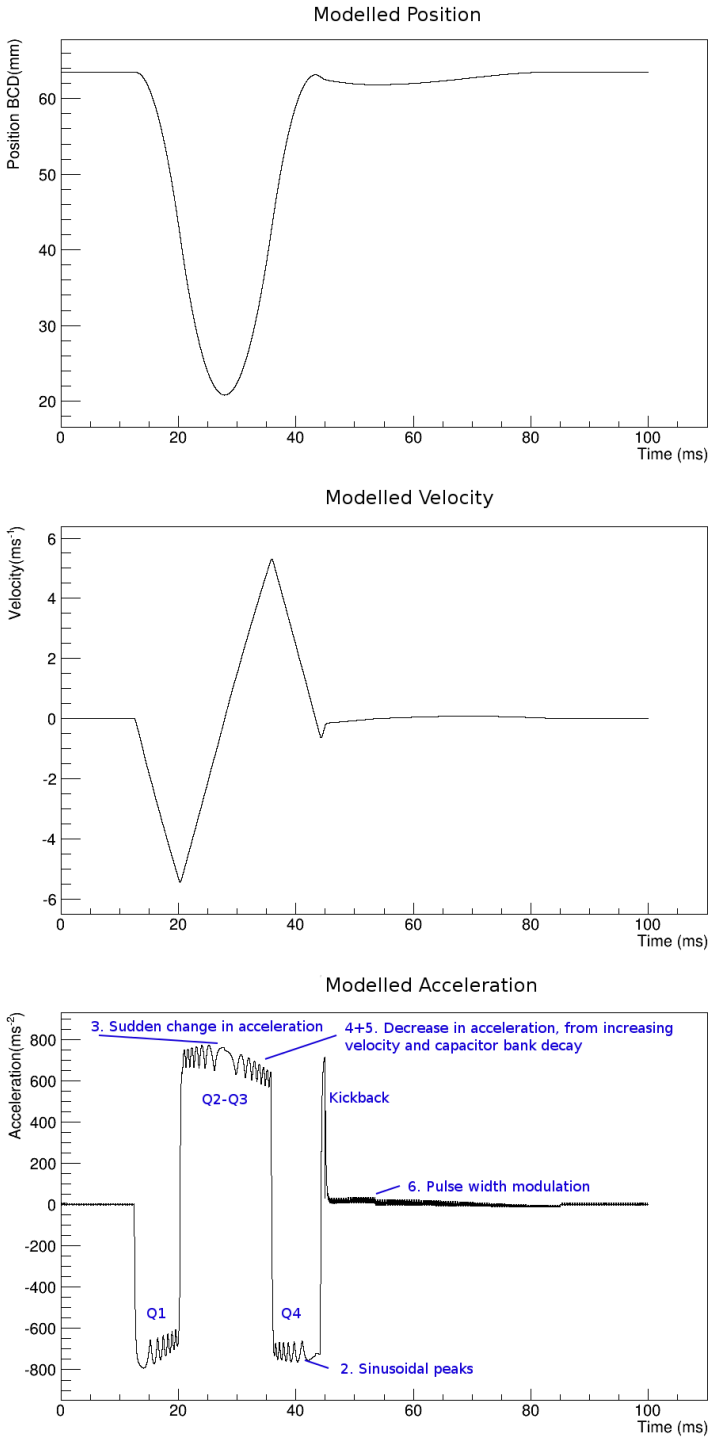

Figure 3.9: Image of the position, velocity and acceleration as a function of time as generated by the numerical model.

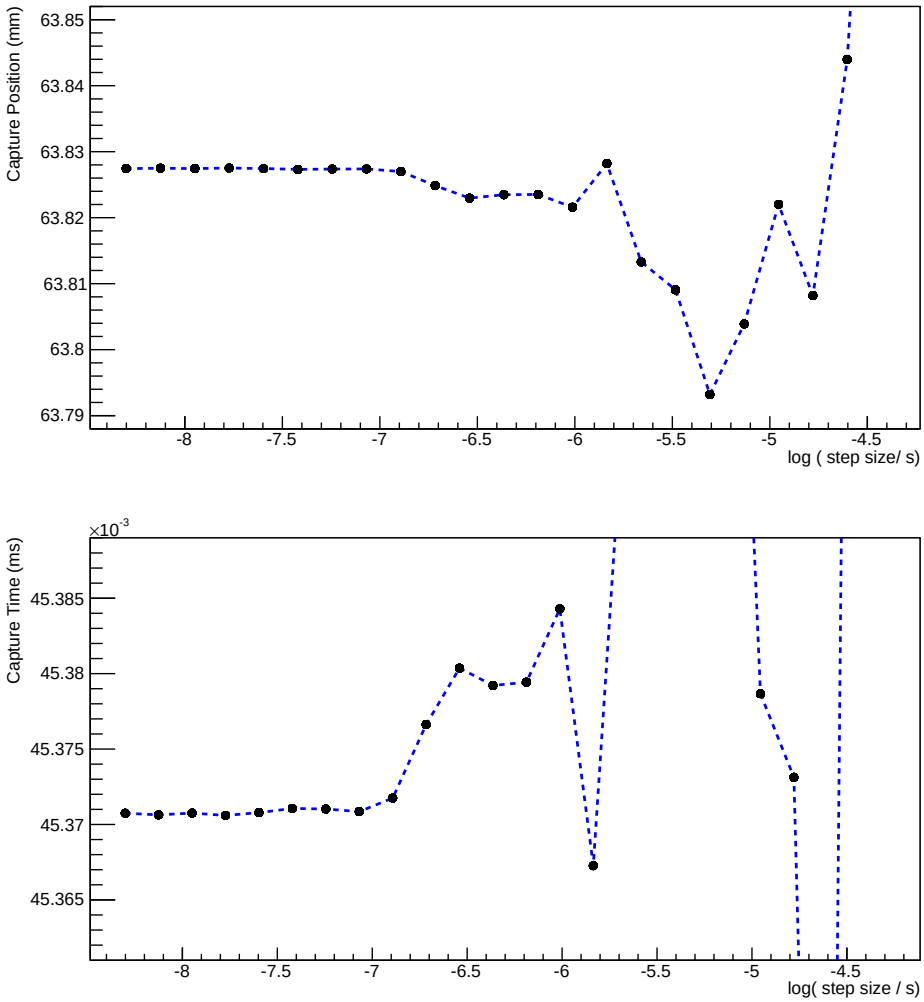

Figure 3.10: Position and time at which the model enters the final capture stage of the actuation as a function of model step size. Step sizes below 100 ns converge to a common value.

Distance (mm)	Time (ms)	Velocity at SP1 ($\mathrm{m\,s^{-1}}$)	$\Delta\overline{a_{\mathrm{SP1}}}$ ($\mathrm{m\,s^{-2}}$)
5.0	3.8	2.6	14
10.0	5.3	3.7	6.7
15.0	6.5	4.6	4.5
20.0	7.6	5.3	3.9
25.0	8.4	5.9	2.7

Table 3.2: Uncertainty of $\overline{a_{\mathrm{SP1}}}$, where $\overline{a_{\mathrm{SP1}}} = 700\ \mathrm{m\,s^{-2}}$ for a number of distances to switch point one.

The small quantity of data collected is ideal for fast analysis and long term monitoring. The second source of data is waveforms collected by the NI card. Each of these waveforms contains thousands of points and requires fitting, making the analysis a more involved task, but also capable of yielding more useful results.

3.3.1 Controller Data

Acceleration

In order to accurately determine the acceleration of the mechanism it is important to select two data points which are well defined in time. For this reason the starting position is used as the first point. The second point is the first switch point, where the high velocity ensures an accurate time measurement. The minimum position is not useful for timing related tasks, since the low velocity and quadrature effects contribute a large amount of uncertainty.

Since the velocity of the target is zero at the start of the actuation the mean acceleration in the first quadrant is simply:

$$\overline{a_{\mathrm{SP1}}} = 2\frac{x_{\mathrm{Start}} - x_{\mathrm{SP1}}}{(t_{\mathrm{Start}} - t_{\mathrm{SP1}})^2} \tag{3.4}$$

where x_{Start} is the start position, x_{SP1} is switch point one position, t_{Start} is the starting time and t_{SP1} is the time at switch point one. The systematic error from both the time and position measurements vary as a function of dip depth and the resultant uncertainty is shown in table 3.2. For nominal R78 running, the uncertainty of the acceleration measurement is $\approx 0.5\%$.

Variation in excess of the systematic error is typically seen during operation and is due to the fact that the acceleration measurement is averaged over time. At the beginning of the actuation the target is initially still and as a result a considerable

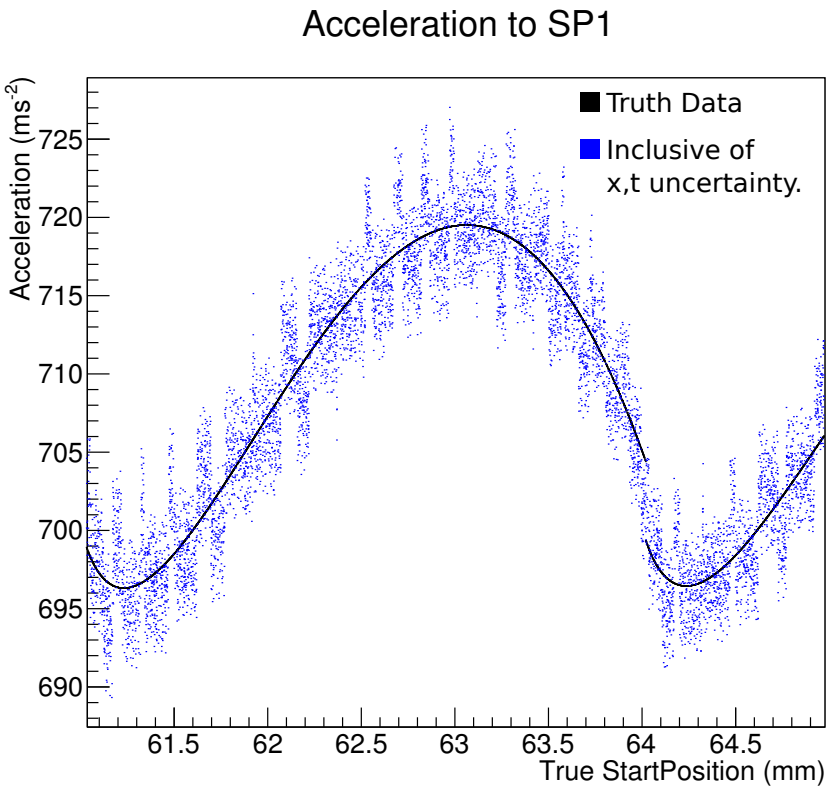

Figure 3.11: Systematic variation observed in mean acceleration to SP1, induced by different starting locations. The x and t uncertainty is a result of quadrature effects. The truth data is the average acceleration calculated from the exact starting position/time and switch point one position/time, recorded internally in the model.

fraction of time is spent in close proximity to the starting position. As previously discussed, the acceleration changes as the vertical position of the shaft changes and this introduces a notable variation in the average acceleration. The effect of this is shown in figure 3.11 which uses the numerical model to demonstrate this in detail. A clear variation of around 20 $\mathrm{m\,s^{-2}}$ is observed in the truth data and is also visible after the application of quadrature effects. This variation can be removed by restricting the range of start positions which are used to calculate the acceleration to set point one.

The driving acceleration of the mechanism is also heavily influenced by the driving voltage and temperature of the coils. The temperature changes depend on the heat load in the coils which in turn changes with the frequency of actuations and dip depth.

To provide a consistent testing environment the coil voltage, the rate of actuations and dip depth are fixed, which in turn fixes the acceleration. This enables the acceleration to be used in the diagnostic for checking the target's health; for instance a sudden change in acceleration can imply a sudden change in the bearing condition. In the case of ISIS operation the target is frequently stopped, or the dip depth changed, which makes changes in acceleration more difficult to interpret.

Start Position Width

In order to aid interpretation of the acceleration a second measurement is needed, which can be determined from the width of the start position. At the end of each actuation the target is trapped in a potential well, within which it oscillates while excess energy is dissipated via friction and eddy currents. For the target to stop completely the restoring force from the potential well must be less than the static friction present within the bearings. This condition can be described with an inequality:

$$F_s > |k(x - x_c) + g| \tag{3.5}$$

where F_s is the static friction, k is similar to a spring constant, x is the position of the target, x_c is the centre of the well and g is the acceleration due to gravity.

The width of the allowed stopping positions can then be written as:

$$W = \frac{2F_s}{k} \tag{3.6}$$

where W is the full width of the finishing position. This equation predicts that an

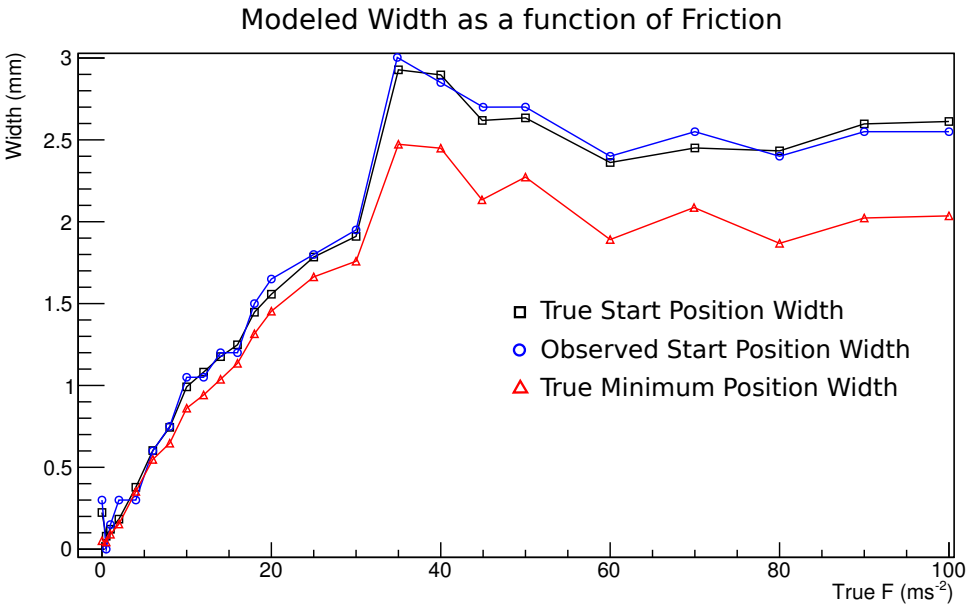

Figure 3.12: Effect of increased friction on the starting position width, showing a steady increase. The width of the minimum position also shows a similar trend.

increasing friction leads to a wider distribution of stopping positions.

The stopping position of the target can be seen at the beginning of the following actuation as the starting position. If the target is caught outside the coil switching pitch a correction is made and the target brought back within this range.

Figure 3.12 shows the starting position width over two thousand actuations for a number of different friction settings simulated by the model. One can see a steady increase as a function of friction up to a maximum of 3 mm, where the width begins to reduce due to the capture correction mechanism included in the model.

A problematic side effect of this is caused by symmetry around the switch point where a higher start position means a greater acceleration distance. This causes a higher velocity at the first switch point and hence requires a larger stopping distance, resulting in an increase in dip depth. An increase in variability around the minimum will also increase the variability of the beam-loss generated by the target, which makes finding an optimum operational point more challenging. To mitigate against this a small addition to the controller firmware was added to fix the acceleration distance by inserting a short variable dead time around the first switch point.

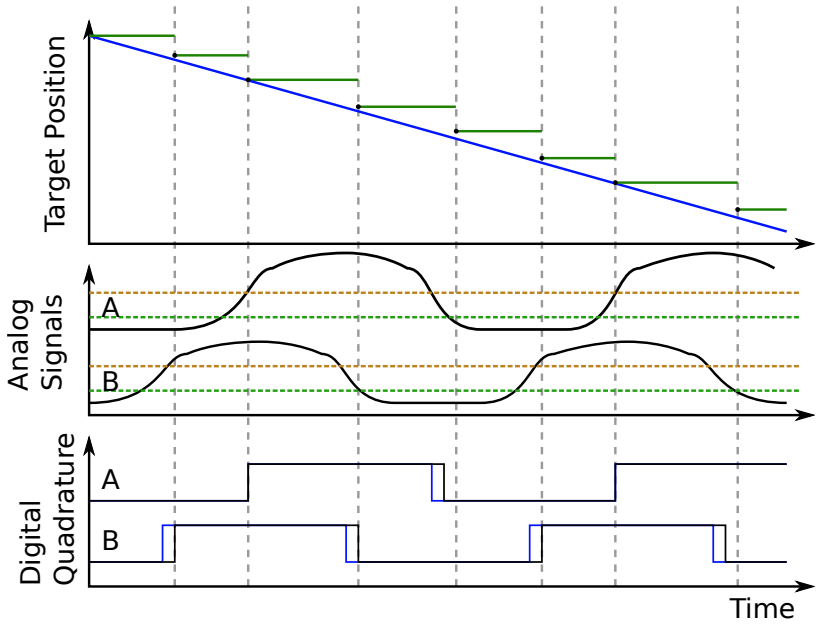

Figure 3.13: Effect of brightness variations on the location of each quadrature edge and hysteresis effects. The ideal waveform is shown in blue.

3.3.2 Waveform Data

In addition to the parametrised actuation data, a waveform of the mechanism is also captured. Fits to this waveform can be used to yield much greater insight into the performance of the mechanism. To gain the most out of the data, it is prudent to first condition and optimise the raw position data acquired before further analysis.

Position Tracking

The digitized quadrature signal is processed before it is used in further analysis, which is motivated by the following:

1. The data is collected by asynchronously sampling a 10-bit integer value at 200 kS/s. If two or more bits change simultaneously there is a possibility that at least one of the bits will be sampled erroneously, which will in turn lead to a bad data point. In practice this tends to cause a rogue data point far from the trajectory.

2. The sampled points contain many data points at identical locations as the target is moving slowly. Over a small fit range this can noticeably bias the fit result,

which is undesirable.

To prevent these issues from influencing the analysis two steps are taken. First each data point is replaced with a median value between each data point and its neighbours. Under normal conditions this has no effect, however in the event of a read error, the point will be replaced with a sensible value.

Next, the waveform is reduced such that only changes between measured positions are kept (since this is the only point where the position can be determined with substantial accuracy). The exact position of the target which causes a quadrature change is affected by many factors, including the focus and alignment of the optical block, intensity of each quadrature channel and hysteresis of the discriminator circuit. These effects distort the timing of the quadrature waveform, as shown in figure 3.13. The distortion is mitigated by applying a correction to each data point, where the same value is always applied to the same quadrature transition. A distinction is made between directions of travel such that two sets of four correction values are required in total. Between each of the four corrections there are only three unique variables, and therefore to constrain the problem the fourth correction is defined to set the sum of all the corrections to zero. As a result 6 unique values are required to apply the correction.

Calibration of the correction is performed by fitting a number of parabolic curves to the uncalibrated points, while using the offsets as free parameters within the fit. The minimiser then computes the optimum value for each offset in order to draw a smooth curve between the data points.

Figure 3.14 shows a comparison between the raw quadrature data, seen as black horizontal lines, while the enhanced quadrature points from this method are in blue. The plot has been fitted with a polynomial trajectory for reference and clearly shows that the new method follows the curve well. The residual between the polynomial fit and enhanced x position over a large number of actuations is shown in figure 3.15. The estimated time and position resolution of both the RAW and enhanced data is shown in table 3.3. The result of this technique is a more clearly defined trajectory, formed from fewer data points which also substantially increases the speed of chi-square fitting functions.

3.3.3 Trajectory Fit

Monitoring the system over an extended period of time enables changes to be identified, but does not enable an immediate measurement of the bearing properties. An

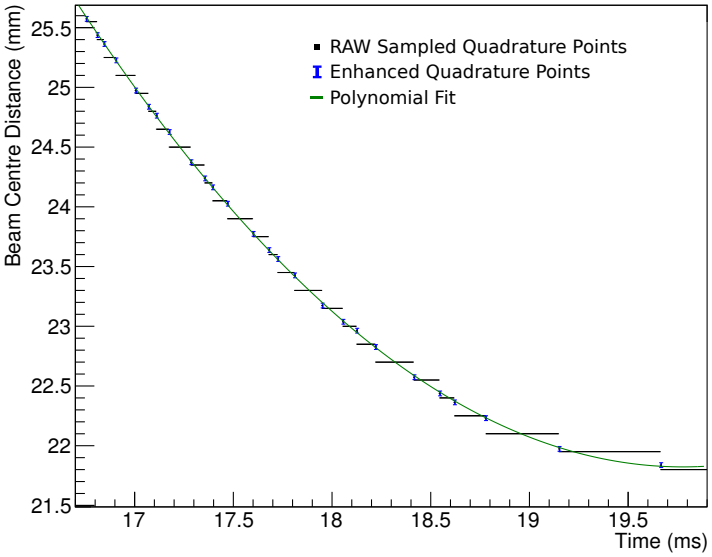

Figure 3.14: Waveform comparison between raw quadrature points and the enhanced calibrated quadrature switch points.

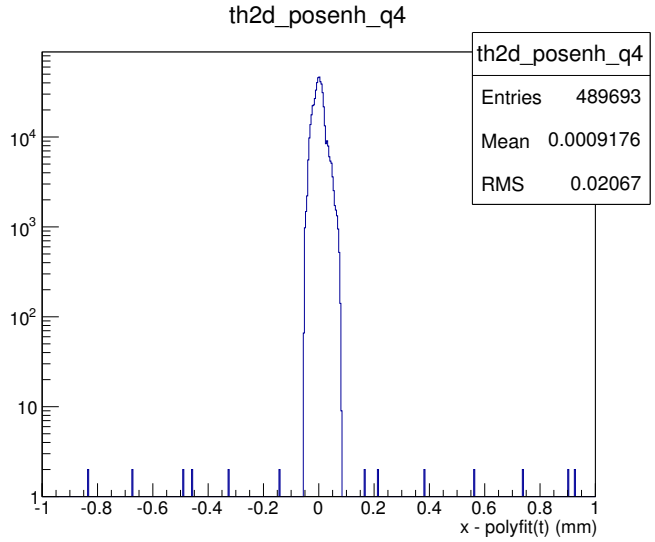

Figure 3.15: Residuals between the estimated position from a parabolic fit and the calibrated enhanced quadrature data points.

Parameter	RAW Data	Enhanced Quadrature
Position Resolution(μm)	43 (granularity: 150)	20
Time Resolution(μs)	1.4	1.4
Data points	10000	<650

Table 3.3: Quadrature comparison between RAW quadrature data and the enhanced data points

ideal measurement to make would be the friction within the bearings at the start of running. This information can be extracted by considering the trajectory and forces present in each quadrant of the mechanism, which this section will explore in detail.

Friction Fit

During the first quadrant of the actuation the target is accelerated downwards, with friction opposing the direction of motion and with gravity always acting downwards. To mitigate systematic effects introduced by eddy currents and coil switching times the acceleration is evaluated when the mechanism is stationary. The acceleration of the mechanism at the stationary points in each quadrant are shown below:

$$a_1(t = t_{\text{start}}) = -a_0\overline{C(v=0)}\exp\left(\frac{-t}{t_{\text{RC}}}\right) + \overline{F(x)} - g \tag{3.7}$$

$$a_2(t = t_{\text{min}}) = +a_0\overline{C(v=0)}\exp\left(\frac{-t}{t_{\text{RC}}}\right) + \overline{F(x)} - g \tag{3.8}$$

$$a_3(t = t_{\text{min}}) = +a_0\overline{C(v=0)}\exp\left(\frac{-t}{t_{\text{RC}}}\right) - \overline{F(x)} - g \tag{3.9}$$

$$a_4(t = t_{\text{end}}) = -a_0\overline{C(v=0)}\exp\left(\frac{-t}{t_{\text{RC}}}\right) - \overline{F(x)} - g \tag{3.10}$$

where a_q is the acceleration in the quadrant (q). a_0 is the absolute maximum driving acceleration. $\overline{C(v=0)}$ is the average fraction of acceleration delivered to the shaft when the velocity of the mechanism is zero. t_{RC} is the capacitor bank decay time constant. $\overline{F(x)}$ is the average friction measured and finally g is the acceleration due to gravity.

Each a_q can be estimated by the fit of a third order polynomial function $(x_q(t))$ to the target trajectory in each quadrant of the actuation. The acceleration is then simply the 2^{nd} derivative of the function, evaluated at each stationary point, specifi-

cally:

$$x_q(t) = \sum_{i=0}^{3} p_{q,i} t^i \tag{3.11}$$

$$a_q(t) = 6p_{q,3}t + 2p_{q,2} \tag{3.12}$$

The third order polynomial fits are capable of measuring a_q at each stationary point to approximately 1.5%, or around 10 $\mathrm{m\,s^{-2}}$ of the nominal 700 $\mathrm{m\,s^{-2}}$ of acceleration. A chi-square minimisation technique is next used to fit estimated values from equations 3.7 to 3.10 to each measured value by varying the parameters a_0, t_{RC}, $\overline{F(x)}$. The purpose and value of $\overline{C(v)}$ is described and calculated in appendix A.

Stator Properties

In addition to parametrising the acceleration and friction of the mechanism it is also possible to extend the method to estimate properties of the stator. In this case it is necessary to consider how the acceleration changes during each quadrant of the actuation. The polynomial fit of the trajectory enables the mean value of this change to be determined by:

$$\frac{d^3x_q}{dt^3} = 6p_{q,3} \tag{3.13}$$

Fit results from the S1 stator indicates that $p_{q,3}$ can be as large as 3700±400 $\mathrm{m\,s^{-3}}$, indicating that the acceleration can change by up to 150 $\mathrm{m\,s^{-2}}$ within a quadrant.

The original fit was to a 3$^{\mathrm{rd}}$ order polynomial, meaning the third derivative is a constant representing the mean change in acceleration across a quadrant. The mean change can be calculated simply by:

$$\overline{\frac{d^3x_q}{dt^3}} = \frac{1}{t_2 - t_1}\int_{t_1}^{t_2} \frac{d^3x_q}{dt^3} dt \tag{3.14}$$

$$= \frac{1}{t_2 - t_1}\left(\frac{d^2x_q}{dt^2}(t_2) - \frac{d^2x_q}{dt^2}(t_1)\right) \tag{3.15}$$

which can then be applied to each quadrant. Both friction and gravity are constant within each quadrant, meaning their values can be neglected. The mean change in

acceleration in each quadrant is as follows:

$$\overline{\frac{d^3x_1}{dt^3}} = \frac{-a_0\left(\overline{C(v_{\mathrm{SP1}})}\exp\left(\frac{-t_{\mathrm{SP1}}}{t_{\mathrm{RC}}}\right) - \overline{C(v_{\mathrm{Start}})}\exp\left(\frac{-t_{\mathrm{Start}}}{t_{\mathrm{RC}}}\right)\right) - E(v_{\mathrm{SP1}} - v_{\mathrm{Start}})}{t_{\mathrm{SP1}} - t_{\mathrm{Start}}} \tag{3.16}$$

$$\overline{\frac{d^3x_2}{dt^3}} = \frac{+a_0\left(\overline{C(v_{\mathrm{Min}})}\exp\left(\frac{-t_{\mathrm{Min}}}{t_{\mathrm{RC}}}\right) - \overline{C(v_{\mathrm{SP1}})}\exp\left(\frac{-t_{\mathrm{SP1}}}{t_{\mathrm{RC}}}\right)\right) - E(v_{\mathrm{Min}} - v_{\mathrm{SP1}})}{t_{\mathrm{Min}} - t_{\mathrm{SP1}}} \tag{3.17}$$

$$\overline{\frac{d^3x_3}{dt^3}} = \frac{+a_0\left(\overline{C(v_{\mathrm{SP2}})}\exp\left(\frac{-t_{\mathrm{SP2}}}{t_{\mathrm{RC}}}\right) - \overline{C(v_{\mathrm{Min}})}\exp\left(\frac{-t_{\mathrm{Min}}}{t_{\mathrm{RC}}}\right)\right) - E(v_{\mathrm{SP2}} - v_{\mathrm{Min}})}{t_{\mathrm{SP2}} - t_{\mathrm{Min}}} \tag{3.18}$$

$$\overline{\frac{d^3x_4}{dt^3}} = \frac{-a_0\left(\overline{C(v_{\mathrm{End}})}\exp\left(\frac{-t_{\mathrm{End}}}{t_{\mathrm{RC}}}\right) - \overline{C(v_{\mathrm{SP2}})}\exp\left(\frac{-t_{\mathrm{SP2}}}{t_{\mathrm{RC}}}\right)\right) - E(v_{\mathrm{End}} - v_{\mathrm{SP2}})}{t_{\mathrm{End}} - t_{\mathrm{SP2}}} \tag{3.19}$$

where $\overline{C(v)}$ is the average fraction of acceleration delivered to the shaft as a function of velocity; note that this quantity also depends on the coil rise time (t_{S}), and is calculated in appendix A. E is a parametrisation of eddy currents and is the deceleration per unit of velocity.

Again a chi-square fit with non-linear minimiser can be used to fit the estimates to the measurements in each quadrant and determine the unknown parameters: E, t_{S}, a_0, t_{RC}. Since two of the unknown parameters are in common with the friction fit, both fits are performed simultaneously by minimising the sum of both chi-square values.

Fit results

The fit results for a test of S1.5 in R78 are shown in figure 3.16, that shows the chi-square from the fit and the constant parameters, E, t_{S}, t_{RC}. The acceleration and friction quantities are time dependent in a real target and have not been investigated here.

The resultant chi-square of the combined friction and eddy current fits are shown in figure 3.16. There were a total of six free parameters in each of the fits, which results in two degrees of freedom. The chi-square probability distribution has also been plotted for 2 degrees of freedom and rescaled vertically to fit the data. Low values of χ^2 fit the distribution well, however there is a clear excess of fits with a high value of chi-square. The excess indicates that not all actuations are modelled well by the fit and a cut is placed to only accept fits with a chi-square less than 20.

The fixed mechanical properties are shown in the remaining plots of figure 3.16.

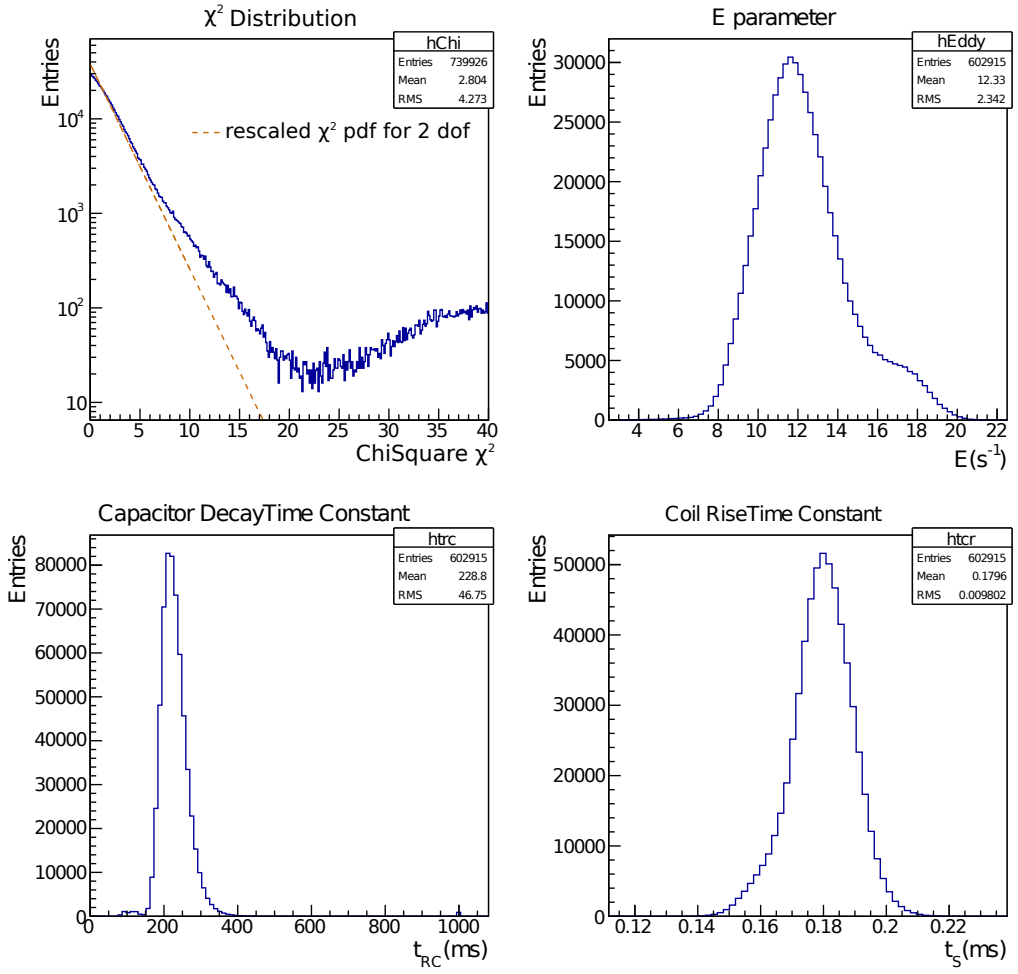

Figure 3.16: Fit of the S1.5 run in R78, showing the resultant chi-square and constant fit parameters. The dotted orange line indicates the expected chi-square probability density function for two degrees of freedom.

Both the capacitor decay time constant and coil rise time constants are centred on a single sensible value, as expected. The fitted eddy current parameter appears as a Gaussian distribution, with a small shoulder present. The shoulder could be the result of a misinterpretation of a quadrant causing a systematic offset, or the manifestation of another velocity dependent parameter which has not been considered by this analysis. An example of such a parameter could be a velocity dependence of the friction term, however such a dependence would be impossible to disentangle outside a controlled test environment.

To understand correlations and systematic offsets within the fitted parameters, the numerical model was run and the output passed to the analysis code. Since both

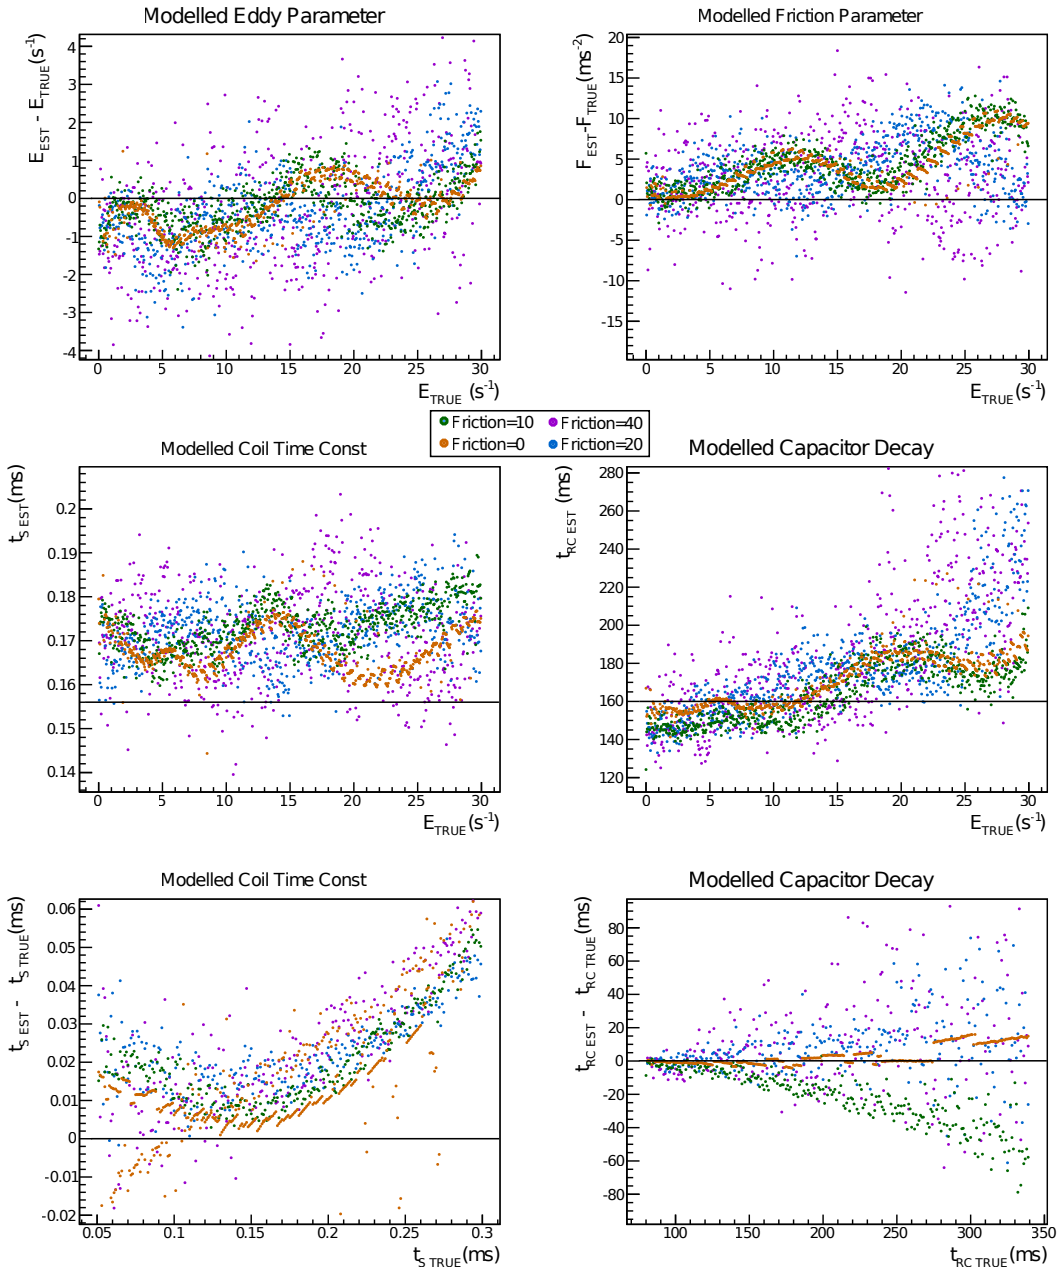

Figure 3.17: Study of fit performance for a total of four friction settings, over a range of E, t_{RC} and t_{S} values.

the friction and eddy currents have a common resistance to the direction of travel, care has been taken to specifically avoid interference between these effects. To verify this the numerical model was run for a range of eddy current settings and four different friction settings.

Figure 3.17 shows the results of this scan and for the most part indicates a good degree of independence between the parameters. The discrepancy between fit and truth for the eddy parameter systematically varies by around $\pm 2\ \mathrm{s}^{-1}$, while the friction varies by around $-1\ \mathrm{s}^{-1}$ to $6\ \mathrm{s}^{-1}$. This variation is particularly noticeable in the zero friction case, the source of this is due to the varying force as a function of position, which is approximated by the $\overline{C(v)}$ function. As the friction is increased, there is an increase in random variation between the points, however this introduces no systematic offset. The source of this variation is the random numbers included by the friction inside the model and can be reduced by averaging over many actuations.

A study of the effect of eddy current and friction was also performed on the coil rise time and capacitor decay time. Both these parameters showed a subtle variation, similar to that discussed previously and induced by the same cause. Overall the coil rise time returned a value between 2% and 15% higher than the true value of 156 μs, the offset could be due to dead time which is not fully accounted for in the $\overline{C(v)}$ approximation. The capacitor decay time registered a value within -20 ms to 60 ms of the true value at 160 ms, the variation included by the friction model increases the spread in points, with bias towards a larger time constant. This effect is because there is only a small reduction in acceleration from this effect during the actuation (20%).

A final study was made of both the coil rise time and capacitor decay values over their typical range. The fitted coil rise time is between 0 μs to 25 μs larger than the true value, with an increasing discrepancy as the true t_s is increased, at around the 10%. A source of this could be that increasing t_s in the model has shifted the ideal points to switch the active coil state, leading to an additional decrease in acceleration. Finally the capacitor decay time seems to vary by around $\pm 10\%$, which worsens as the decay time increases, which is due to the measured effect becoming smaller and harder to measure.

Ultimately the analysis code is parametrising the behaviour of the target mechanism and then using this to study a number of small effects. This study has shown that these parameters can be measured reasonably well, with systematic offsets around 10 % to 15 %. These offsets are caused by approximations used in the fit which must be present to reduce the number of fit parameters to a manageable number. It

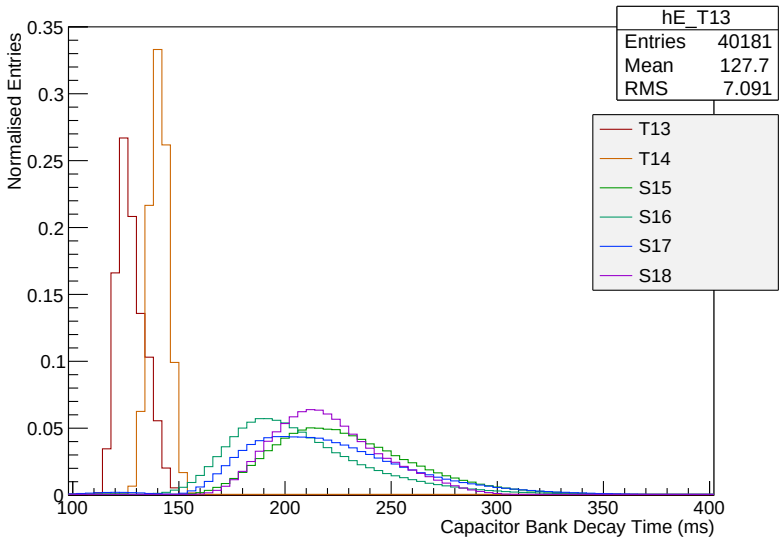

Figure 3.18: Capacitor bank decay time as measured for different runs of the T1 and S1 stators in R78.

may be possible to improve these by using a more sophisticated evaluation of $C(v)$, however this would require precise knowledge of the points which the motor states switch.

3.4 Stator comparison

During the course of the target system development a number of different stators have been constructed. Direct measurements of the properties of these devices are involved and would require the construction of a number of dedicated test rigs. The techniques developed in the previous section enable a brief study of some of these properties.

Figure 3.18 shows the value of the decay time for a number of runs of the T1 stator and the S1 stator. The T1 stator has a decay time around 140 ms, and a simple estimate from the coil stack direct current resistance and capacitor bank size yields 100 ms. For the S1 stator the fit returned a value of around 200 ms, while the simple estimate yields 160 ms. A probable source of the discrepancy is that the simple estimate does not consider alternating current effects, which limit the average current draw and increase the effective resistance.

Figure 3.19 shows a similar plot for the rise time of the coils within the stator's

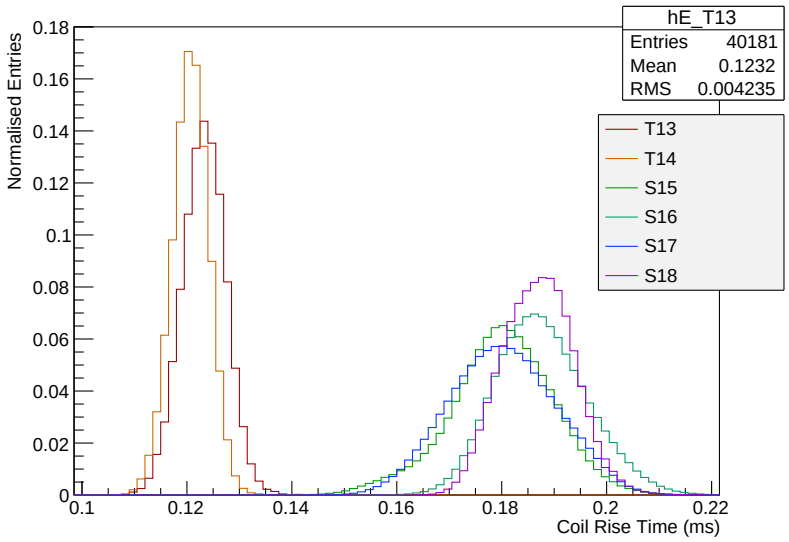

Figure 3.19: Coil rise time measurement for different runs of the T1 and S1 stators in R78.

coil stack. The fit suggests that T1 targets have a rise time around 120 μs, while the S1 stators have an increased value at 180 μs. Again the coil resistance and inductance measurements estimate a rise time of 146 μs and 156 μs for the T and S stators respectively. Note that the simple estimate does not include cable or transistor response times.

The final parameter in the study was the effect of eddy currents on different stators, shown in figure 3.20 using the same runs as previous figures. Here the variation of the eddy parameter is larger, however there is a noticeable difference between the two stator types. The original T1 stator fits with an eddy parameter between 3 s^{-1} to 7 s^{-1}, while the S1 stator has a value between 7 s^{-1} to 17 s^{-1}.

The results are summarised in table 3.4 and show a distinct difference in performance of each stator. The older T1 stator has more responsive coils and lower damping effects from eddy currents. Increased eddy currents in the S stator can be explained by a reduced inner bore radius, placing the permanent magnets in closer proximity to conductive metals. Similarly, the increased coil rise time is likely caused by increased number of windings and smaller wire cross section.

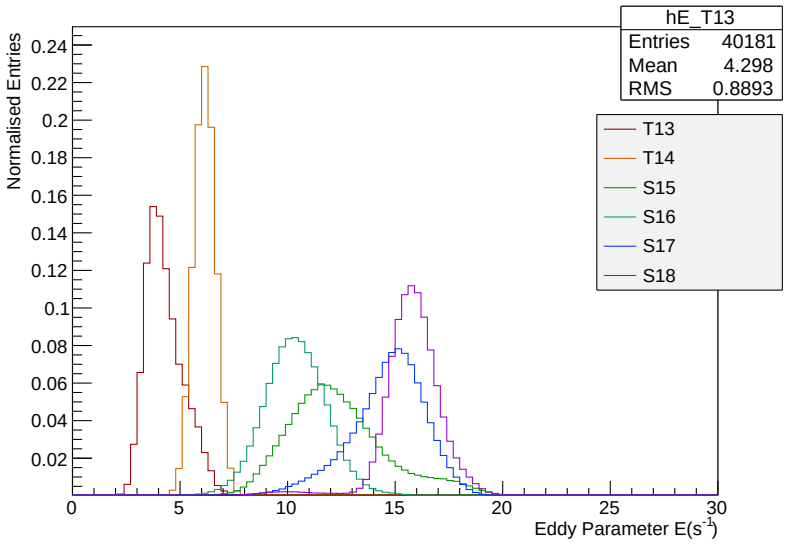

Figure 3.20: Eddy current parameter comparison for the T1 and S1 stators.

Target	Capacitor Decay Time (ms)	Coil Rise Time (μs)	E Parameter (s^{-1})
S1.4	172 ± 20	185 $^{+0}_{-25}$	10.6 ± 2
S1.5	216 ± 21	181 $^{+0}_{-25}$	11.8 ± 2
S1.6	192 ± 20	187 $^{+0}_{-25}$	10.4 ± 2
S1.7	205 ± 21	180 $^{+0}_{-25}$	15.1 ± 2
S1.8	215 ± 22	187 $^{+0}_{-25}$	15.8 ± 2
T1.3	123 ± 20	123 $^{+0}_{-25}$	3.8 ± 2
T1.4	141 ± 20	121 $^{+0}_{-25}$	6.1 ± 2

Table 3.4: Performance of the T1 and S1 stators in a number of tests in R78. The notable source of uncertainty is from systematic effects in the fitting procedure and is estimated from the numerical model. The asymmetric errors are adopted because the model indicates this systematic offset to only be an over estimate.

3.5 Target Performance Evolution

The primary goal of developing the waveform analysis was to directly measure changes to the bearing's performance, as a means to detect early signs of failure. To study in detail the frictional changes the free parameters t_{RC}, t_{CR}, g and E from the previous sections are fixed. This leaves the driving acceleration (a_0) and friction (F) as free parameters in the fit.

3.5.1 Initial running

Figure 3.21 shows the performance changes during the initial 150 thousand actuations of the S1.6 run. These actuations were performed as part of the normal testing regime, where the target is dipped at a rate of 0.83 Hz, with a dip depth of 20 mm BCD and ran continuously, with a one hour stop after each day. The magnitude of the acceleration is shown in the first plot. Over time the magnitude of acceleration in quadrants 2 and 4 increased, while a decrease was observed in quadrants 1 and 3. This change is interpreted as an increase in friction by the fit, meaning the driving acceleration (a_0) remains relatively constant. The second plot in the figure shows the estimated friction over an identical range of actuations and shows a steady increase to a maximum of 3 N. The final plot shows the measured width of both the starting and minimum position; as expected both show an increase as the measured friction increases.

For comparison with a better performing target figure 3.22 shows the initial actuations of the S1.8 target test which concluded successfully after 9 million actuations. The test was terminated to allow S1.9 to be fitted and tested prior to installation on ISIS. Inspection of the S1.8 target showed little dust, and was deemed satisfactory by ISIS. During the initial running the acceleration remained reasonably constant and a maximum friction of 0.3 N was observed.

Finally the results show a consistent difference between the acceleration in the first quadrant as measured by the third order polynomial fit (a_1) and that measured by the controller ($\overline{a_{\mathrm{SP1}}}$). The source of this discrepancy is from velocity dependent effects present in the controller measurement, which are not present in the third order polynomial fit technique.

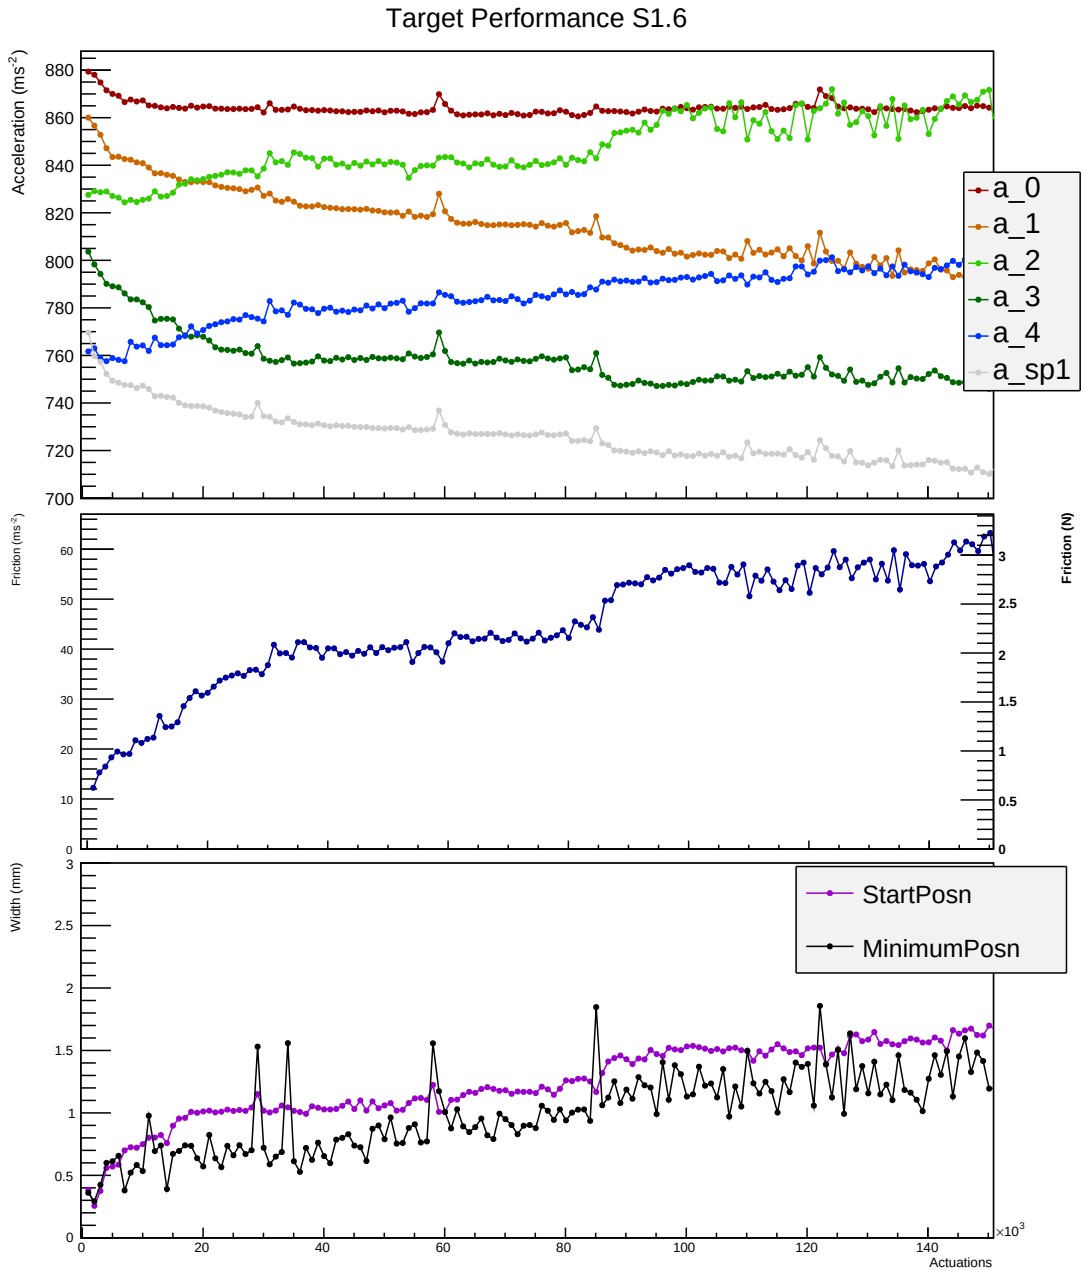

Figure 3.21: The acceleration magnitude, friction and start/minimum position width for target S1.6 during the first 150k actuations. The plotted value for both acceleration and friction is the median of 1000 samples recorded. The width is also measured over this range. Errors are not shown, but are dominated by systematic effects ($^{+2}_{-5}$ m s^{-2})

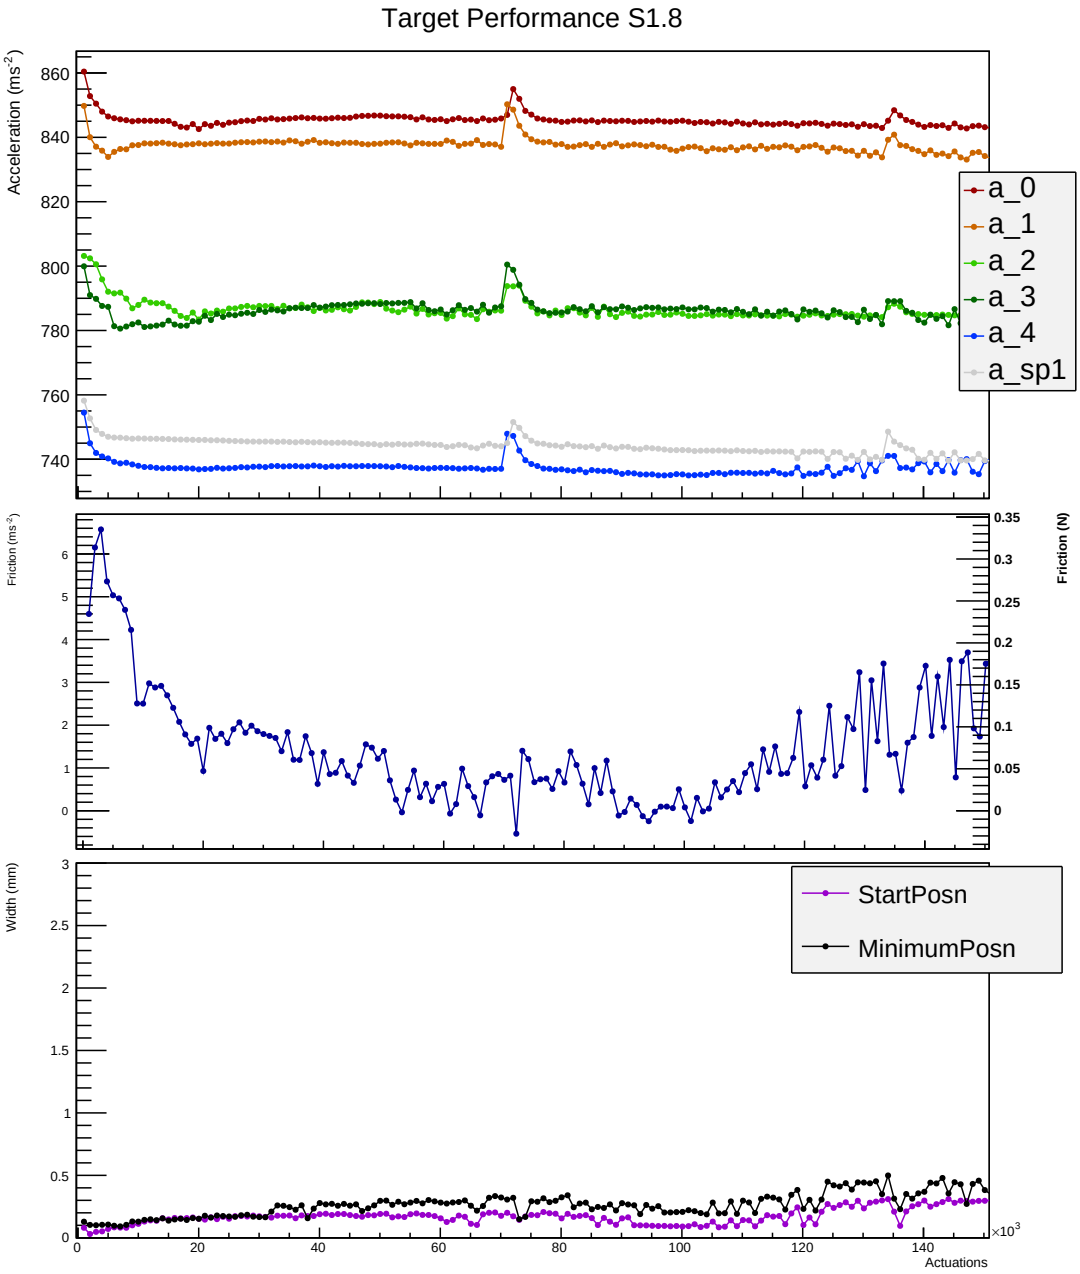

Figure 3.22: The acceleration magnitude, friction and start/minimum position width for target S1.8 during the first 150k actuations. The plotted value for both acceleration and is friction the median of 1000 samples recorded. The width is also measured over this range. Errors are not shown, but are dominated by systematic effects ($^{+2}_{-5}$ m s^{-2})

	Fail Parameters			Ideal Parameters		
Range	A_0	V_0	V_1	A_0	V_0	V_1
Friction	3×10^5	0	$75\,\mathrm{m\,s^{-2}}$	1×10^6	0	$40\,\mathrm{m\,s^{-2}}$
Acceleration Change	3×10^5	0	$-75\,\mathrm{m\,s^{-2}}$	1×10^6	0	$-40\,\mathrm{m\,s^{-2}}$
Start Width	3×10^5	0	3 mm	1×10^6	0	2 mm
Start Width (LR)	3×10^5	0.3 mm	3 mm	1×10^6	0.3 mm	2 mm

Table 3.5: Table showing the parameters which are used to determine the limits for failure and ideal boundaries. The LR case is adopted for reduced start position resolution from the controller data.

3.5.2 Quality Assurance for Running

A critical part of a reliable mechanism is the Quality Assurance (QA) process, which must validate each target before installation on ISIS. To that end the three monitoring techniques described in this chapter are used: friction (S1.4+), change in acceleration to SP1 and starting position width. From the test programme, it has also been possible to identify regions which deliver an acceptable performance as the target ages. These regions are arbitrary, however seem a good visual approximation to the wearing of the target mechanism. For completeness, the boundaries of each region are estimated by:

$$V(A) = V_0 + V_1 \exp(A/A_0) \tag{3.20}$$

where A is the number of actuations completed. V_0, V_1, A_0 are parameters, the values of which are specified in table 3.5.

Figure 3.23 shows the measured friction, acceleration change and starting position width over both the initial 50 thousand actuations and a longer period of five million actuations. The initial running is used in the QA process.

Both the S1.4 and S1.8 tests performed well, with S1.4 reaching a total of ten million actuations, while the S1.8 test concluded successfully after nine million actuations. The sudden decrease in friction of S1.4 coincided with a partial inspection of the mechanism, which did not disassemble the bearing assembly. The S1.5 and S1.6 tests showed a rapid degradation in performance and were rejected as production targets by the QA procedure. The S1.7 test was comparable to the S1.4 test in many ways, however it was terminated after 1.5 million actuations due to issues during the capture step of the actuation. The post-mortem of the S1.7 run showed a decreased temperature during levitation, which implies that the levitation current

was unintentionally reduced. This would increase the width of the finishing position, and cause issues during capture, which is compatible with the results seen in testing.

Figures 3.24, 3.25 and 3.26 show the process repeated for the comparable S1 tests, T1 tests and T2 tests respectively. Only controller data was available for this analysis, which prevented the friction measurements and also reduced the resolution of the start position measurement. The start position width regions have been adjusted to account for the decreased resolution.

The results of the QA tests are detailed in table 3.6, which indicates the result of each test for each target and the overall suitability as a production target. On the whole both friction and start position width are very clear discriminators of a good target. The acceleration change is influenced by many effects which ultimately reduces the clarity of the test result.

A total of 12 targets were studied, which contained two failures that did not reach 1 million actuations. Both the S1 and T1 stators have demonstrated that accepted target tests operated for in excess of 1.8 million actuations. In addition a number of these tests have demonstrated in excess of 5 million actuations. Note that testing to in excess of 5 million actuations has only been adopted recently in the test programme due to the time (75 days) required for the test.

3.6 Conclusion

This chapter has described in detail the numerous components of the MICE target mechanism and how it operates. From these descriptions a number of diagnostic measurements have been developed. These diagnostics use information directly from the control system to monitor the acceleration and friction using the starting position width. In addition the digitised position waveform has been subjected to an involved fitting procedure in order to study stator and mechanical properties. These diagnostics have been studied with a specifically built numerical model of the target and on the whole show fairly good agreement, although some systematic offsets are present. The friction measurement, which is of particular interest to monitoring the performance over extended periods, has demonstrated a resolution of 7 $\mathrm{m\,s^{-2}}$, or $\approx 1\%$ of the total acceleration.

Finally, these diagnostic techniques have then been applied to the QA process in R78, where each bearing test has been analysed. A total of two QA failures have been observed, and the accepted test targets each have delivered in excess of 1.5 million actuations, while some targets have demonstrated 10 million actuation lifetimes.

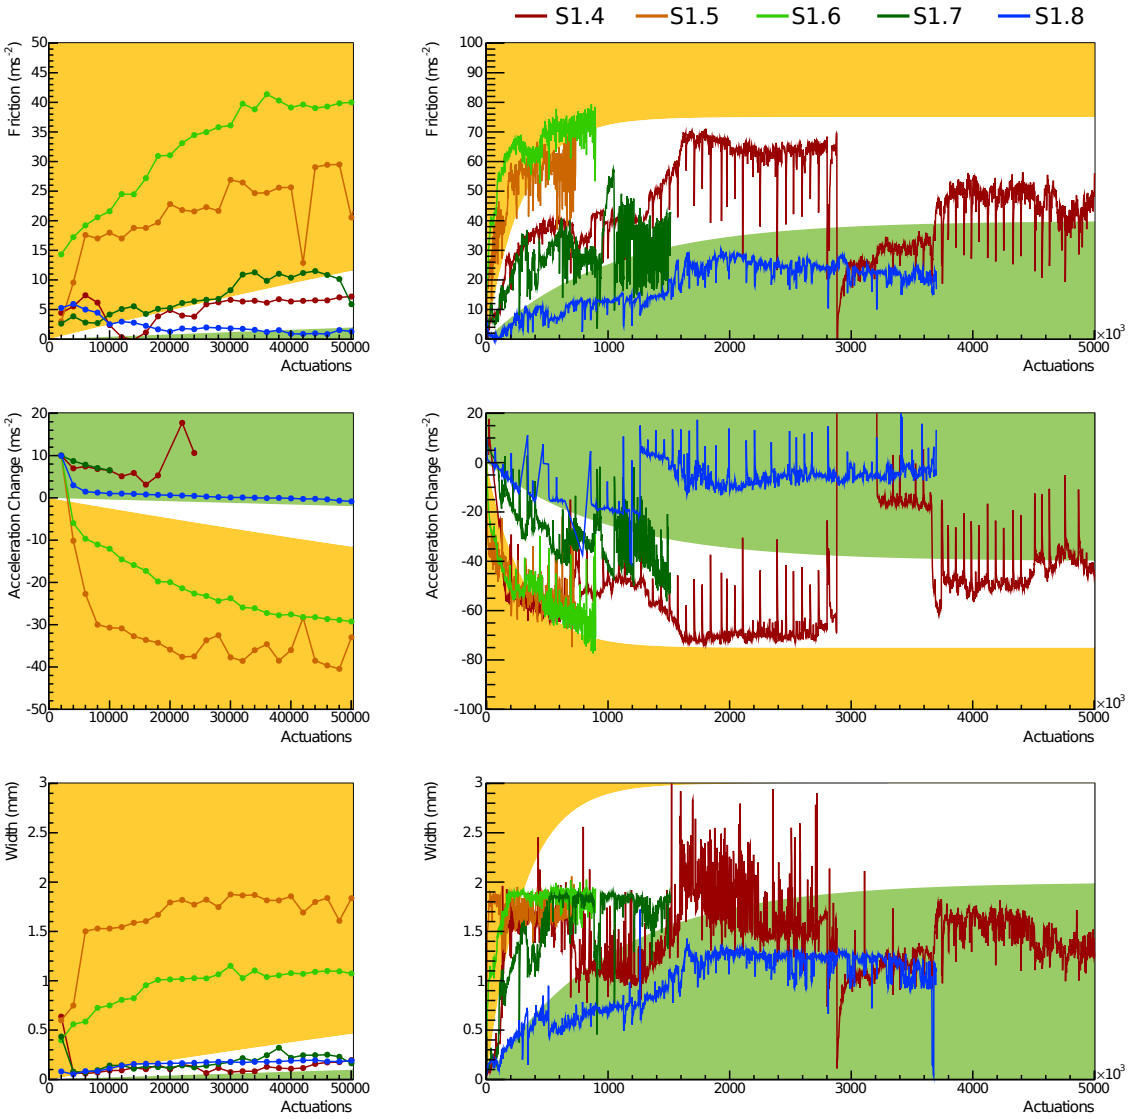

Figure 3.23: QA (left) and normal running of the S1 stator(right). The friction, acceleration change and start position width are shown. The yellow band indicates the unacceptable performance region and the green band identifies the ideal region. Missing acceleration change data points are due to a cut, which enabled only a single start position to measure the acceleration.

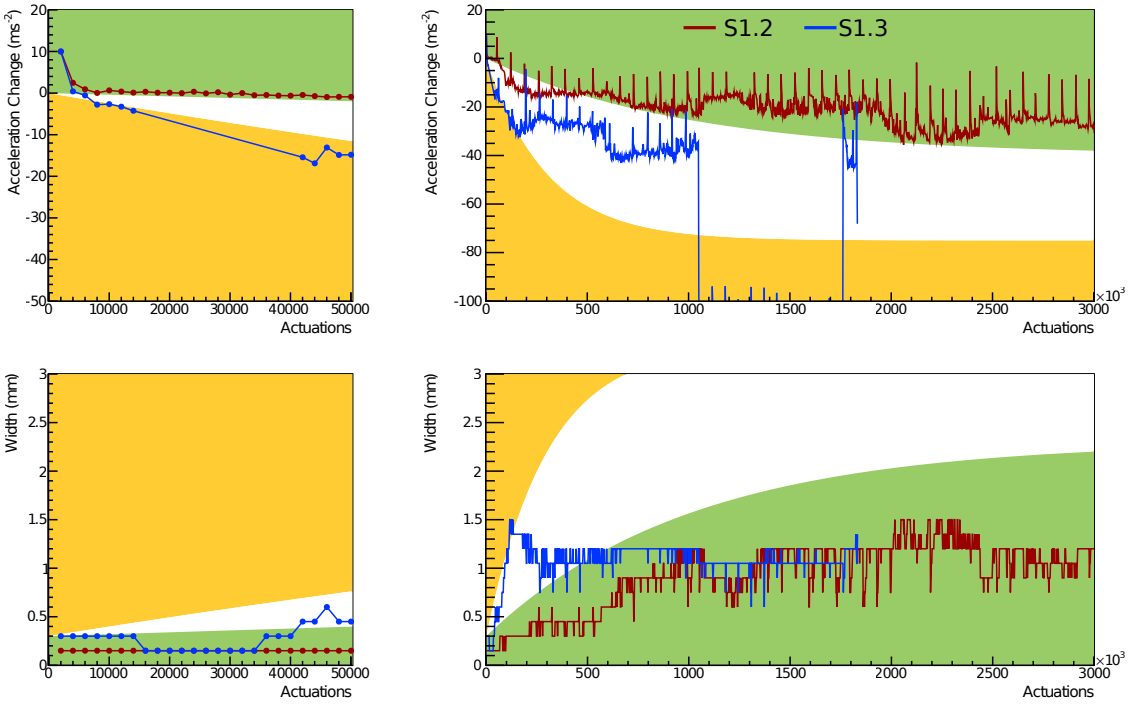

Figure 3.24: QA (left) and normal running of the S1 stator(right). The acceleration change and start position width are shown. The yellow band indicates the unacceptable performance region and the green band identifies the ideal region.

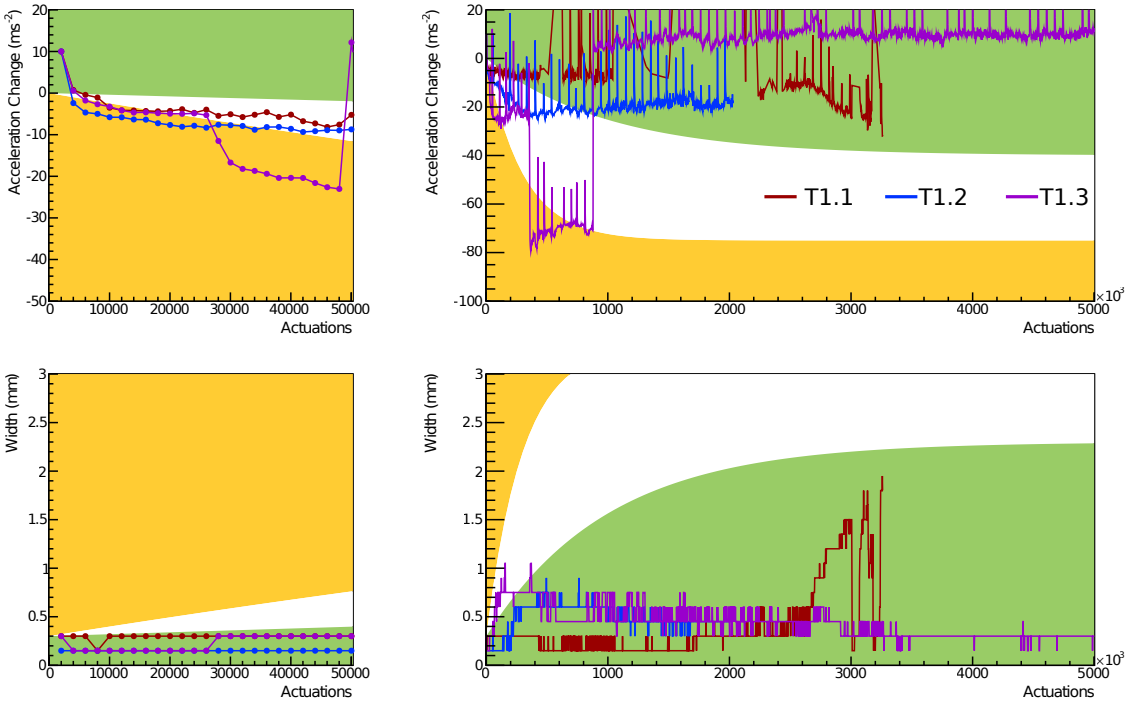

Figure 3.25: QA (left) and normal running of the T1 stator(right). The acceleration change and start position width are shown. The yellow band indicates the unacceptable performance region and the green band identifies the ideal region.

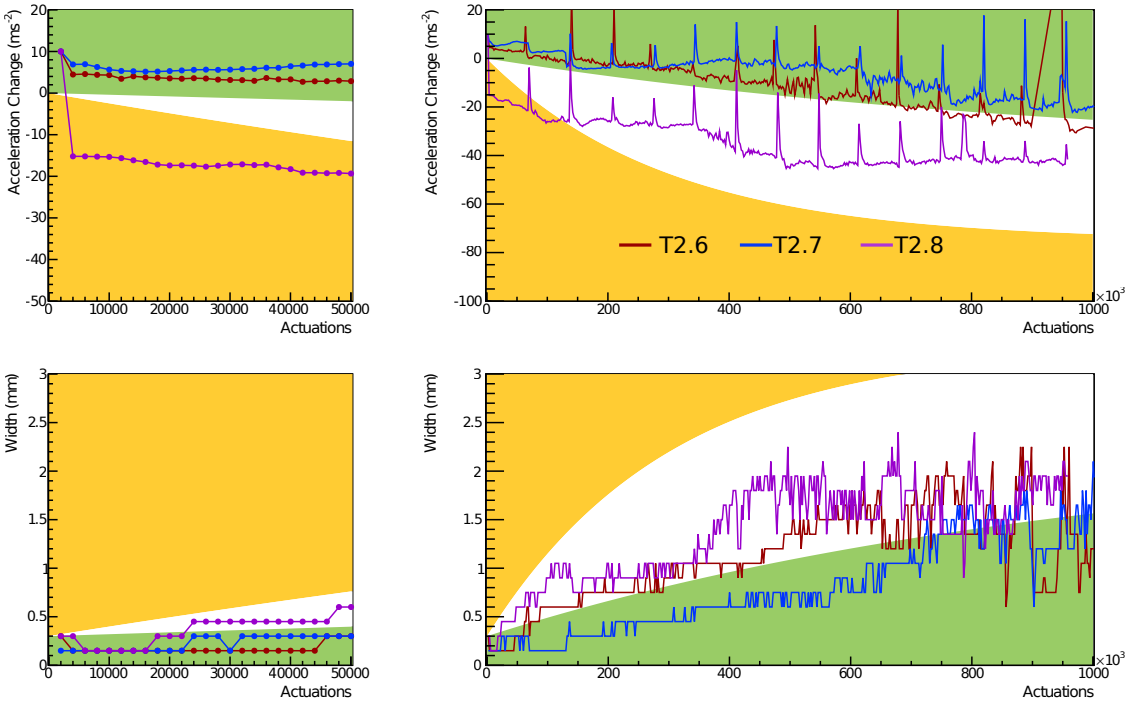

Figure 3.26: QA (left) and normal running of the T2 stator(right). The acceleration change and start position width are shown. The yellow band indicates the unacceptable performance region and the green band identifies the ideal region.

	50K tests				Number of
Target	Friction	Acceleration	Width	Accepted	Actuations (million)
S1.8	P	P	P	Y	8.5+
S1.7	p	-	P	Y	1.5
S1.6	F	F	F	N	0.9
S1.5	F	F	F	N	0.7
S1.4	P	-	P	Y	10
S1.3	-	f	P	Y	1.8
S1.2	-	P	P	Y	3.1
T1.3	-	f	P	Y	5
T1.2	-	P	P	Y	2
T1.1	-	P	P	Y	3.2
T2.8	-	F	P	Y	2.5
T2.7	-	P	P	Y	1.3
T2.6	-	P	P	Y	1.1

Table 3.6: Summary of stator QA results, in date-descending order. P(p) = Pass(borderline). F(f) = Fail(borderline). 1 million actuations corresponds to 14.8 days of continuous running at 50/64Hz (MS/64).

This page is intentionally left blank.

Chapter 4

The MICE Spectrometers for Step IV

The emittance of the MICE muon beam will be measured using a pair of spectrometer solenoids and fibre trackers, mounted before and after the cooling channel. As particles enter each solenoid a large B_z field causes the particle to follow a helix, where the radius is proportional to transverse momentum. Inside each solenoid a high granularity tracking detector is installed to record the path of the particle, allowing a measurement of x, y, z, px, py, pz. Measurements from each particle can then be combined in the covariance matrix, to allow a determination of emittance. Furthermore each tracker has a very high space point finding efficiency, which aids in accurate measurement of particle transmission though the cooling channel.

4.1 Spectrometer Design

The MICE spectrometers can be divided into four distinct components: a spectrometer magnet that generates a uniform B_z field in the tracking region, a scintillating fibre tracker which measures the position of particles passing through the solenoid, the supporting electronics needed to read out a scintillating fibre tracker, finally the software needed to reconstruct the particle tracks. These components are described in detailed in the following sections.

4.1.1 Spectrometer Solenoid

The spectrometer solenoid used in MICE is designed to provide a 4 T solenoidal field along the beamline, within which a tracking detector is to be installed. The field

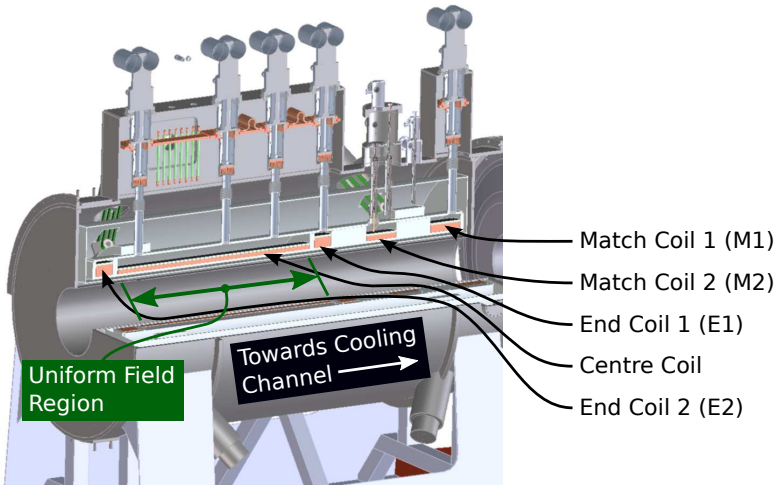

Figure 4.1: A cross section of the spectrometer mounted before the cooling channel, showing the position of the five internal coils and the uniform field region, within which the tracking detector will be installed.

within this region is specified to have a variation of less than 0.1% and is generated using a long centre coil surrounded by two smaller end coils as shown in figure 4.1. A further two coils are placed between the end coil (E1) and cooling channel to allow the beamline to be matched.

Each coil within the solenoid is wound from Nb-Ti superconductor, which has a copper to superconductor ratio of four[63]. The original design of the magnet specified the use of three pulse tube cryocoolers providing a total of 4.5 W of cooling at 4.2 K. Testing of the magnet revealed several problems and an additional two coolers were added to the design to increase the excess cooling at full load to 3.5 W[64].

4.1.2 Tracker Design and Construction

Within the 40 mm bore of the solenoid is a tracking detector. Particles are detected using 350 μm fibre that is doped with 1.25% by weight para-terphenyl (pT)[65] and emits scintillation light at a wavelength of 350 nm. A secondary dopant of 3-Hydroxyflavone (3HF) is added with a concentration of 0.25% by weight to shift the wavelength of the emitted light to 525 nm for better transmission along the fibres and a higher detection efficiency.

The fibres are moulded onto doublet layers as shown in figure 4.2a and the two layers are offset such that there are no gaps between fibres. The scintillating fibres are

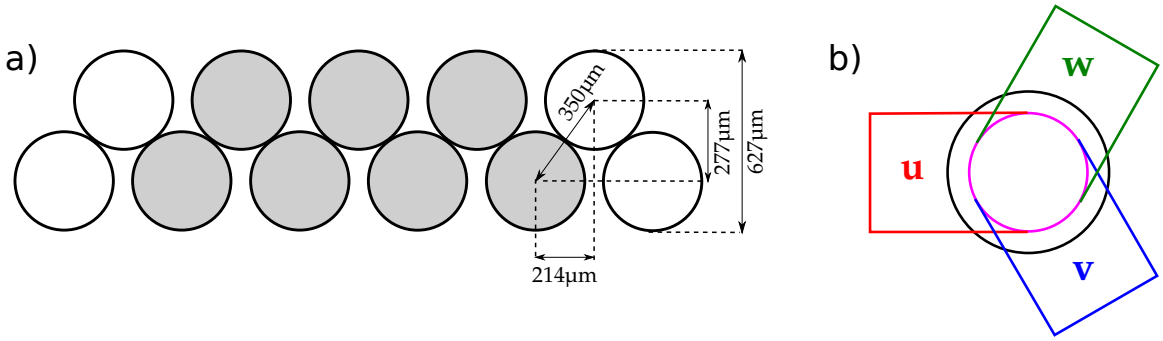

Figure 4.2: a) Drawing of the layout of the doublet fibre layer. b) the layer layout in the station.

grouped seven to a readout channel, to reduce the total number of fibres leaving the tracker and amount of readout required while still maintaining a thin layer to minimise scattering. One end of each fibre is connected to an undoped "clear" fibre lightguide, which connects to the photon conversion devices outside of the tracker (through a bulkhead connector). The remaining end of the scintillating fibre is converted to a mirror using aluminium vapour deposition, with mean reflectivity of 0.75[66]. Each doublet layer requires between 212 and 214 clear fibres to be readout.

The doublet layers are mounted on a carbon fibre station which has an inner diameter of 30 cm and is specially shaped to limit the minimum bend radius of the fibres. A total of three layers are mounted onto a carbon fibre station at 120 ° to each other, in the order W, U, V (see figure 4.2b). Between each layer a 25 μm thick layer of Mylar was added.

There are five stations mounted inside a single tracker, with separation starting at 20 cm and incrementing by 5 cm to produce a tracking volume 110 cm long. The stations are held in place with six support rods which run along the length of the tracker. To add mechanical strength to the supportive frame four cylindrical supports were installed between the stations. Figure 4.3 shows an image of a constructed tracker module. Light from each station is brought to a patch panel at the end of the tracker furthest from the cooling channel using clear fibre lightguides.

4.1.3 Readout and Data Acquisition

To digitise the optical signals from the clear fibres, MICE has opted to re-use the already built electronics used in the central fibre tracker of the DØ experiment. This decision removes the need to build a specific front end and has the advantage that

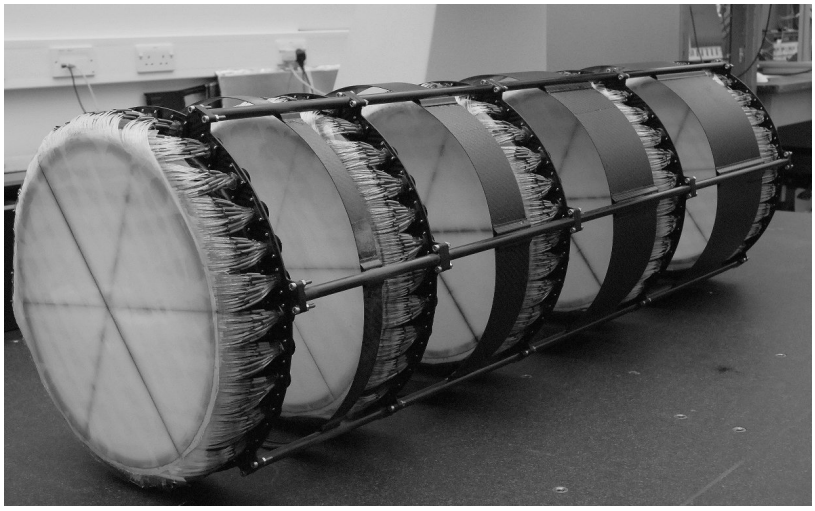

Figure 4.3: Image of a MICE fibre tracker after assembly

the front end electronics have been thoroughly tested. The downside is that the electronics have features specific to DØ which need to be accounted for.

Photon Conversion

The photon detection device used is the Visible Light Photon Counter (VLPC) and is a form of avalanche diode, which has a typical quantum efficiency of about 80%. The VLPC relies on the avalanche effect within a semiconductor diode and requires a bias voltage between 6 V and 8 V in order to operate effectively. Furthermore the devices must be operated at 9 K in order to provide a gain between 20,000 to 60,000[67] depending on manufacture conditions.

The VLPCs are grouped into cassettes, each containing a total of 128 VLPC modules with a similar gain. Each VLPC module contains 8 VLPCs, providing a total of 1024 channels per cassette. Each cassette also contains a number of temperature probes and heaters to accurately control the VLPCs temperature.

In MICE a total of four cryostats (two per tracker) each containing a pair of cassettes provide 8096 readout channels for the fibre trackers. The light from the patch panel is routed through external lightguides to the cryostat, the length of the internal and external lightguides combined being 4.0 m for all channels.

Front End Electronics

The VLPC chips are connected to a sophisticated circuit board, named the Analogue Front End IIt (AFE-IIt)[68]. The AFE-IIt board provides a variety of functions to support and read out the VLPC devices.

- **Control:** The AFE-IIt board supplies eight bias voltages, each one supplying 64 VLPCs with the same voltage. Selection of the correct bias voltage is critical, since a too high voltage will send some devices into breakdown and insufficient voltage will result in a loss of efficiency. A pair of AFE-IIt boards are connected to a single cassette to control the 1024 channels. One of the pair of boards is also responsible for adjusting the heater power to maintain the temperature of the VLPCs to 9.0±0.1 K.

- **Charge Collection:** Charge from the VLPCs is collected by an Application Specific Integrated Circuit (ASIC), called the TriP-t or Trigger and Pipeline with timing. An externally synthesised integration gate (or window) defines the time during which charge is collected and integrated. Any charge arriving outside this gate is ignored by the electronics, making it vital that this gate is synchronised to charge arrival. A discriminator compares the integrated charge with a programmable threshold, to provide a digital output. Once the discriminator is triggered a small capacitor also begins charging to provide timing information. At the end of the integration gate, both the charge and analogue timing information are stored in an analogue pipeline, which can be accessed later. In DØ the AFE-IIt boards were designed to generate an integration gate synchronised to a fixed frequency 53 MHz clock. In MICE, particles come from a target in the ISIS Synchrotron, and the bunches have a constantly increasing frequency due to the acceleration, making synchronisation to a fixed frequency clock impossible. To solve this problem the AFE-IIt run a modified firmware which allows an asynchronous externally generated integration gate to be used. The 53 MHz clock is still required by the boards and is distributed from a single source to each tracker.

- **Digitisation:** A digitisation (L1Accept) trigger is sent to the AFE-IIt board by encoding a bit onto the 53MHz clock distributed to the boards. After receiving this bit the boards wait until the closing of the integration gate, making it vital that the encoded trigger does not straddle this time. Once the gate is closed, the AFE-IIt board looks back a number of "buckets" (typically 4) in the analogue

pipeline and digitises the analogue values. To digitise all 512 channels on a board the process takes approximately 6μs. During this time the AFE is dead and unable to acquire any more data. In MICE this will limit the maximum muon rate; the reduction of this time has been investigated however not yet implemented. The digitised data is then packaged, check-summed and zero suppressed by Field Programmable Gate Arrays (FPGAs). The packaged data is sent across Low Voltage Differential Signalling links to a VME crate, where the data is buffered before being collected at the end of the MICE spill.

The control and configuration of the AFE-IIt is handled independently of the data collection, using a 1553 serial data bus. The bus connects each cryostat to a VME crate, which is then connected to a computer. Software running on the computer enables a complete configuration of the AFE-IIt, while relaying selected information to EPICS as process variables.

Calibration

For efficient operation of the detector the front end electronics of the AFE-IIt must be correctly calibrated. This is done in stages, starting with the optimisation of the photon detectors, followed by the electronics and finally the determination of constants required to convert from electronics counts to physical values.

For each group of 64 VLPCs a bias voltage must be found which balances a high quantum efficiency with a low noise rate without sending channels into saturation. A value is found by sweeping a number of different voltages, without an injected signal. The optimum bias voltage is selected where the probability of a single photo-electron from dark noise arriving inside the charge integration gate is close to 2% for the VLPCs[69].

Once the bias voltages are fixed, it is then possible to find the conversion constants required to compute the number of photo-electrons from the raw data collected. The constants are found by injecting light into the VLPCs and histogramming the integrated charge as shown in figure 4.4. By identifying the pedestal peak and following peaks it is then possible to find the pedestal and gain of each counter.

A similar procedure is also required for selecting the optimum discriminator and TDC configuration, as well as the appropriate software conversions.

In MICE the AFE-IIt readout boards will be located inside the MICE experimental hall, where there is substantial temperature variation, large stray magnetic

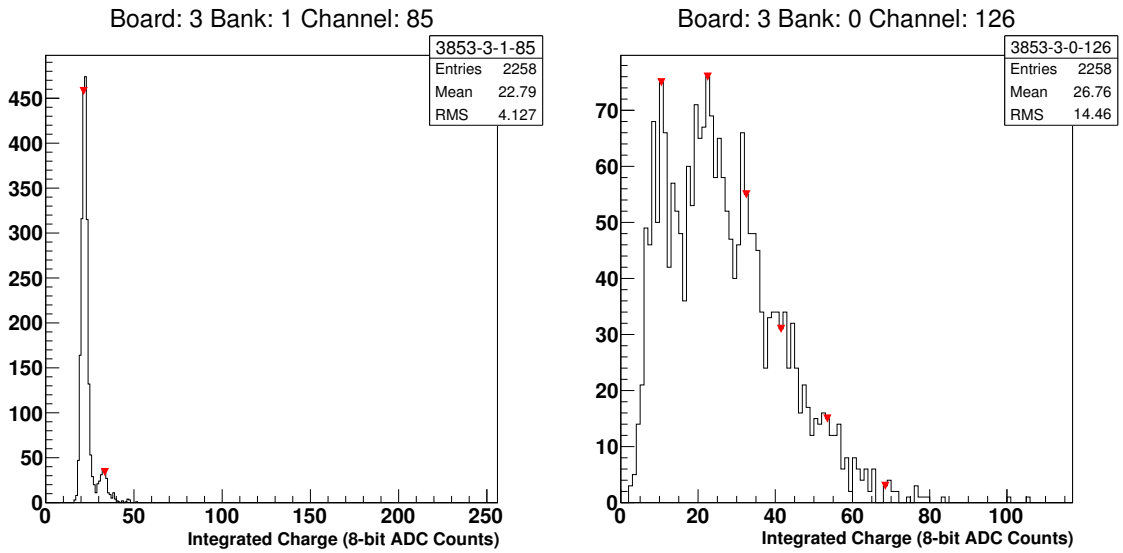

Figure 4.4: The integrated charge of a calibrated VLPC. Left: without light injection, Right: with light injected from an LED source, showing photon peaks. Note these images were taken during the same run, where the LED pulser was attached to a single cassette.

fields and high power RF. For these reasons it is anticipated that the readout boards will require recalibration on a regular basis to achieve optimum performance.

4.1.4 Reconstruction

The reconstruction process converts the raw detector data into particle tracks containing the track and measured momentum in a number of stages. First the integrated light for each channel is converted into a photo-electron count using a calibration. Next a cabling file is used to map the front end channels to a tracker, station, plane and fibre number. A digit is then made from each fibre which produced more than 2 photo-electrons in the front end electronics. Next neighbouring digits are combined into clusters to rejoin particles which were split across two fibres. Triplet space points are built by identifying the intersection of three clusters on separate fibre planes. Remaining (unmatched) digits are then also converted into duplet space-points. The computed space-points for each station are then passed to a pattern recognition routine, which identifies the most likely straight and helix tracks from the points. The final stage in the reconstruction is to pass the recognised tracks to a Kalman filter which provides a measurement of the particle's position and momentum, taking into account multiple scattering and misalignment of each station.

4.2 Synchronisation to ISIS 1RF

As discussed in the previous section, the charge collected from the VLPCs must be integrated by the TriP-t within a integration window. Step I of MICE had no concept of a bunched muon beam, although the underlying time structure is present. Inclusion of the tracker for step IV+ will require a system to synchronise the MICE trackers to the bunched muon beam.

As MICE is a single particle experiment, most "bunches" are empty. It is therefore necessary to use diagnostics from the ISIS accelerator to define when the bunches will arrive. The chosen signal is the reference signal from the ISIS RF cavities (ISIS-1RF). The choice to use this signal is that it is a well behaved sine wave which is phase locked to the protons inside ISIS and it is always available (even when ISIS is in shutdown), for the use of other diagnostic equipment.

This section describes the main features of the ISIS-1RF and how it can be integrated with the rest of MICE.

4.2.1 RF Frequency and Particle Correlation

The first task was to understand the behaviour of the ISIS RF and verify it is well correlated with the particles seen in the MICE detectors. For simplicity the necessary data was collected using a Tektronix DPO7000 oscilloscope, which had sufficient internal memory to digitise 4 channels at 200 MS/s for 20 ms or 500 MS/s for 5 ms. This allowed an entire ISIS spill to be recorded, including the following signals:

- **ISIS-1RF Reference** The reference signal provided from the ISIS Main Control Room, patched to the MLCR using a spare target diagnostic link.
- **TOF0 Hit** A coincidence between two PMTs fixed on a single slab of TOF0
- **Particle Trigger from TOF1** A particle trigger generated from TOF1, which adds an additional 50 ns delay compared to a direct coincidence as in TOF0.
- **Spill Gate** The spill gate which limits the time which the MICE detectors record data.

The raw signal data collected by the oscilloscope was stored to csv files, to allow a ROOT script to perform a more sophisticated analysis later. The calculation of frequency and time of each period was accomplished by fitting a sine function to the raw data.

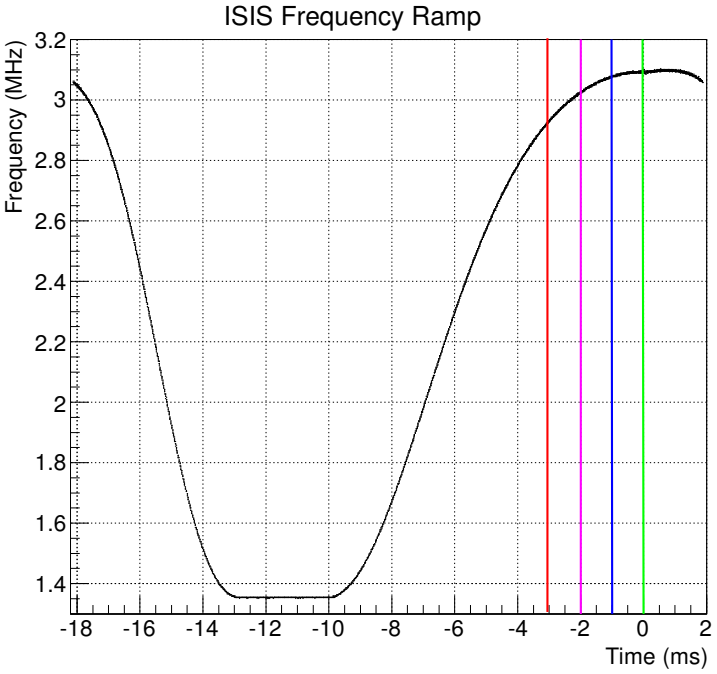

Figure 4.5: The frequency of the ISIS-1RF across the entire ISIS spill. The oscilloscope was triggered at the end of the spill, such that injection occurs at -10 ms and extraction is at 0 ms. The red, purple and blue lines indicate the opening of 3 ms, 2 ms and 1 ms spill gates, while the green line indicates the closing of the spill gate.

Spill Gate Width (ms)	Δf (kHz)	Δt (ns)
3.03	168	18.5
2.00	70	7.48
1.00	17	1.78

Table 4.1: Table of frequency and period change over data collection for typical spill gate widths. The opening and closing times of the spill gate are indicated by the coloured lines in figure 4.5.

The frequency of the ISIS-1RF during the ISIS spill is shown in figure 4.5. At injection the RF frequency begins at 1.37 MHz and is increased to a maximum of 3.09 MHz at the end of the acceleration cycle. Within MICE readout triggers are restricted to a specific time within the ISIS spill using a spill gate. In step I this gate was located during the final 3 ms of the ISIS spill and this is likely to remain unchanged for step IV. At step VI the inclusion of the high power RF within MICE will restrict the active time of the experiment to 1 ms. Table 4.1 gives the frequency and period variation for different possible spill gates. Note that a 3 ms spill gate introduces almost a 20 ns time shift across the spill gate.

The correlation of the particle triggers was accomplished by time stamping both the NIM signals from the particle triggers and the zero crossing rising edges of the ISIS-RF. Once both signals had been stamped it was then trivial to measure the delay between each trigger and the previous RF edge. During data taking the decay solenoid was unavailable and this reduced the overall particle rate and caused a noticeable number of noise triggers. To reduce these a cut was applied to the data to require a hit in TOF0 50 ns to 300 ns before the particle trigger request was seen from TOF1.

The measured delays for a range of muon beams are shown in figure 4.6, which shows triggers arriving around 350 ns. To prevent the distribution being split, the delay axis range was set between 200 ns to 530 ns, delays which fell outside of this had the RF period added/subtracted as necessary.

Both muon and pion beams showed a strong correlation with the ISIS-1RF, with the majority of triggers being received within 280 and 420ns, demonstrating that the ISIS-1RF signal can be used to synchronise the MICE trackers to particle arrival.

4.2.2 Trigger Veto

Triggers from the TOF system can arrive at any time, which can occur outside of the integration gate for the tracker front end. This will introduce a subtle loss of particles within the tracker, which introduces further complications to measurements of the cooling process. To prevent this it is necessary to apply a veto to the particle triggers to only allow readout of particles which were integrated completely by the front end electronics. This veto is named the "alive window", as this is the time the electronics are able to receive triggers.

To find the fraction of triggers which could be accepted by an alive window equation 4.1 was used to test each delay (t_d) measured previously and histogrammed

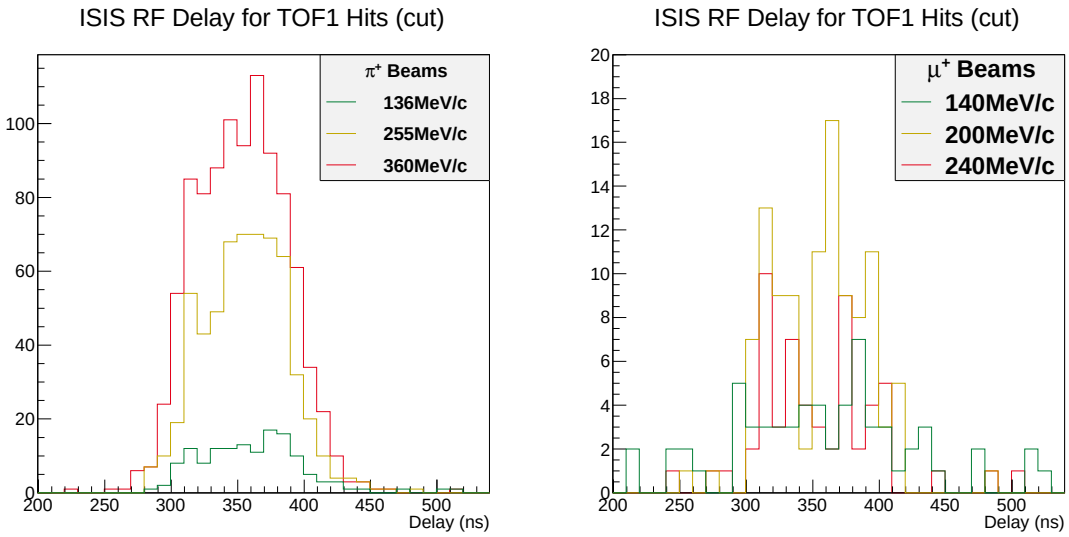

Figure 4.6: The delay between RF edge and particle trigger arrival for a range of pion and muon beams. Beam momentum defined at D2.

in figure 4.6. If a delay met the criteria in the equation it was accepted. To study the effects of different alive window widths, a_w was swept and the window centre, a_c was adjusted to maximise the number of accepted particles. In reality it is unlikely to perfectly align the window, so an additional case where a_c was offset by 10ns from the optimal position was also calculated.

$$a_c - \frac{a_w}{2} < t_d < a_c + \frac{a_w}{2} \tag{4.1}$$

The results from the study for a 360 MeV/c pion beam are shown in figure 4.7. For an alive window with width less than 100 ns the fraction of triggers drops linearly, which is to be expected as bunches inside ISIS are around 100 ns long. Conversely when the width is greater than 120 ns there is little gain in the overall rate. For MICE an expected alive window width is 120 ns, which will result in at least 88% of coincident triggers being digitised even if there is a 10 ns window misalignment.

4.2.3 Window Generation

The next task was to design and build a system which would allow programmable synchronous gates to be generated from the ISIS-1RF. For the inclusion of both trackers a total of three windows are required (one integration gate for each tracker and an alive window). Furthermore it must be possible to generate each window at

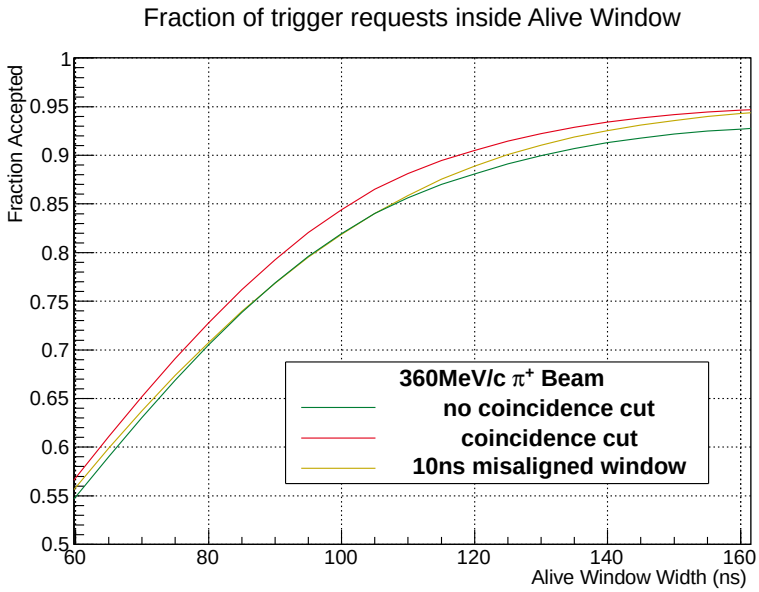

Figure 4.7: Fraction of accepted particles as a function of alive window width. The coincidence cut refers to the requirement of observing a hit in TOF0.

any point on the RF signal, since the precise cable lengths and timing are not known.

The adopted solution was to trigger a programmable delay directly from the ISIS-1RF signal, specifically a pair of Highland Technology T560 Delay Generators[70] were used. Each generator has four outputs, each one capable of generating a independent programmable duration pulse after a programmable delay for each channel. Once a generator is triggered it cannot be re-triggered until all outputs have completed timing, plus some additional dead time (60 ns). The latest possible time a window could be located is at 320 ns (the end of the ISIS RF period) and the maximum width for the integration window is 200 ns, which would need a total of 600 ns including the generators insertion delay and dead time. Hence, it is not possible to use a single unit to generate a window on each crossing. Instead a pair of delay generators were connected to the RF with the second gated to only trigger when enabled from an output on the first. The outputs from the two generators were then passed to a logical OR to combine the signals. Figure 4.8 shows a timing diagram for this application. This approach is convenient since the timing configuration is identical for both units and only a small change is required to the gate configuration.

This technique works well at the design frequency, however at frequencies below 2.5 MHz a timing subtlety causes only alternate windows to be generated. As this is located outside of the spill gate (see figure 4.5), no particle data will be collected at

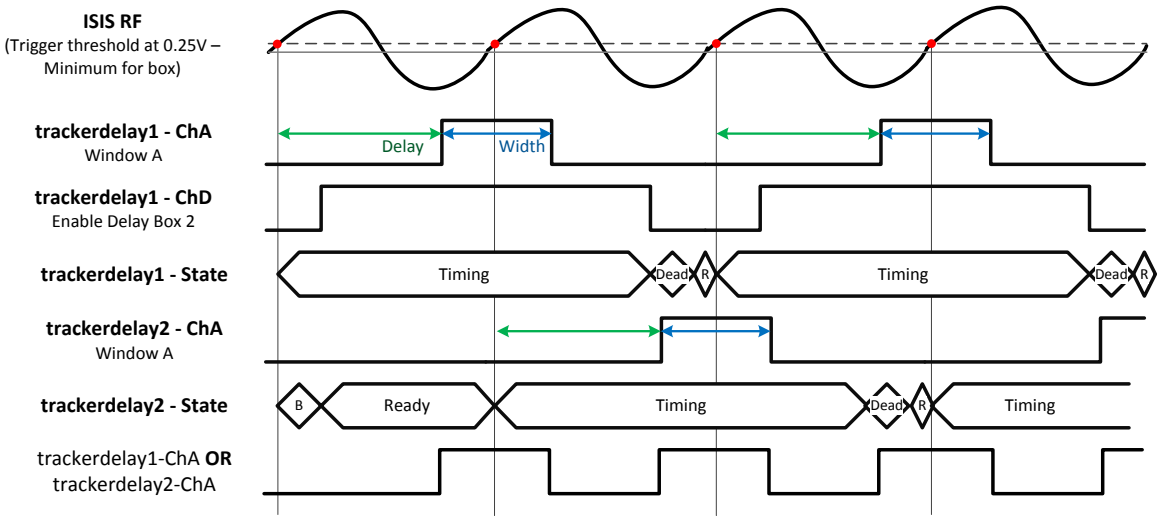

Figure 4.8: Timing diagram for the use of T560 Digital delay boxes to generate windows for charge integration and trigger gating from each crossing of the ISIS RF.

this time so it is not anticipated to adversely affect operation of the tracker.

The delay generators are triggered directly from the RF signal to ensure they are always synchronised, however this makes them more susceptible to noise introducing timing jitter. To minimise this effect the delay units are triggered close to the zero crossing, where dV/dt is largest (around 40 mV/ns). The amount of additional jitter introduced by this is expected to be small in comparison with the timing variation from the changing RF frequency.

4.2.4 Inclusion with MICE

The next step was to integrate the gate generation with the MICE DAQ in order to correctly signal the tracker readout electronics. Figure 4.9 shows the location and signalling of the electronics elements. The alive window is connected directly to the MICE trigger veto, which inhibits triggers to all detectors globally. A pair of independent integration windows are sent through a patch panel to the MICE hall, where each is fanned out to the necessary AFE-IIt boards.

The particle trigger signal is distributed to the trackers through a third programmable delay unit. This is used to add an additional degree of freedom to ensure that the triggers arriving at the AFE-IIt do not straddle the closing of the integration window.

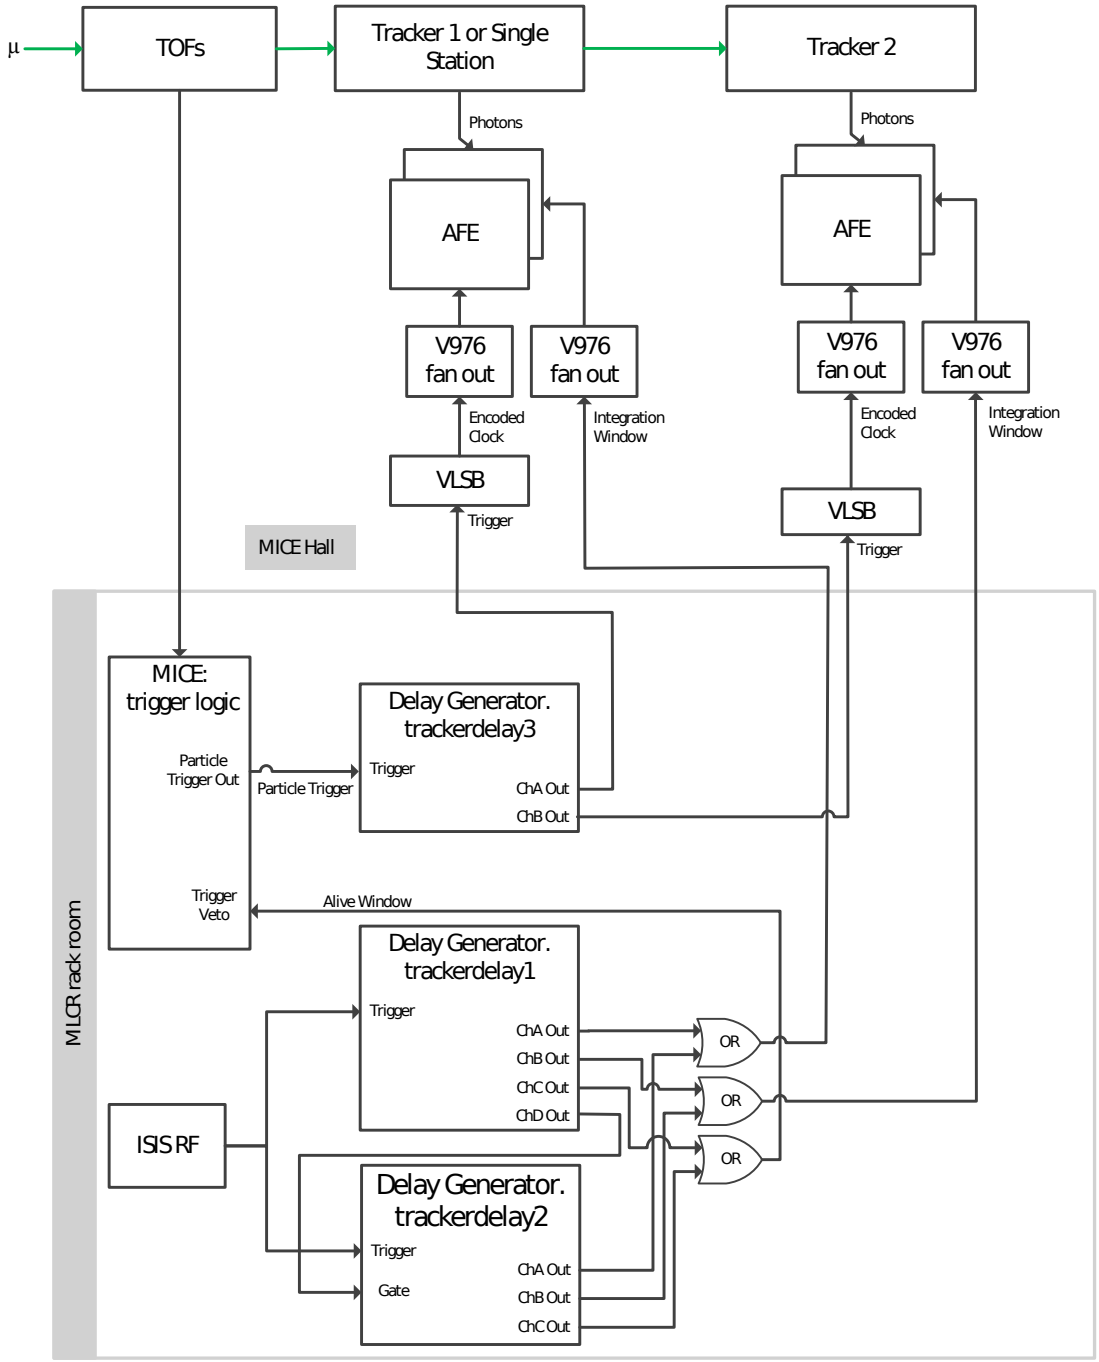

Figure 4.9: Figure of additional signal distribution and logic required for tracker operation.

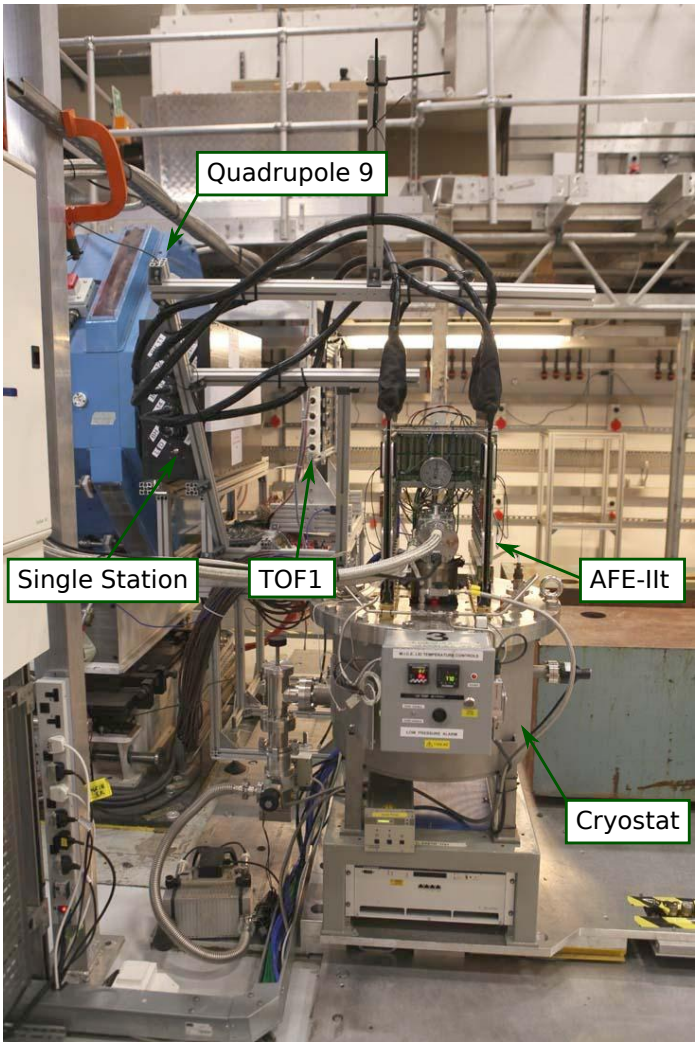

Figure 4.10: Image of the single station test

All delay generators used are programmable, since the timing of the trigger arrival can shift depending on if TOF0 or TOF1 is used to trigger the experiment. The T560 delay generators used in MICE have an Ethernet connection, which accepts text commands sent over a TCP socket. For initial testing the configuration was handled by a simple python script, however for step IV a full EPICS interface is being written. This will enable the configuration for the units to be stored in the MICE configuration database with other sensitive parameters.

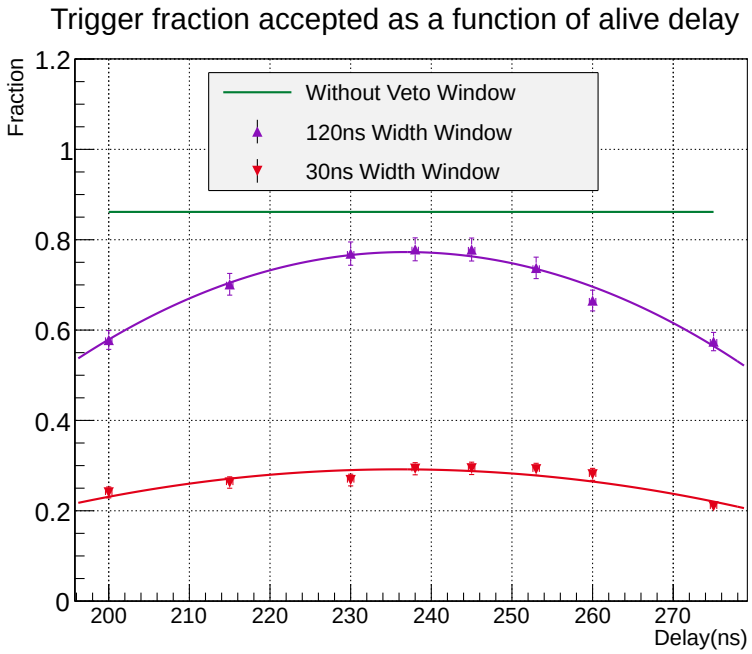

Figure 4.11: Ratio of accepted triggers requests, for different alive window delays.

4.2.5 Beam Test with a Single Station

The integration of the MICE tracker electronics was tested with beam in the MICE hall. For the test a single unused tracker station was mounted inside a light tight box and fitted between Q9 and TOF1 as shown in figure 4.10. The station contained a fraction of the total number of fibres, making it possible to use a single cryostat for readout.

The configuration of the timing circuits was completed in steps, the first of which was to ensure that the alive window was correctly aligned to particle triggers in the DAQ. This was checked by using a fixed window width and then sweeping the delay to determine the value which maximises the ratio of accepted triggers to trigger requests. Figure 4.11 shows the result of the timing scan, where an optimal location for the window occurred at 238 ns. The sweep used a high rate beam and consequently the dead time of the experiment limited the maximum number of triggers to around 86%, without any alive window in place. Once this is taken into account, the fraction of accepted triggers is 89%, and compares well with estimates from figure 4.7.

Once the alive window was fixed it was then necessary to adjust the trigger delay to ensure that triggers arriving at the AFE boards were not straddling the closing of

the integration window. To check this a diagnostic from the VLSB was used, which counted the number of clock cycles between encoding a trigger and the digitised data being returned. The variation of this time should be around the width of the alive window, however in the event of window straddling the variation of this time becomes much larger ($+ \approx$ ISIS-1RF period).

The final task was to align the integration window on the AFE-IIt to the time when charge arrived and select the correct point in the pipeline to read out. This optimum location occurs when the light yield and space-point finding efficiency are the highest.

The light yield of the station was determined by investigating each cluster forming a space-point, shown in figure 4.12. One of the VLPC cassettes used was of high gain and resulted in saturated channels around 10 photo-electrons, while the other cassette was significantly lower gain, requiring that saturation effects must be accounted for. The light yield was therefore measured using the same technique as for the cosmic data[66]. This technique rejects all saturated channels (where the 8-bit ADC value was 255) from the analysis, which is then compensated for by rescaling each bin by the probability of being rejected due to a saturated channel. This technique is possible because the low gain VLPCs are capable of measuring a larger number of photo-electrons.

The efficiency of finding a space-point was calculated by requiring a single particle to pass though both TOF0 and TOF1. If at least one space-point was detected in the single station this was counted as a hit, conversely if nothing was observed this counted as a miss. The efficiency was then determined as the fraction of hits to hits and misses.

A starting estimate for the location of the integration window was made by computing the trigger latency and cable delays. To maintain a fixed time between integration window and trigger arrival, both the trigger delay and integration window delay were always adjusted simultaneously. To optimise the location of the window a scan of the integration delay was then performed. Figure 4.13 shows the result of the scan and indicates an optimal location for the window to be around 275 ns. For the test a much less sophisticated analysis was adopted and a value of 250 ns was chosen. This reduced the overall light yield by around 3.6% and had a 95% efficiency for producing a space-point. A summary of the delay unit settings adopted by the single station run is presented in table 4.2.

The efficiency of the station was significantly less than observed during cosmic ray tests (99.8 ± 0.1 %). To verify that the integration window was aligned cor-

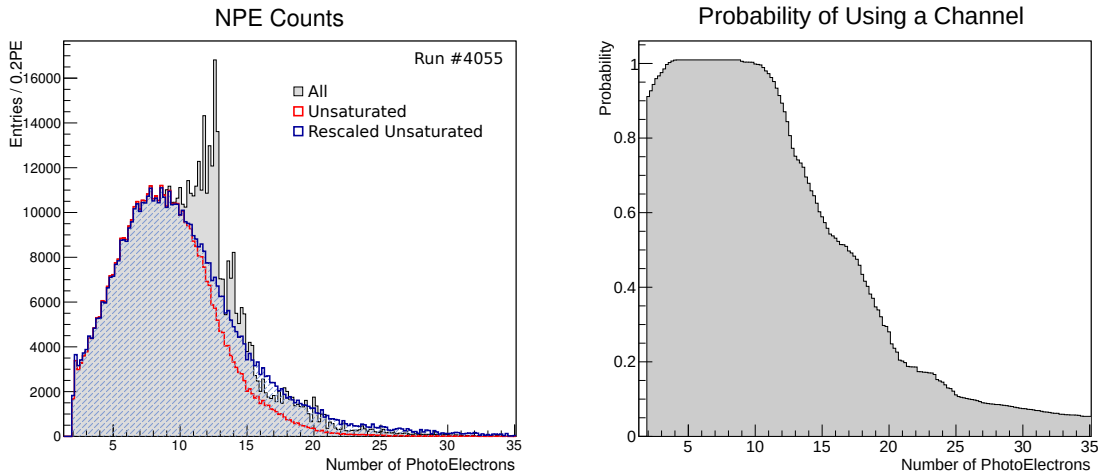

Figure 4.12: left: Number of photo-electrons collected by the tracker front end electronics, showing the raw data, unsaturated channels, and rescaled by dividing by the probability of a light yield being used. right: Probability of a light yield being included in the analysis.

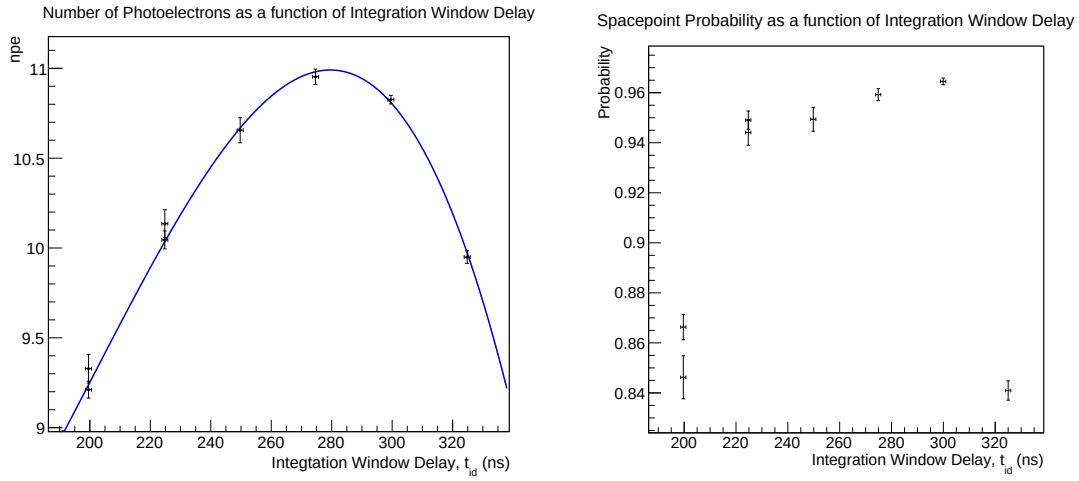

Figure 4.13: left: Mean number of photo-electrons collected by the tracker front end electronics, for different integration gate delays (t_{id}). Fitted with a generic 3rd order polynomial. right: Probability of finding a space-point in the station as a function of integration gate delay.

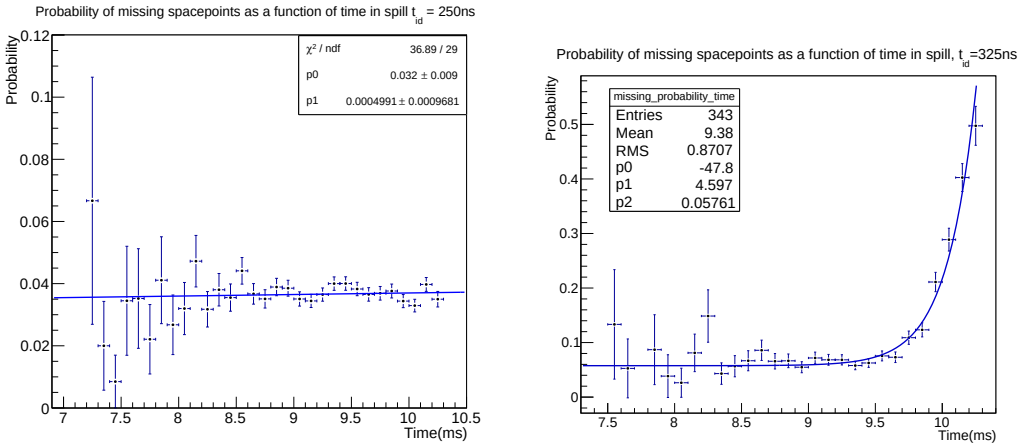

Figure 4.14: Probability of not observing a space-point in the spill gate. left: correctly configured integration gate. right: late integration window causing a loss in space-points at the end of a spill.

Parameter	Value
Alive Window Width	110ns
Alive Window Delay	238ns
Integration Window Width	200ns
Integration Window Delay	250ns
Trigger Delay	320ns
Pipeline Depth	4

Table 4.2: Table of the parameters used in the single station test.

rectly during the entire 3 ms spill gate, the probability of missing a space-point was investigated as a function of time within the spill. Figure 4.14 shows the missing space-points for a run where the integration delay was too large and caused a loss of space-points at the end of the spill. For the chosen integration delay the fraction of missing space-points was constant across the spill, indicating that the window was correctly aligned.

To verify that the particles were passing through the station the probability of finding a space-point was found for each pixel within the TOF1 detector. The result of the computation is shown in figure 4.15, where the core of the beam shows a considerably higher efficiency than that seen around the edges. The space-point finding efficiency using only the centre 9 pixels of the TOFs is 99.4 ± 0.4%, which compares well with the cosmic ray result. This indicates the reduction in efficiency

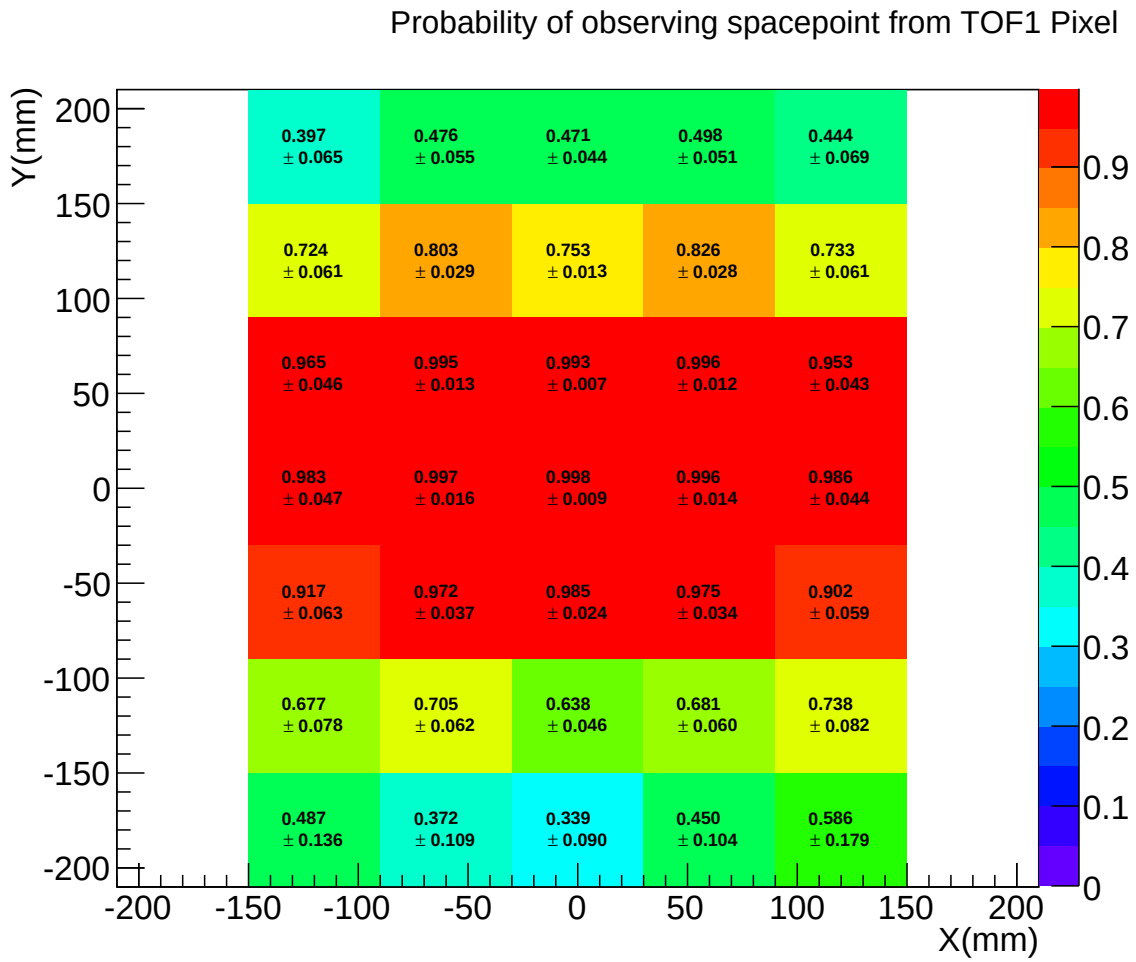

Figure 4.15: Space-point finding efficiency for each pixel in TOF1

was primarily caused by particles passing through the TOFs but missing the single tracker station.

4.2.6 Tracker Dead Time

The inclusion of the tracker to the MICE DAQ adds two sources of dead time into the experiment. The first arises from the alive window, which reduces the number of particles by around 10%. The second comes from the 6 μs digitisation time of the front end electronics and becomes notable at high rates. To study the dead time it was necessary to investigate the raw data from the scalars in the trigger logic, as these count independently and contain a negligible amount of dead time. Figure 4.16 shows the number of accepted triggers as a function of the total number of trigger requests, for a very high rate pion beam. The data was fitted to the following equation:

$$y = \frac{\alpha x}{1 + \beta x} \tag{4.2}$$

where α is the probability of a trigger being accepted due to the alive window. β is the dead time as a fraction of total data collection time. The beamloss and therefore instantaneous trigger rate is not constant during data collection, which overestimates β. Note that a typical muon beam is limited to only 40 particles per spill, allowing in excess of 85% of particles to be digitised. If the rate is increased to more than 100 particles per spill, the fraction which can be digitised is less than 65%.

This work has allowed a realistic estimate of the possible data collection rate the MICE instrumentation is capable of. This has allowed MICE to construct a run plan which finds a compromise between maximising the number of particles in each measurement and the number of measurement points which can be obtained. Finally, MICE are also considering mechanisms to decrease this dead time, specifically by running the tracker with discriminators only.

4.2.7 Improvements from the Single Station Test

The single station test provided a thorough test of the electronics needed to operate the fibre trackers inside the MICE hall. During the test a number of improvements were suggested, which include:

- **Independent Reference Signal** The ability to use an independent copy of the ISIS RF, which does not change the specific cabling (which in turn changes timing), would greatly assist diagnostics during setup.

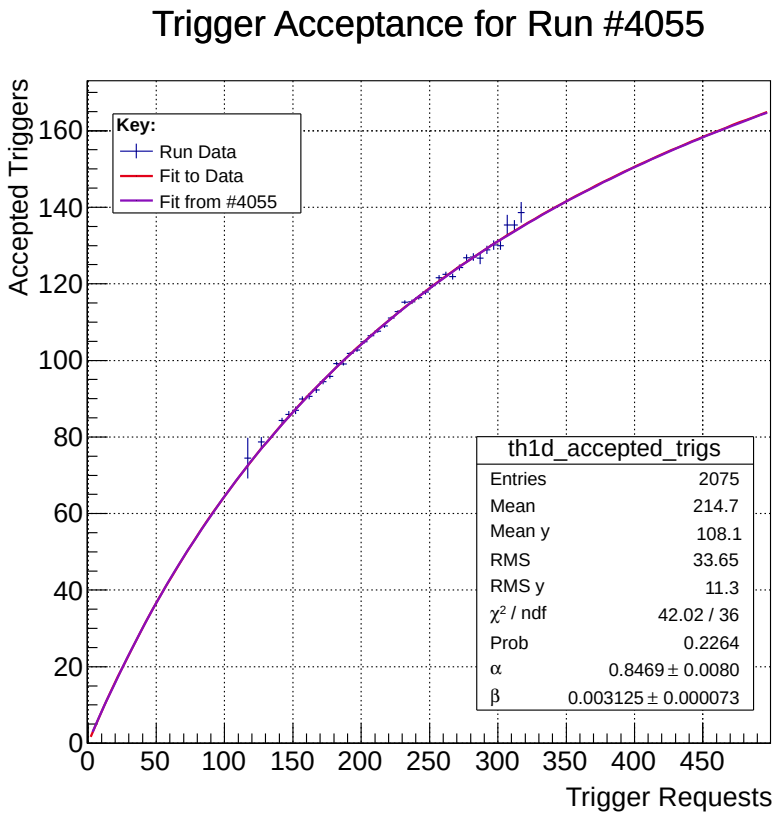

Figure 4.16: Effect of dead time on accepted particle rate.

- **Distribution of the RF** The fan-out of the ISIS 1RF to the delay units and oscilloscope distorted the incoming signal substantially. This would ideally be avoided.
- **Delay Unit Gating** The gating used on the delay boxes had a subtle glitch, which it would be nice to avoid in future.

The solution to these issues is to correctly buffer the signal inside the local control room and provide a gate independently of the delay units. These requirements were fulfilled with the construction of a buffer module. The module provided four buffered copies of the ISIS 1RF and an additional gate output which inverted with each positive crossing of the RF signal.

4.3 Light Injection System

A complete calibration requires the ability to inject light onto the VLPCs. For the cosmic ray tests of the tracker, the external lightguides were removed and replaced with a LED pulser fixture. This procedure adds significant time to the calibration process and is not without risk to both the lightguides and fibre map. For operation in the MICE hall it is therefore necessary to inject light into the trackers, with the external lightguides in-situ. The fast calibration procedure, which will validate the setting of the VLPC bias voltages and maintain the calibration constants required for converting to an integrated charge in photo-electrons.

In addition the ability to inject light into the trackers allows timing diagnostics to help identify the edges of the charge integration gate. This can then identify misalignment of the timing and trigger gates which must be corrected for efficient operation. Furthermore, injecting light through the entire collection and readout system allows a complete check of the detector.

Ultimately the LED system will perform a major component in validating that the MICE trackers are operating efficiently and allow that efficiency to be maintained, which is required for a high precision, unbiased transmission and emittance measurement.

The procedure for light injection is based on part of the fibre QA process, where a highly focused 405 nm LED was used to illuminate each fibre in a doublet layer[71]. This directly excites the 3HF wavelength shifter and causes the emission of 525 nm light. For the light injection system an unfocused light source will be used to illuminate an entire tracker station.

This section describes the design, manufacture and control of the LED system implemented in the fibre tracker.

4.3.1 Light Injection Requirements

The light injection system must conform to the following requirements to be suitable for installation in the trackers:

- **Light Uniformity:** The light intensity across each station should be as uniform as possible to reduce the number of different intensity runs needed for calibration of the VLPCs. It is worth noting that each fibre is of different length, causing variation in light level even for a perfectly uniform illumination. Furthermore the light uniformity is not required to be precisely known since the photon counters have by definition the ability to resolve single photons.
- **Stability:** The experiment has a running schedule up to 2020. In order to verify the trackers are not degrading, the light injection system must also not degrade over this period.
- **Reliability:** Once the spectrometers are fixed within MICE, removal will be unlikely, therefore components of the system installed within the solenoid must continue to function for the duration of the experiment.
- **Installation Risk:** The MICE trackers are built and there is a limited number of spare parts. The complexity and installation of the system should be designed to minimise risk to the already built and tested detectors.
- **Operation:** The system must have the ability to illuminate each side of every station independently, to minimise the current draw of the front end electronics and provide useful debugging information. In addition the system must only be operated in pulsed mode, as exposure of the active fibres to UV will reduce the light yield.
- **Magnetic Compatibility:** The system will be installed within the 4T uniform field region of the spectrometer solenoid.

4.3.2 Considered Approaches

To find the best solution to the problem several solutions were considered.

Passive Internal LEDs

The initial design used was to directly connect a small number of LEDs to a 50 Ω coaxial cable. The LEDs would then be mounted inside the tracker and driven by an externally connected pulse driver. A mock up of this approach was tested and it was found that in order to acceptably match the brightness of each LED the line had to be partially impedance matched. The matching required 150 Ω resistors in series with each LED followed by a terminating resistor on the final module. Additionally the pulse shaping caused by each LED module required a minimum pulse length of 40 ns, which was adequate but not desirable.

Furthermore the design required a pulse driver per channel capable of delivering a pulse between 6 V to 10 V into the cable over a 40 ns period. It was not possible to locate a commercial unit which supplied close to the number of necessary output channels (20) for a reasonable price.

The overall decision was not to use this design. The design would function however it would carry a large cost, and only provide adequate performance.

Driven Internal LEDs

The next iteration of the design was to move the pulse driving electronics close to the LEDs and make use of a commodity LED pulser circuit. The advantage is that the the cost of driving the LEDs is significantly reduced as the pulser can be driven by a small amount of control electronics. This approach also provides a more uniform light output and improved pulse characteristics. The disadvantage is that active electronics are installed inside a large magnetic field, restricting the use of certain components. In addition the limited access to the tracker would make modification or repair of the active components extremely difficult. For these reasons the decision was made not to place custom active electronics in the spectrometer solenoid.

External Driven LEDs

The final solution was to migrate the LEDs to outside the tracker body and use clear plastic fibre to inject the light inwards. This approach has two key advantages over installing LEDs within the MICE tracker:

1. The electrical components of the system can be upgraded or repaired at a later date without requiring access to the inside of the solenoid.

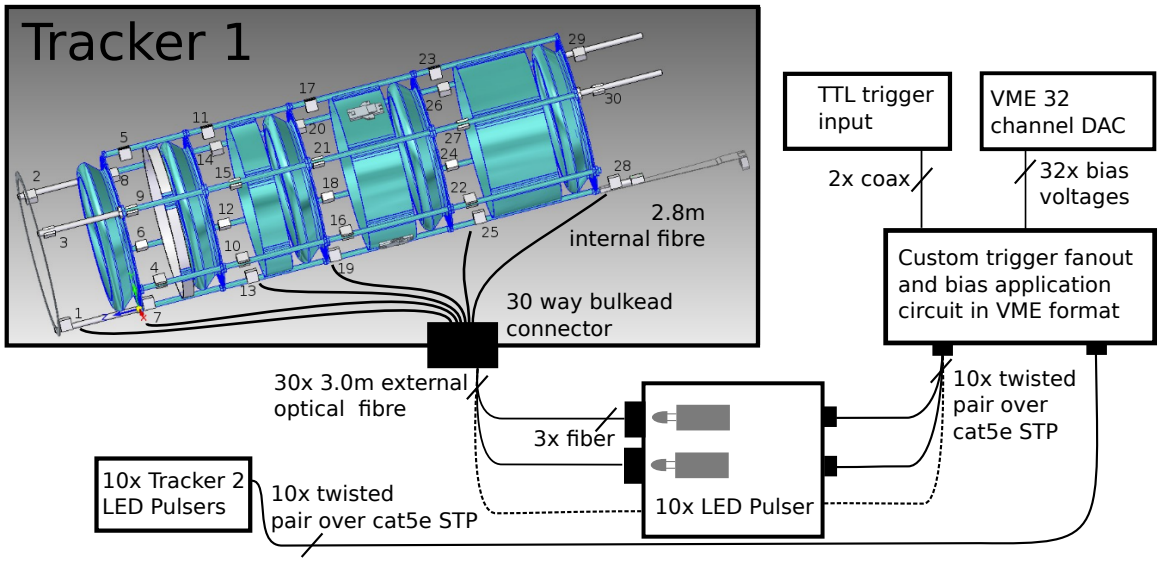

Figure 4.17: Overview of the light injection system, showing the joining of key components. The optical modules are shown on the tracker and labelled 1-30.

2. The LEDs can be connected directly to a pulse driver circuit, reducing the overall cost and improving pulse characteristics.

4.3.3 System Design and Construction

The light injection system is constructed from several components, each with a specific function. An optical module is fitted within the tracker and serves to diffuse the light from an optical fibre across an entire tracker station. A LED pulser is used to produce a short pulse of light into the fibre and is mounted outside the tracker. Finally a signal conditioning, control and trigger distribution subsystem for the LED pulsers has been designed. Figure 4.17 shows an overview of the components used in the system.

Optical Module

The limited space within the tracker body restricts the placement of the internal modules to be in close proximity to the scintillating fibre stations. The space limitations also prevent all the modules being placed at an identical distance from the station, however a second location was found for the modules which did not fit in the first. In order to deliver a sufficiently uniform distribution of light at the station required highly divergent light, which the optical fibre alone was incapable of producing. The

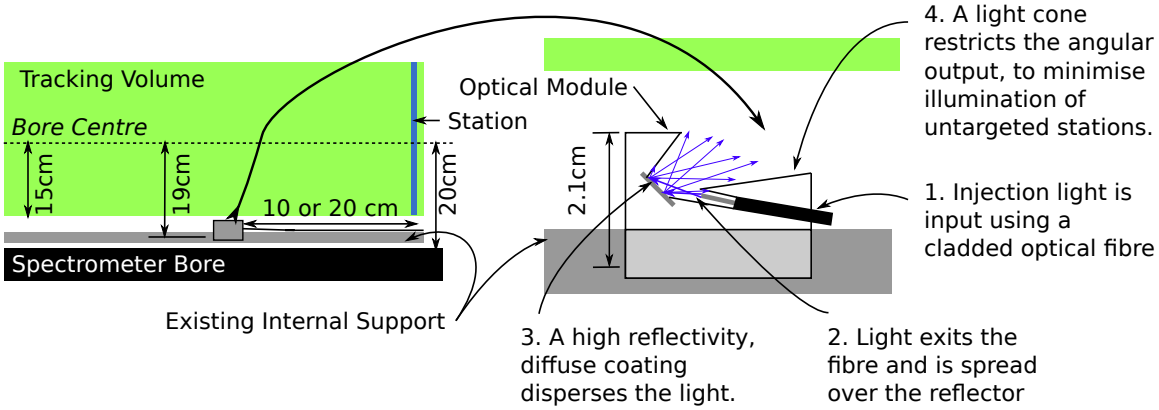

Figure 4.18: Sketch of the optical module, showing the key functional components and positioning in the spectrometer bore, with dimensions.

primary purpose of the optical module is to diffuse the light using a reflective surface similar to that found in integrating spheres. A sketch of an optical module is shown in figure 4.18.

To optimise the design of the module and determine the flux of light though the station a toy Monte Carlo model was used which assumed an almost Lambertian distribution of light coming from the module. In particular this model provided a tool to identify the optimum angle for the reflective surface. The expected fraction of light passing though the station is shown in figure 4.19. The XY profile indicates there will be dark areas, which are caused by the shadow of the carbon fibre station body and only occur on modules directed towards the inner (u) plane of the station. The amount of light expected in the fibre plane is very uniform in the centre of the station, however the light around the edges is much less due to the much shorter scintillating fibres. Unfortunately this is a design limitation and is extremely difficult to overcome without a significantly more complex light distribution method.

The modules were designed with an open source CAD package, OpenSCAD which constructs a model from the addition and subtraction of primitive shapes described in a file. The design combines several simple features to clamp to both the optical fibre and tracker supports, as shown in figure 4.20a. A light cone is also built into the module to improve strength and restrict the light output to the active area of the target station to reduce unwanted illumination.

For manufacture the modules are split into two parts and printed from black Acrylonitrile Butadiene Styrene (ABS) on a 3D printer. A printed module can be seen in figure 4.20d, where the reflector has been painted white. The rapid proto-

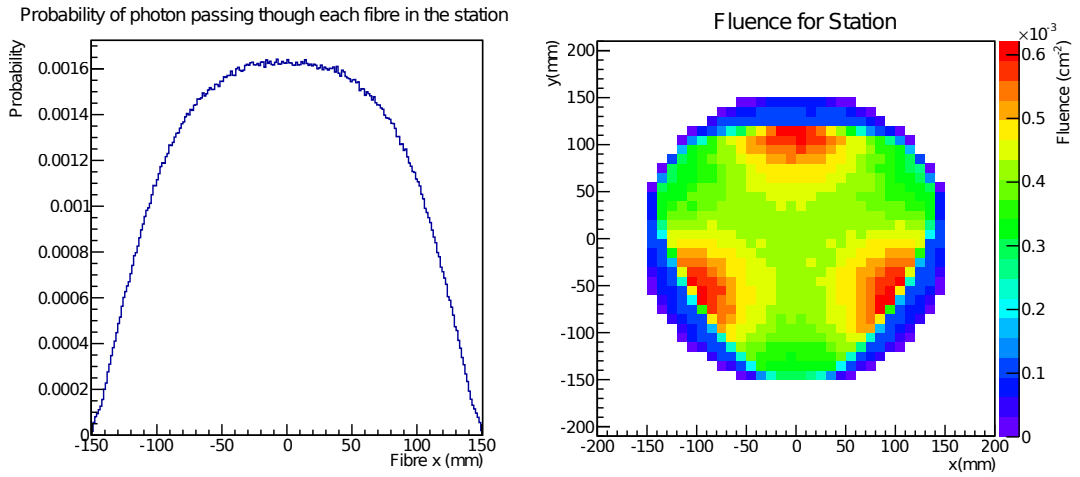

Figure 4.19: Expected probabilities for a photon emitted from the optical module passing through the station, modelled for a module 100 mm from the fibre plane, directed towards the inner (u) plane of the station

typing technology allowed several minor modifications to the design before the final production. The modifications included the optimisation of reflector size and sizing of the fibre optic path to tightly clamp the fibre. The position and size of the 3mm holes used to hold the modules together were also adjusted to allow the overall size of the module to be minimised. This optimisation process was then repeated for the second type of module which is placed further from the target station.

The next stage in the manufacture process was to apply a diffuse reflective coating to the module. The paint mixture comprised mostly of barium sulphate, which is commonly used in integrating sphere paint for its optical properties [72]. The barium sulphate (80% by weight) is mixed with Kraton elastomer (4% wt) to act as a binding agent and white spirit (16% wt) to act as a solvent[73]. The mixed paint was then applied to the modules in a thick coating as shown in figure 4.20d.

A total of 30 optical modules are required to illuminate each tracker, 24 of which were placed 100 mm from the station and a further 6 were placed 200 mm from the station.

Optical Fibre and Routing

Light to each of the modules is connected through an optical fibre. The fibre is split at the bulkhead connector of the tracker in a similar way to the existing lightguides. The internal optical fibre has a length of 2.8 m, which was chosen to match the length

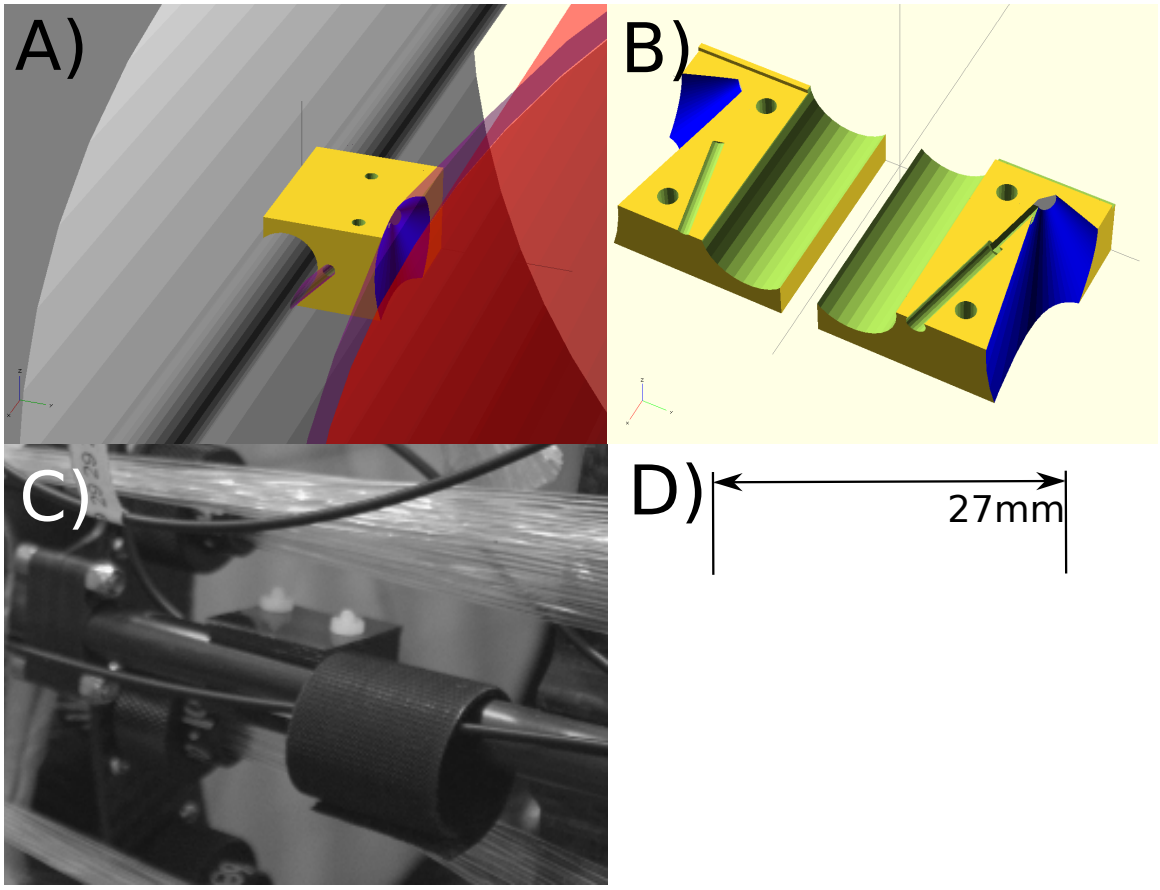

Figure 4.20: Images of the optical modules. a) Shows the module fitted on the tracker. b) Shows the module ready to be printed. c) The module mounted in the Tracker. d) The printed and painted module.

Loss Mechanism	Attenuation (dB)	Transmission (%)
LED Coupling	-9 (estimated)	12
External Optical Fibre	-0.45	95
Fibre to Fibre Coupling	-3	50
Internal Optical Fibre	-0.42	95.3
Total to Module	-13	5.1

Table 4.3: Estimated Optical Losses

of the internal lightguides. There is also 3.0 m of optical fibre external to the tracker, to allow the LED pulse unit to be placed a distance away from the tracker, due to the limited space around the cooling channel.

The fibre selected for this application was 1 mm diameter polymer fibre, which had a high cross section to maximise LED coupling and a low attenuation ($0.15\ \mathrm{dB\,m^{-1}}$) of 405 nm light [74]. In addition the fibre was cladded to prevent the illumination of unwanted channels. The estimated optical losses for the system are shown in table 4.3. The main source of loss is expected to be from the coupling of the LED to the optical fibres.

LED Pulser

The chosen LED pulse driver for the system is a low cost and high performance circuit, derived from the Kapustinsky [75] design. The specific variant was modified to use more modern transistors, reduce noise and be driven over a 100 Ω transmission line[76]. The intensity of the pulse is controlled by applying a bias voltage from 0 V to 20 V across the input connections. The pulse can then be triggered by a rising edge coupled to the bias voltage on the input pins.

The pulser circuits were originally fitted with a blue LED, which required replacing with a 405 nm LED[77]. Typical manufacturer data-sheets provide no information on the performance of the product when used for ultra short pulse applications, requiring that the LED to be tested with the pulser. First the pulse duration was tested using a fast (0.8 ns rise time) photomultiplier placed after an optical attenuator to protect the photomultiplier. The test demonstrated that the pulser could provide an acceptable pulse length (< 3.5 ns Full Width Half Maximum). The pulse length measurement was limited by photomultiplier and oscilloscope response.

Next the output power of the pulser was measured, using an optical power meter in close proximity to the LED. Figure 4.21 shows the pulse power to increase

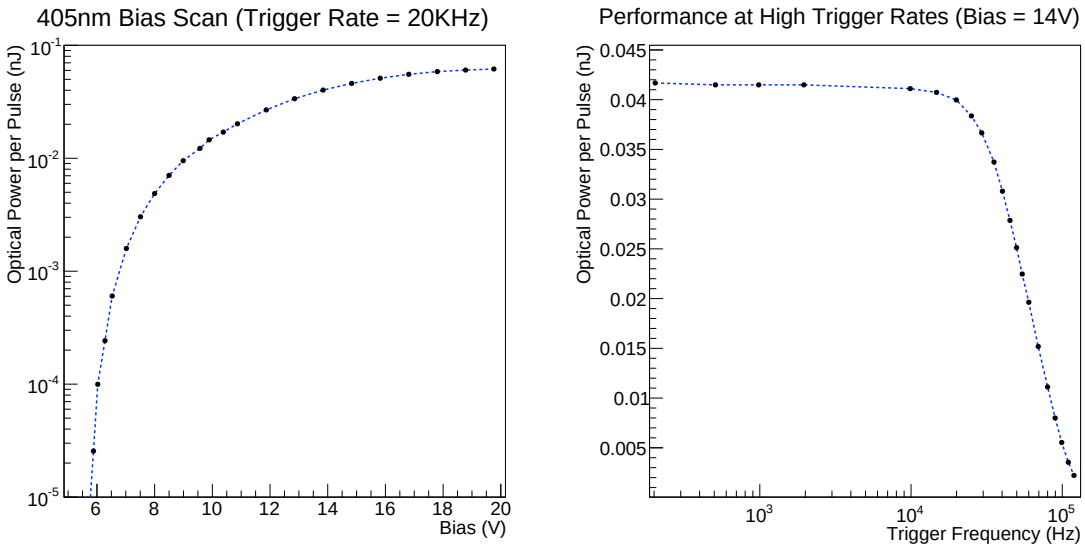

Figure 4.21: Performance of the 405 nm LED in the pulser

sharply as the bias voltage is raised from 6 V to 15 V, after which the LED begins to saturate. The number of photons produced was estimated to between 10^5 and 10^8, depending on bias voltage.

Finally the pulse power as a function of trigger frequency was measured to determine the maximum trigger rate of the pulse circuit. Figure 4.21 shows a sharp reduction in pulse power above 20 kHz. The limitation is caused by the finite recharge time of the capacitor used to drive the LED. The trigger frequency is much lower than the maximum trigger rate for the front end electronics, however it is not anticipated to significantly increase the calibration time.

Each LED pulser will be connected to the three optical fibres facing each side of each station and a total of 10 pulsers per tracker will be needed. The bias and trigger for each pulser is connected using a separate pair of a Category 5e Shielded Twisted Pair (Cat5e STP) cable, used commonly for computer networking.

Control and Timing

It is a request that all equipment operating in the MICE experimental hall should be remotely controlled wherever possible. The integration of the LED system with the existing control and monitoring infrastructure is also desirable.

Each LED pulser is required to have an individual bias control, to allow different modules to provide different intensity light (and to disable unwanted pulsers). There-

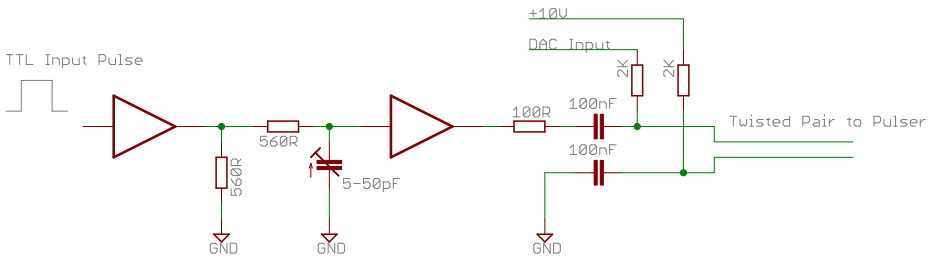

Figure 4.22: Simplified trigger fan-out and bias coupling circuit.

fore to include both trackers a total of 20 programmable bias voltages are required. The bias voltages are generated by a pair of Industry Pack DACs[78], which provide a total of 32 programmable outputs over a range of ± 10V. The packages are installed inside a VME Industry Pack carrier card[79], to allow the DACs to be fitted inside the existing VME crate used for tracker control and monitoring.

The output from the VME carrier is connected to a custom non-intelligent VME card which applies the bias voltages to the first wire of each pair. The other wire is connected to a common bias of $+10$V to allow operation over the full bias range. Figure 4.22 shows a simplified schematic for operation of a single channel. In addition to bias coupling the circuit also features a de-skewing element for each channel, to remove unwanted skew caused by different twisted pair lengths. Each de-skewed signal is capacitively coupled to the first wire of each twisted pair.

The non-intelligent VME card is powered from the VME bus, however it does not partake in any bus operations in order to maintain a simple design. Design of the final card has been expanded to allow up to 32 pulsers and the trigger for groups of four pulsers to be selected by means of an on board jumper.

The relative timing of each of the pulsers within a tracker will be set to minimise station to station skew. The time between trackers will be set to imitate the passage of a muon. The trigger source for the LED system will be provided externally and is expected to be generated and gated within the MLCR rack room trigger logic.

4.4 Conclusion

The MICE spectrometer has promise to deliver a high accuracy measurement of emittance, when installed in the beamline. For this to be accomplished a number of additional components need to be integrated. This chapter has described in depth the implementation of two of those components.

Tracker	Optical Module	Distance to station (mm)	Reflector angle	Pulser module	Bias channel	Trigger input
1	1,2,3	100	45	1	1	1
1	4,5,6	100	45	2	2	1
1	7,8,9	100	45	3	3	1
1	10,11,12	100	45	4	4	1
1	13,14,15	100	45	5	5	1
1	16,17,18	100	45	6	6	1
1	19,20,21	200	20	7	7	1
1	22,23,24	100	45	8	8	1
1	25,26,27	200	20	9	9	1
1	28,29,30	100	45	10	10	1
2	1,2,3	100	45	11	17	3
2	4,5,6	100	45	12	18	3
2	7,8,9	100	45	13	19	3
2	10,11,12	100	45	14	20	3
2	13,14,15	100	45	15	21	3
2	16,17,18	100	45	16	22	3
2	19,20,21	200	20	17	23	3
2	22,23,24	100	45	18	24	3
2	25,26,27	200	20	19	25	3
2	28,29,30	100	45	20	26	3

Table 4.4: Summary of module to pulser to trigger and bias mapping for the LED calibration system

The correlation of particles in MICE to an ISIS diagnostic signal has been verified. A system to convert the ISIS-1RF into a programmable gate has been designed, built and integrated with MICE. The integration of the system was tested with a subset of the detector electronics, demonstrating for the first time tracker hardware operating in the MICE beamline. The results from the test have been analysed and indicate similar performance to that seen during cosmic ray tests.

In addition a light injection system has been designed to allow in-situ calibration of the fibre trackers. The system is expected to provide an acceptable level of light to the majority of each station and provide a controllable intensity to help calibrate fibres which receive insufficient light during the first pass. The calibration system will also provide a short pulse of light to serve as a debug and diagnostic tool, in particular when setting up the trackers for MICE Step IV.

Chapter 5

Alignment Studies of the Step I Beamline

As part of the commissioning process the Step I MICE beamline was aligned and surveyed nominally to better than 0.2 mm, although operational observations have indicated that a potential misalignment could be causing an offset in the beam position at TOF2. The aim of this study is to verify the alignment of the individual quadrupole magnets within the Q789 triplet to the beam using a beam-based technique and the detectors available in the present beamline. In addition the scan will also be used to estimate the terms of the beam's covariance matrix at each quadrupole in an attempt to determine the emittance of the beam.

5.1 Strategy

5.1.1 Quadrupole Alignment

The relative beam position at each quadrupole can be measured by varying the selected quadrupole current. In the case where the beam is aligned to the quadrupole then symmetry around the axis ensures that on average there is no deflection. Any misalignment breaks this symmetry and causes an on average deflection of the beam, which depends on the magnet's focus strength.

This study will exploit this by varying the current in each quadrupole, while the others in the triplet remain unpowered. The beam position downstream of the quadrupole will then be measured in order to determine the amount of deflection. Unfortunately, without downstream tracking it is only possible to measure the beam

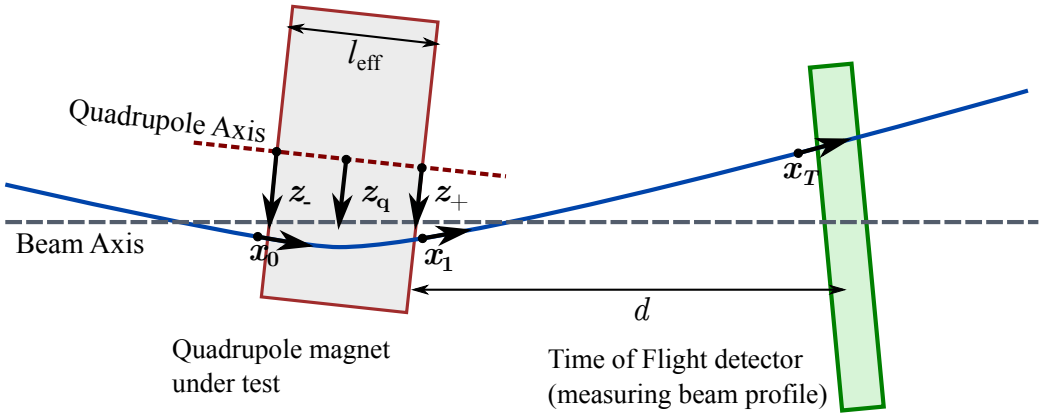

Figure 5.1: Quadrupole alignment to the beam axis; $\boldsymbol{z_-}$ is the misalignment at the entrance to quadrupole and $\boldsymbol{z_+}$ is the misalignment at the exit. $\boldsymbol{z_q}$ is the misalignment at the centre of the quadrupole.

position and width.

Using two dimensional linear beam optics, a particle can be traced through a magnet by multiplying the initial state vector by the appropriate transform matrix. In the case of a misaligned magnet the initial state ($\boldsymbol{x_0}$) is offset by a misalignment ($\boldsymbol{z_-}$) before the multiplication. The misalignment then is subtracted at the exit of the quadrupole to reconnect the particle with the original beam axis as shown in figure 5.1.

$$\boldsymbol{x_1} = \hat{F}(k, l_{\text{eff}})(\boldsymbol{x_0} + \boldsymbol{z_-}) - \boldsymbol{z_+} \tag{5.1}$$

where $\hat{F}(k, l_{\text{eff}})$ is the quadrupole focusing matrix, k is the focusing strength of the magnet (which can be scanned). l_{eff} is the effective length of the MICE type IV quadrupoles in this study and was previously determined to be 0.68 m[46]. Considering that the quadrupole could have an angular misalignment, the entry and exit offsets related by:

$$\boldsymbol{z_\pm} = \hat{O}(\pm l_{\text{eff}}/2)\boldsymbol{z_q} \tag{5.2}$$

where $\boldsymbol{z_q}$ is the misalignment at the centre of the quadrupole and it is the aim of this study to measure this value. $\hat{O}$ is the drift matrix, which propagates a state in empty space by extrapolating a straight line trajectory from the state. In this case the drift matrix is used to propagate the misalignment backwards and forwards from the centre of the magnet. The particle state can finally be propagated to the detector with:

$$\boldsymbol{x_T} = \hat{O}(d)\boldsymbol{x_1} \tag{5.3}$$

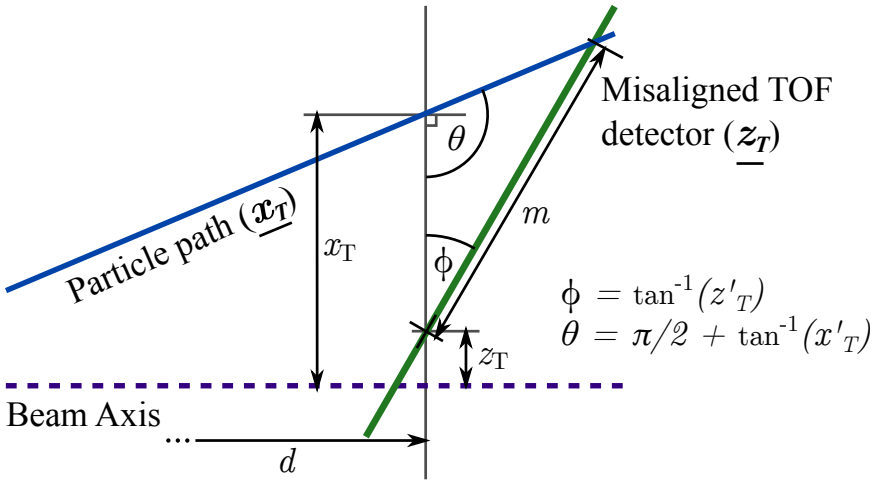

Figure 5.2: Sketch of a misaligned TOF position monitor from the beam axis. m is the value measured by the detector. Vectors are both bold and underlined for clarity, while their components as defined in equations 5.4 and 5.5 have been used to determine θ and ϕ. Specifically, the position offset of the detector is z_T, while the angular offset of the detector is z'_T.

where $\boldsymbol{x_T}$ is the particle state at TOF1 and d is the distance between the end of the quadrupole and the nominal location of the TOF.

The vectors used in this study to represent particle state ($\boldsymbol{x}$) and misalignment ($\boldsymbol{z}$) are as follows:

$$\boldsymbol{x_N} = \begin{pmatrix} x_N \\ x'_N \end{pmatrix} \begin{matrix} \text{Particle position} \\ \text{Particle angle from beam axis} \end{matrix} \tag{5.4}$$

$$\boldsymbol{z_T} = \begin{pmatrix} z_N \\ z'_N \end{pmatrix} \begin{matrix} \text{Magnet/Detector offset position (0 is aligned)} \\ \text{Magnet/Detector offset angle (0 is aligned)} \end{matrix} \tag{5.5}$$

5.1.2 Detector Alignment

It is also important to consider the alignment of the detector used to measure the beam position. If left unchecked this could introduce substantial systematic error, particularly if the location of the detector during data collection was not accurately known. Figure 5.2 shows a sketch of a misaligned time of flight detector.

Using the sketch and simple geometry the measured position(m) at the detector can be written as a function of incoming particle state ($\boldsymbol{x_T}$) and detector misalign-

ment ($\boldsymbol{z_T}$):

$$\begin{aligned} m &= (x_\mathrm{T} - z_\mathrm{T})\frac{\sin(\theta)}{\sin(\pi - \phi - \theta)} \\ &= (x_\mathrm{T} - z_\mathrm{T})\frac{\cos(\theta - \frac{\pi}{2})}{\cos(\frac{\pi}{2} - \phi - \theta)} \\ &\approx (x_\mathrm{T} - z_\mathrm{T})\frac{2 - x'^2_\mathrm{T}}{2 - (x'_\mathrm{T} + z'_\mathrm{T})^2} \end{aligned} \tag{5.6}$$

where x_T is the particle position at the detector plane, x'_T is the angle between the particle and the beam axis at the detector plane, z_T is the detector offset from the beam axis and z'_T is the angular offset of the detector from alignment. The approximation arises from the use of small angle approximations. An offset in the position of the detector has a linear impact on the position. The position of the detector will be considered a free parameter in the misalignment fitting procedure to reduce the systematic error it introduces. If the detector also has an angular offset a scaling error is introduced. Since the effect is introduced at second order even a relatively large angular offset is suppressed significantly and no attempt will be made to further reduce this. Additionally a rotation along the beam axis would result in a similar effect as would a skew of the detector, both of which are ignored by this study.

5.1.3 Mean Position

The mean position of the beam measured by the detector with no angular misalignment can be found by combining the above equations into:

$$\overline{m(k)} = \frac{1}{n}\sum_{i=1}^{n}\left(\hat{O}(d)\left(\hat{F}(k)\left(\boldsymbol{x_{0,i}} + \hat{O}(-l_{eff}/2)\boldsymbol{z_q}\right) - \hat{O}(+l_{eff}/2)\boldsymbol{z_q}\right)\right)_1 - z_T \tag{5.7}$$

where the mean is taken over n particles and the 1 denotes that only the first element of the resulting vector is used. Since the linear beam optics equation only consists of linear terms, the sum can be migrated to a sum over just $\boldsymbol{x_{0,i}}$. The first moments of a particle beam on axis evaluate to zero, meaning the mean initial state is:

$$\frac{1}{n}\sum_{i=1}^{n}\boldsymbol{x_{0,i}} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \tag{5.8}$$

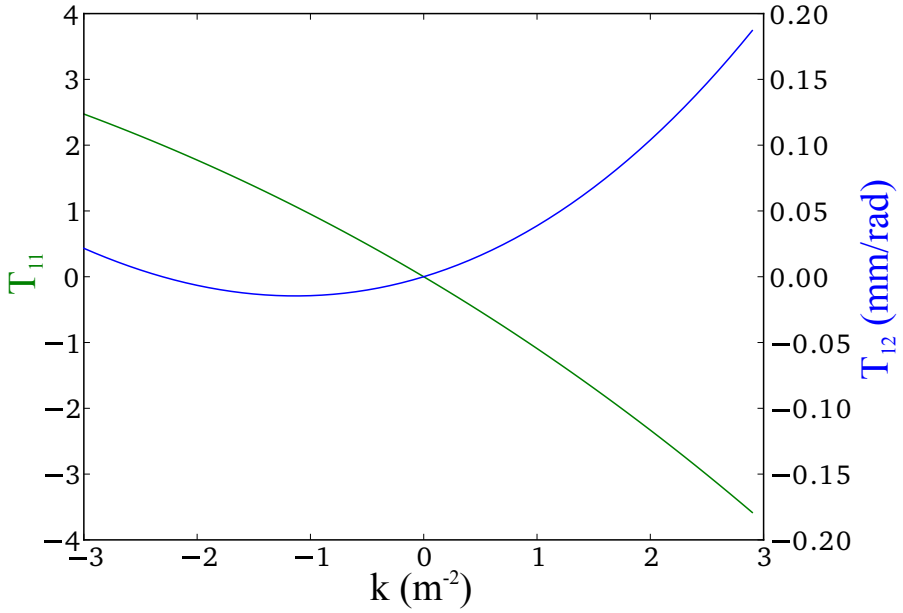

Figure 5.3: Estimated values of $T_{11}(k)$ and $T_{12}(k)$ for a scan of quadrupole 7 (T is defined in equation 5.10). When $k < 0$ the quadrupole is focusing.

and the equation becomes:

$$\overline{m(k)} = T_{11}(k)z_q + T_{12}(k)z'_q - z_T \tag{5.9}$$

This equation has a total of three unknown quantities: z_q the position of the beam relative to the quadrupole centre, z'_q the angle of the beam relative to the quadrupole, z_T the offset of the TOF detector from the beam. These parameters can be estimated by fitting the equation through the evaluation of the matrix $\hat{T}(k)$:

$$\hat{T}(k) = \hat{O}(d)\left(\hat{F}(k)\hat{O}(-l_{eff}/2) - \hat{O}(+l_{eff}/2)\right) \tag{5.10}$$

where $\hat{T}(k)$ is comprised entirely from knowledge of the beamline geometry and quadrupole settings. Figure 5.3 shows the value of two of the matrix elements for a scan of quadrupole 7. Since T_{11} is not proportional to T_{12} it is possible given sufficient data to disentangle the angular offset from a position offset.

5.1.4 Emittance Measurement

In addition to measuring alignment, a quadrupole scan can also be used to measure emittance[80]. This relies on the principle that the covariance matrix can also be transformed from an initial to final state, specifically:

$$\Sigma_T(k) = \hat{R}(k)\Sigma_Q\hat{R}^T(k) \tag{5.11}$$

where $\Sigma_T(k)$ is the covariance matrix of the beam at the TOF and Σ_Q is the covariance matrix just before the scanned quadrupole. The transfer matrix $\hat{R}(k)$ for this study is:

$$\hat{R}(k) = \hat{O}(d)\hat{F}(k) \tag{5.12}$$

Using solely the particle position to measure the beam width it is possible to determine the first element of the Σ_T matrix. As the strength of the quadrupole is changed, elements within the transfer matrix change strength, revealing the three unknown quantities in the initial Σ_Q matrix. The transfer matrix can be computed with knowledge of the beamline settings. The three remaining unknown parameters inside Σ_Q matrix can be found by fitting the predicted Σ_T element to the measured Σ_T.

Once the complete covariance matrix is known the emittance calculation is trivial using equation 2.7.

5.2 Beam Selection and Simulation

The anticipated limiting factor for the study was the accuracy with which the mean position of the beam could be obtained, which is limited by the number of particles studied. In order to maximise the particle rate a "pion" beam was selected for study, for which both dipoles in the beamline are set to a common momentum causing maximum particle rate. The initial study was undertaken during the Step I data campaign, which included all Step I diagnostics, plus an additional beam monitor. The beam monitor was placed upstream of TOF1, but did not generate any useful data. Once the data campaign was complete the beamline was surveyed, not however before TOF1 was disturbed. In addition the configuration contains a possible miscalibration of the detector which could introduce a 2.4% decrease to the measured time difference between detectors[81].

In order to clarify issues within the 2010 dataset the data was re-taken during the summer of 2013. Unfortunately the decay solenoid was unavailable for the scan, which

	2010 data	2013 data	Monte Carlo
Software	MAUS 0.7.0	MAUS 0.8.1	G4Beamline 2.14
Decay Solenoid (A)	100	0	100
Proton Absorber Thickness (mm)	83	78	83
TOF0 (z) (m)	21.0705	21.0688	21.0705
Q7 (z) (m)	25.2937	25.293	25.293
Q8 (z) (m)	26.4537	26.4537	26.4537
Q9 (z) (m)	27.6137	27.6137	27.6137
TOF1 (z) (m)	28.7931 [a]	28.754	28.7735
TOF1 fine resolution (mm)	12.7	12.1	12.7

Table 5.1: Essential configurations for the alignment study data analysis. Distance measured along beam axis from the target.[84, 85]

[a] Replaced before survey, accurate to 1 cm

significantly reduced the number of particles collected. The number of scan points was reduced in order to maximise the number of particles at each point, enabling a more accurate parametrisation. Furthermore a single slab of the proton absorber jammed, just prior to the run, in turn fixing the thickness differently to the initial study.

A simulation of the initial study was computed using G4Beamline and O. Hansens[82] deck, with magnet currents adjusted to match the collected data. Detector truth hits were smeared before use in the analysis according to table 2.2. In addition a number of virtual detector planes were placed upstream of each quadrupole in order to enable a truth measurement of the beam parameters. The post-processing of the virtual detector hits into a mean position and covariance matrix was handled by a post-processing tool, XBOA[83].

Settings for the beamline in each dataset are shown in table 5.1.

5.3 Data selection and Parametrisation

5.3.1 Particle Selection and Identification

The measurement of the quadrupole focusing strength (k) relies on knowledge of the momentum of the beam. To study the momentum spread the velocity of each particle

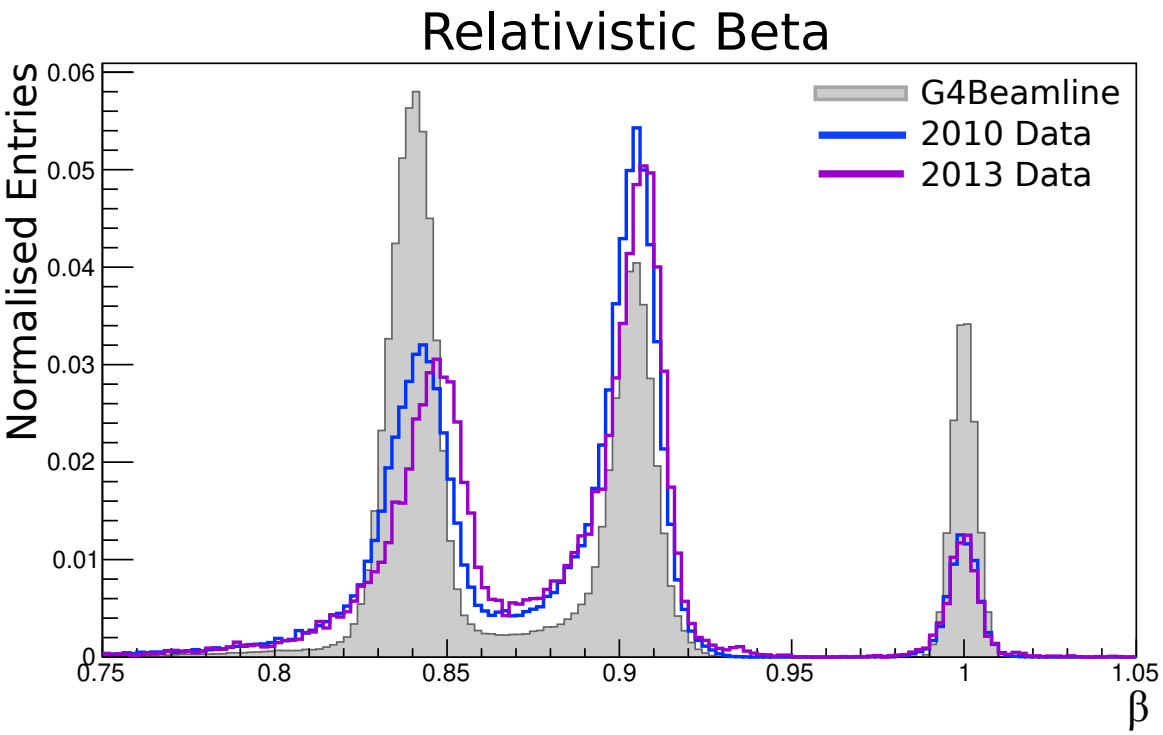

Figure 5.4: Calculated β for each dataset. Plotted on a normalised histogram.

was calculated, between the TOF0 and TOF1 detectors, using:

$$\beta = \frac{\overline{t_e}}{t} \tag{5.13}$$

where $\overline{t_e}$ is the mean time of flight for an electron and t is the particle time of flight. This simple calculation does not consider different particle path lengths, a feature introduced by the beamline magnets. However using the electron time of flight as a reference significantly reduces uncertainty from the detector location and time scaling.

The constrained momentum range of the particles downstream of the dipole magnets enables the time of flight measurement to operate as a form of particle identification. This can be seen in figure 5.4, which shows the computed relativistic beta for all data used in the study. The three peaks indicate the presence of electrons, muons and pions in the beamline and the integral of each peak can be used as a rough guide to the fraction of each particle species in the beamline. A quick observation is that the Monte Carlo both underestimates the fraction of muons within the beam and the background hits between the peaks. The two data taking campaigns show a good degree of consistency, with the newer data carrying a slightly higher momentum. A

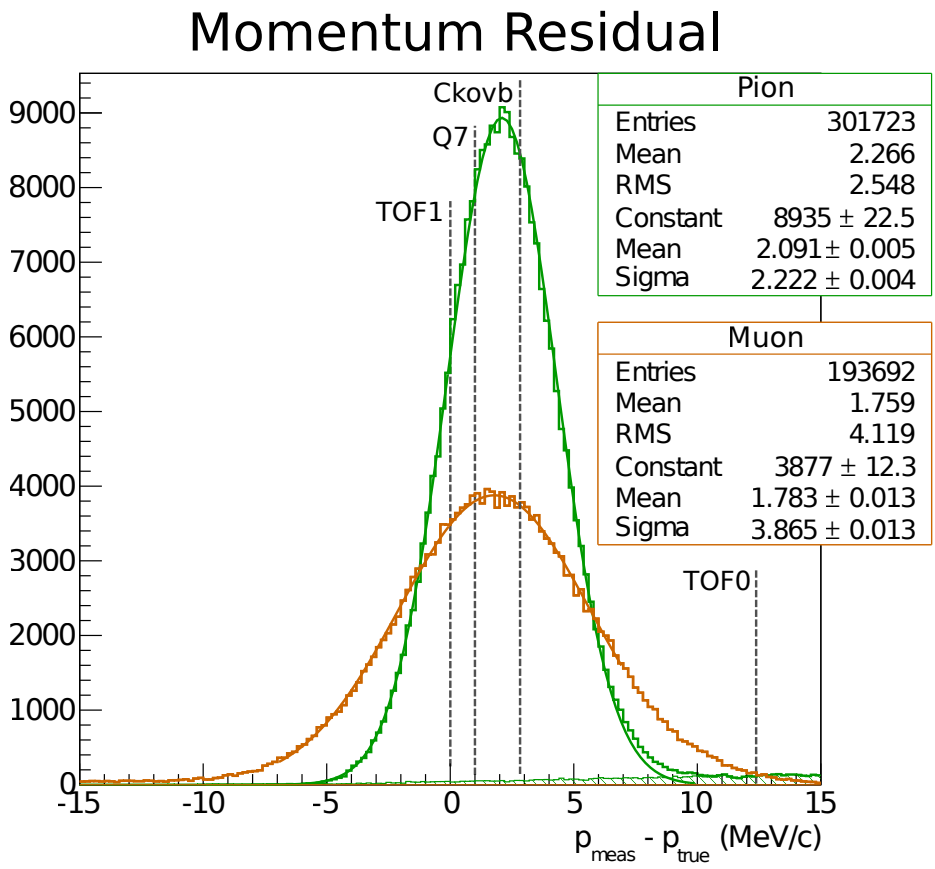

Figure 5.5: Residual momentum between estimation and the true momentum at TOF1. Matter effects reduce the momentum along the beamline, such that the difference for muons is indicated

possible source could be less material in the MICE beamline, caused by the removal of the beam position monitor and a thinner proton absorber used in the study.

Particle identification is performed by placing cuts on β, enabling the mass to be input and therefore momentum to be calculated. Figure 5.5 shows the residual between the estimated longitudinal momentum and the true value at TOF1 for both pions and muons. Matter effects along the beamline cause the particles to lose momentum, meaning that the residuals are not centred on zero. Furthermore the momentum at the detector is also subtly different from that at each quadrupole magnet ($< 1\%$). These effects were estimated using the MICE materials budget spreadsheet[86] and a correction factor applied for each quadrupole during the estimation of each quadrupole focusing strength. The residual width is a combination of different path lengths between the detectors and the resolution of the time of flight measurement. The pions have a larger time of flight, which reduces the error in the

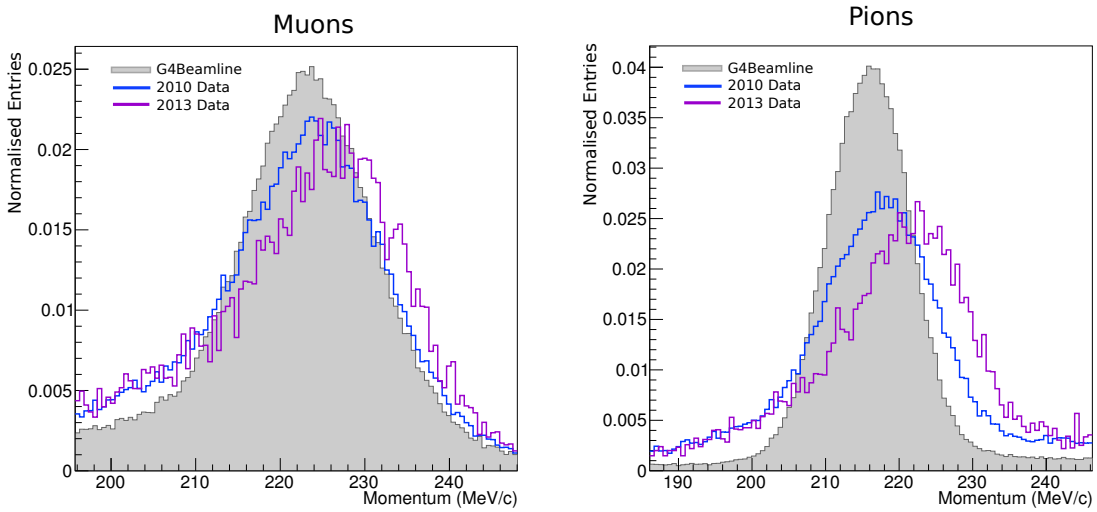

Figure 5.6: Longitudinal momentum of pions and muons, to be used in the calculation of k. Only particles within the momentum range displayed is kept from figure 5.4, which alters the relative normalisation. The particles to used in the study are cut, accepting only a longitudinal momentum between 205 and 235 MeV/c.

velocity calculation and reduces the overall residual.

The momentum of both the pion and muon segments of the beam are shown in figure 5.6. To avoid the background events between the peaks, only particles within the momentum range 205 and 235 MeV/c are accepted by the study.

5.3.2 Time of Flight stations as Beam Profile Monitors

The time of flight detectors used in MICE were originally intended for precision timing measurements of each particle travelling the beamline. The detectors themselves were not intended for accurate profiling of the beam. The original set up forms a detector of square pixels 60 mm wide in the case of TOF1. A coarse pixel based detector has insufficient resolution for this study.

As part of the Step I reconstruction a technique was developed and integrated with the G4MICE analysis package which enabled a differential time measurement along the bar. This technique enables a more precise (fine) position measurement, which had easily understood Gaussian like residuals. These features make the TOF system a viable beam position monitor and are considered a necessity for this study.

The G4MICE analysis code was been superseded by the MAUS analysis pack-

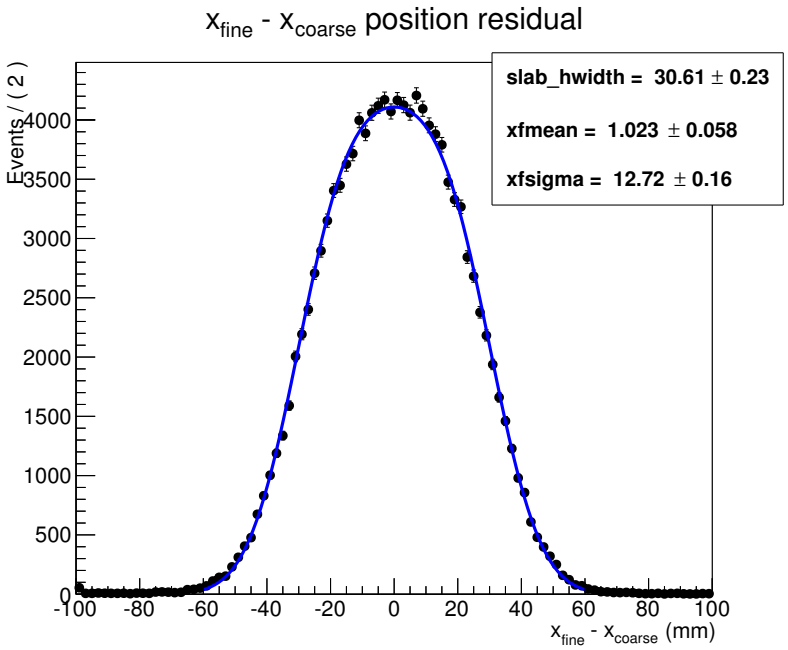

Figure 5.7: TOF position residual between the coarse (perpendicular bar) and fine (differential time) measurements, for the 2010 x axis data. The data is fitted with a convolution of a uniform distribution, with width (`2slab_hwidth`) and a Gaussian with mean (`xfmean`) and width (`xfsigma`). The estimated resolution of the fine position measurement is `xfsigma`.

age and is now unsupported. To enable the use of the fine position measurement a short script was written to re-implement the methods used in G4MICE. A calibration technique was also written in order to determine the effective speed of light in the scintillator bar and each slab's differential cable delay. The calibration was verified by plotting the difference between the coarse position, derived from the perpendicular slab hit and the fine position derived from the differential time measurement. Figure 5.7 shows the result of the test and allows an estimation of the fine position resolution.

The understood residuals of the fine position estimation indicate that both the calculation of the position and calibration steps were performed correctly, in turn enabling the fine position of each particle to be combined into a beam profile, for use in the study.

	Beam RMS Width (mm)		
Quadrupole Centre	X	Y	Selected Fraction
Q7	132	134	1.3σ
Q8	183	182	0.96σ
Q9	239	234	0.74σ

Table 5.2: Estimated beam width at each quadrupole's centre, projected from the virtual detector just upstream from Q7 in G4Beamline assuming no field. The selected fraction of the beam used is given as a fraction of beam width(σ), assuming a normal distribution.

5.3.3 Beam Profile Parametrisation

The next step was to use the measured fine position to parametrise the beam profile at the detector for each magnet setting. During the scan the focal length of each quadrupole magnet changed dramatically, which in turn significantly altered the width of the beam at the downstream detector. In many scans a significant proportion of the beam was lost beyond the range of the detector and in the quadrupoles. In order to accurately parametrise the beam at all magnet settings a normal distribution was fitted to particle hits on each axis. To ensure only the beam distribution was fitted, and not shadows of the quadrupole, the fit range was constrained by the following factors:

1. The physical detector size. The active length of the detector is 420 mm, of which the central 380 mm is used.

2. Losses downstream of the scanned quadrupole, induced by the physical aperture (376 mm diameter). The fit range was restricted to the central 240 mm of the beam.

3. Losses up to and including the scanned quadrupole, caused by the increase in beam size along the quadrupole triplet. The restriction on fit range is specified as a fraction of the beam width seen at the detector, since only the beam edges are lost. The increasing beam width at each successive quadrupole required a different cut for each quadrupole, shown in table 5.2.

In order to obtain the best fit results for each magnet current an unbinned likelihood fit was performed, removing the need to determine the optimum bin width for each scan setting. Each axis was fitted independently and a number of fits can be

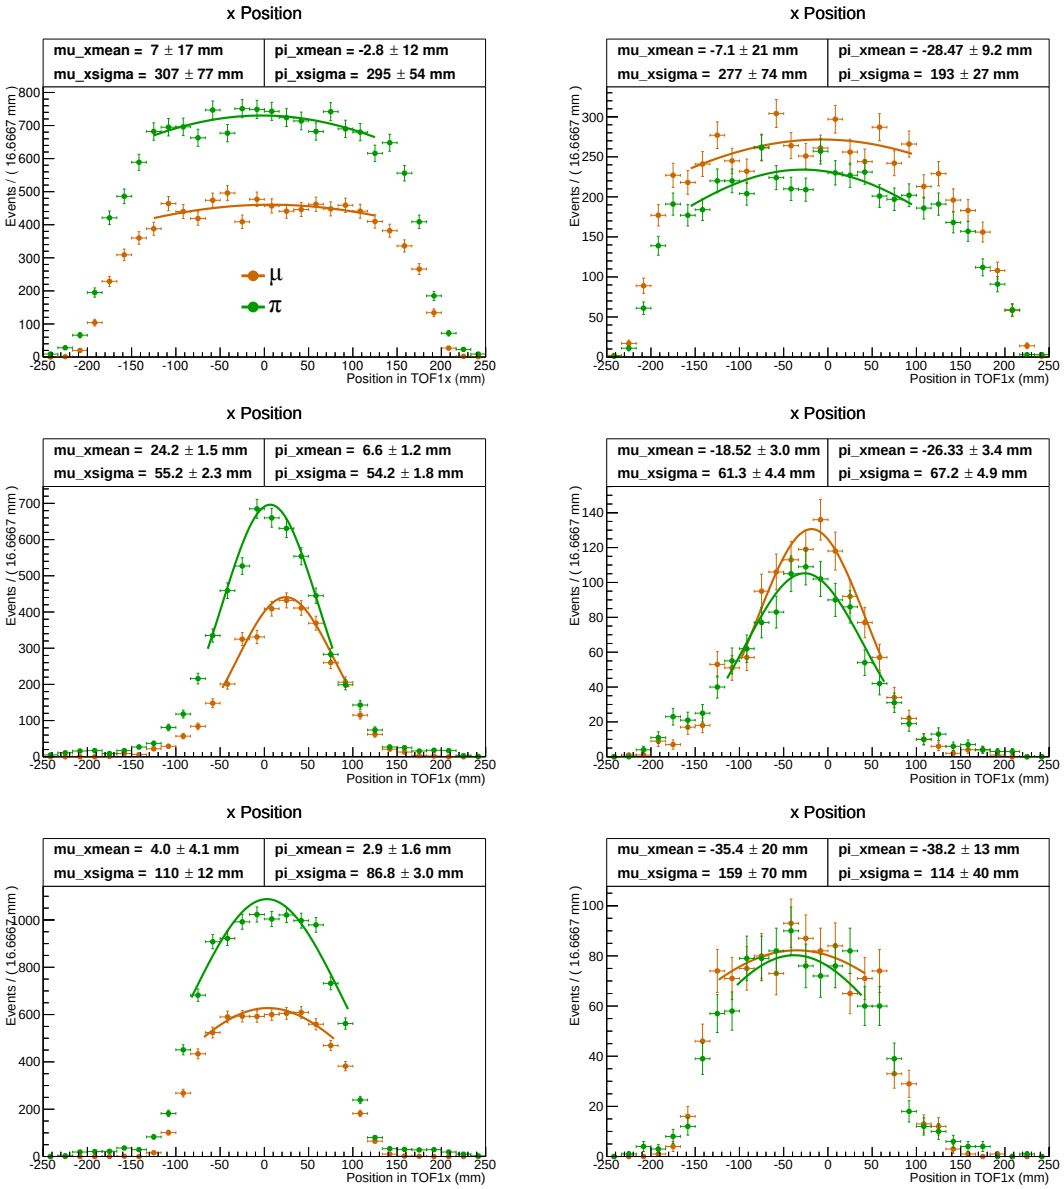

Figure 5.8: Fits to beam projections. Left Monte Carlo, Right 2010 dataset. Top: Q7,8,9 = 0A, Middle: Q7=120A Q8,9 = 0A, Bottom:Q7,8 = 0A, Q9=120A

seen in figure 5.8. A large amount of error is incurred from the fit range constraints and is most notable when the beam is defocused in the detector. The fit parameters are then passed to the next stage of analysis, in which the mean position is used to determine the magnet alignment and the beam width which is needed to determine the beam covariance matrix.

5.4 Alignment Studies

The final stage of the alignment study was to estimate the misalignment of each quadrupole using the parametrised data. A chi-square fit to equation 5.9 was used to determine the misalignment of the quadrupoles and TOF. Since the misalignment of the TOF detector was common to each scan, the three scans were fitted simultaneously. In addition quadrupole 8 defocuses x, significantly increasing the error in the resulting fit parameters, the same is true for y in quadrupoles 7 and 9.

The fit process demonstrated a correlation exists between the position and angular offset of each quadrupole magnet, indicating that both parameters could not be accurately determined. In addition the correlation enabled the minimiser to find non-physical minima, which had angular offsets in excess of 0.4 radians. The large correlations are caused by a substantial variation in error on each point, induced by a changing beam width from varying quadrupole strengths. For a focusing quadrupole the effect of an angular offset is relatively very small, such that a 5° offset would move the beam by substantially less than a 1 mm offset. Since a large correlation exits between the two measurements the most simple solution was to fix a free parameter, specifically the angular offset, to zero.

The fits for each scanned magnet are shown in figures 5.9 and 5.10, with the fit results in table 5.3. The first two columns enable a comparison between an estimation derived from the virtual detectors inserted into the Monte Carlo at each quadrupole to measure the beam profile and offset (truth) with the the fitted scan results. In all cases the fitted results compared well with the truth-derived values and demonstrated a viable technique. No additional offsets were introduced into the simulation and the offsets observed in Monte Carlo truth are from an offset between the beam position and nominal beam axis.

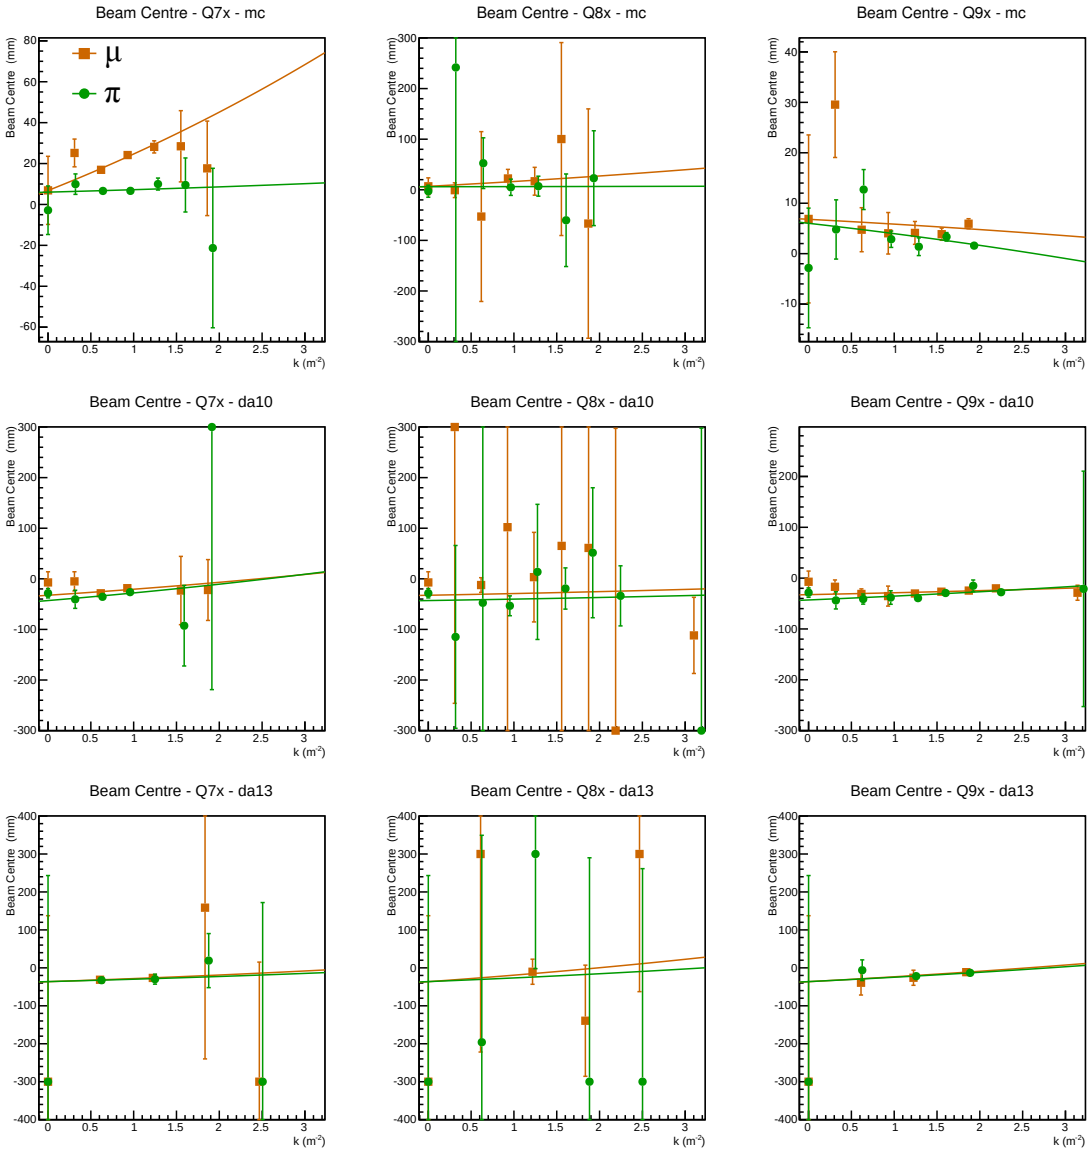

Figure 5.9: Fits to the x axis for Q7 to Q9 (left column to right column). Top: Monte Carlo, middle: 2010 data, bottom 2013 data.

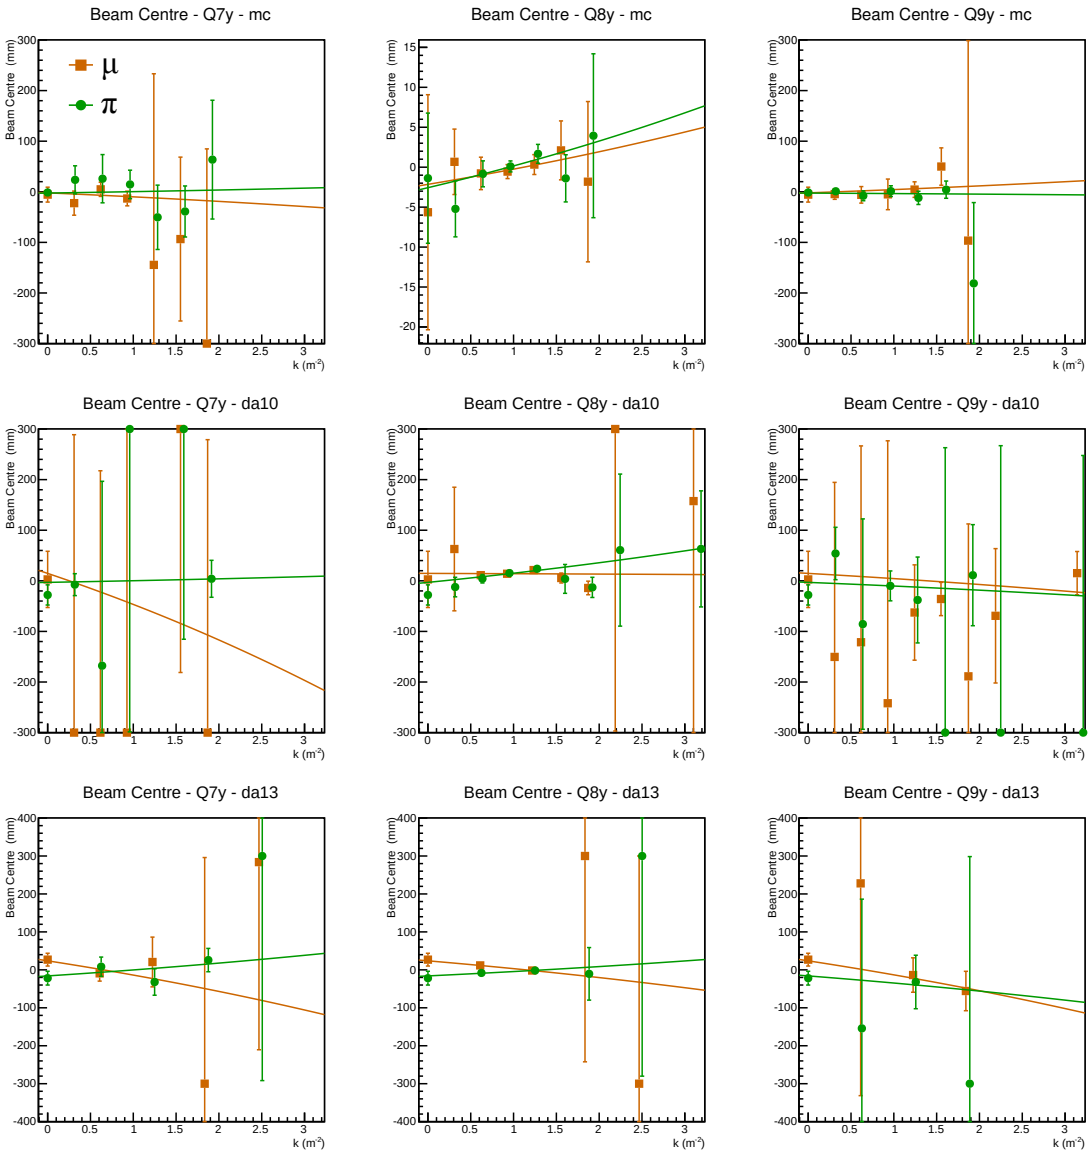

Figure 5.10: Fits to the y axis for Q7 to Q9 (left column to right column). Top: Monte Carlo, middle: 2010 data, bottom 2013 data.

			Monte Carlo		Data	
Axis	Parameter		Truth	Fit	2010	2013
x	z_t (mm)	μ	-	6.8 ± 2.6	32.5 ± 4.6	36 ± 11
		π	-	6.0 ± 1.8	43.0 ± 5.0	36.2 ± 7.2
	z_7 (mm)	μ	-4.1 ± 0.9	-7.0 ± 1.2	-4.7 ± 2.5	-3.2 ± 5.8
		π	-2.0 ± 0.7	-0.4 ± 0.8	-5.9 ± 2.7	-2.5 ± 4.4
	z_8 (mm)	μ	0.9 ± 2.0	-5.4 ± 8.2	-2.0 ± 9.3	-10 ± 15
		π	3.0 ± 1.7	-0.1 ± 6.3	-1.6 ± 7.8	-6 ± 79
	z_9 (mm)	μ	5.4 ± 2.9	1.1 ± 1.8	-4.3 ± 3.2	-15.0 ± 6.6
		π	2.0 ± 2.1	2.4 ± 1.3	-8.8 ± 3.6	-13.5 ± 4.9
	χ^2/n.d.f	μ	-	15/14	12/17	2.9/7
		π	-	11/14	10/17	3.6/7
y	z_t (mm)	μ	-	-2.1 ± 2.5	-14.6 ± 6.8	-24 ± 11
		π	-	-2.6 ± 2.0	-3.1 ± 8.2	16 ± 10
	z_7 (mm)	μ	-0.2 ± 1.0	3.1 ± 5.1	24 ± 75	15 ± 13
		π	-4.1 ± 0.7	-1.1 ± 6.9	-1.3 ± 7.1	-6.2 ± 5.7
	z_8 (mm)	μ	-2.4 ± 2.0	-1.1 ± 1.4	0.4 ± 3.9	12.1 ± 5.7
		π	-1.8 ± 1.4	-1.6 ± 1.1	-10.6 ± 4.7	-6.7 ± 6.5
	z_9 (mm)	μ	-1.4 ± 2.8	-8 ± 11	12 ± 12	43 ± 26
		π	-2.5 ± 2.0	1.1 ± 6.9	8 ± 29	21 ± 64
	χ^2/n.d.f	μ	-	8.5/14	14/17	2.4/7
		π	-	8.1/14	13/17	2.5/7

Table 5.3: x and y axis alignment data. Note that Q8 in x is defocusing and Q7, Q9 are in defocusing in y, which degrades the ability to find the beam centre. The truth values are determined from virtual detectors placed in the Monte Carlo at each quadrupole.

5.4.1 Interpretation of Results

The measurement technique obtains the relative offset between the beam centre and each quadrupole centre. An offset can then be observed when the quadrupoles are aligned to the surveyed beam axis, but there is an offset between the beam centre and surveyed axis, which can be seen in the Monte Carlo data.

The fit results indicate a possible x offset in quadrupole 9, present in both the 2010 and 2013 datasets, where three of the four scans indicate a larger than 2σ offset of the beam from the quadrupole axis. There was no significant sign of an offset which was consistent between both particle species, or data campaigns in the y axis data.

In addition to the quadrupoles, the scans clearly show an offset of the TOF1 detector, which has been observed elsewhere. The Step I analysis observed an offset(x=30.2 mm, y=−17.6 mm [1])[87], which seems consistent with the observed μ offset. The latest survey of the beamline indicates a small offset (x=13.25 mm, y=1.84 mm)[85], which does not fully account for the offset observed.

A possible explanation is that the beam path is notably different from the (surveyed) beam axis, which results in an off centre position in both the Q9 quadrupole and TOF1 detector. Consider now that both the beam path and beam axis cross in the second dipole (D2), from here the TOF1 misalignment can be estimated from a quadrupole ($\mathrm{Q}n$) geometrically:

$$z_{\mathrm{T}}(n) = -z_{\mathrm{n}}\frac{D(\mathrm{D2} \rightarrow \mathrm{TOF1})}{D(\mathrm{D2} \rightarrow \mathrm{Q}n)} + h_{\mathrm{TOF1}} \tag{5.14}$$

where $D(\ldots)$ is the z displacement between beamline objects and h_{TOF1} is the surveyed location of the detector. Table 5.4 shows these estimated positions from both Q9 and Q8 in the x and y axes respectively. The estimated values are consistent with the measured values, however note the substantial uncertainty in both values.

5.5 Covariance Studies

The final part of the study was to estimate the covariance matrix for the beam just upstream of each quadrupole. The true width of the beam will be smeared by the

[1] Once converted into right-handed G4BL co-ordinates, from left-handed G4MICE co-ordinates by inverting the x axis. Also note the step one studies used μ beams, which use a different D1 current.

Data	Quadrupole(Qn)	Fraction ($\frac{D(...)}{D(...)}$)	Estimated z_{T}(mm)	Residual (σ)
$\mu\ x$	Q9	1.13	30.2 ± 7.5	0.3
$\pi\ x$	Q9	1.13	28.5 ± 5.5	0.6
$\mu\ y$	Q8	1.26	-13.4 ± 7.1	0.6
$\pi\ y$	Q8	1.26	10.3 ± 8.2	0.3

Table 5.4: Estimated values of z_{T} from misaligned beam and survey data, using equation 5.14 and 2013 fit results

resolution of the TOF detector. Since both the TOF residual and beam width are a normal distribution the observed width of the beam is a quadrature sum of both:

$$\sigma_m^2(k) = \Sigma_{T,11}(k) + \sigma_R^2 \tag{5.15}$$

where σ_m is the width of the fitted normal distribution, $\Sigma_{T,11}$ is the variance of the beam at the detector and σ_R^2 is the detector resolution. Since the detector resolution is already known, it is possible to subtract this effect in order to obtain the original beam width. Next incorporating the $\Sigma_{T,11}(k)$ results with equation 5.11 and a fit allows the covariance matrix at each quadrupole can be determined. The fit is only effective for focusing quadrupoles, as this is the only case where the beam comes to a waist, which reveals the unknown terms in the matrix and therefore restricts this part of the study to those magnets.

In order to also determine the covariance matrix directly from the Monte Carlo, hits from virtual detectors located just upstream of each quadrupole were processed with XBOA. The covariance matrix was next propagated to the same position in z as the Σ_Q matrix, specifically 34 cm upstream from the quadrupole centre.

The fit to each quadrupole is shown in figure 5.11, with the fit results in table 5.5 for completeness. As with the alignment study a lack of well known data points added a substantial error to the fit. In particular a lack of data points around the waist of the scan enabled some fits to have negative beam widths, a clearly non-physical result. An attempt was made to also calculate the emittance, however substantial error on each of the covariance terms made a meaningful measurement impossible.

A comparison between the output from XBOA and the fitted Monte Carlo indicate an acceptable agreement for the pion data at quadrupole 7. The muon data however shows a notable underestimation of the fit parameters. Beyond the first magnet, the XBOA calculation significantly differs from the fitted measurement. This is to be expected since XBOA directly evaluates the moments of the beam, the

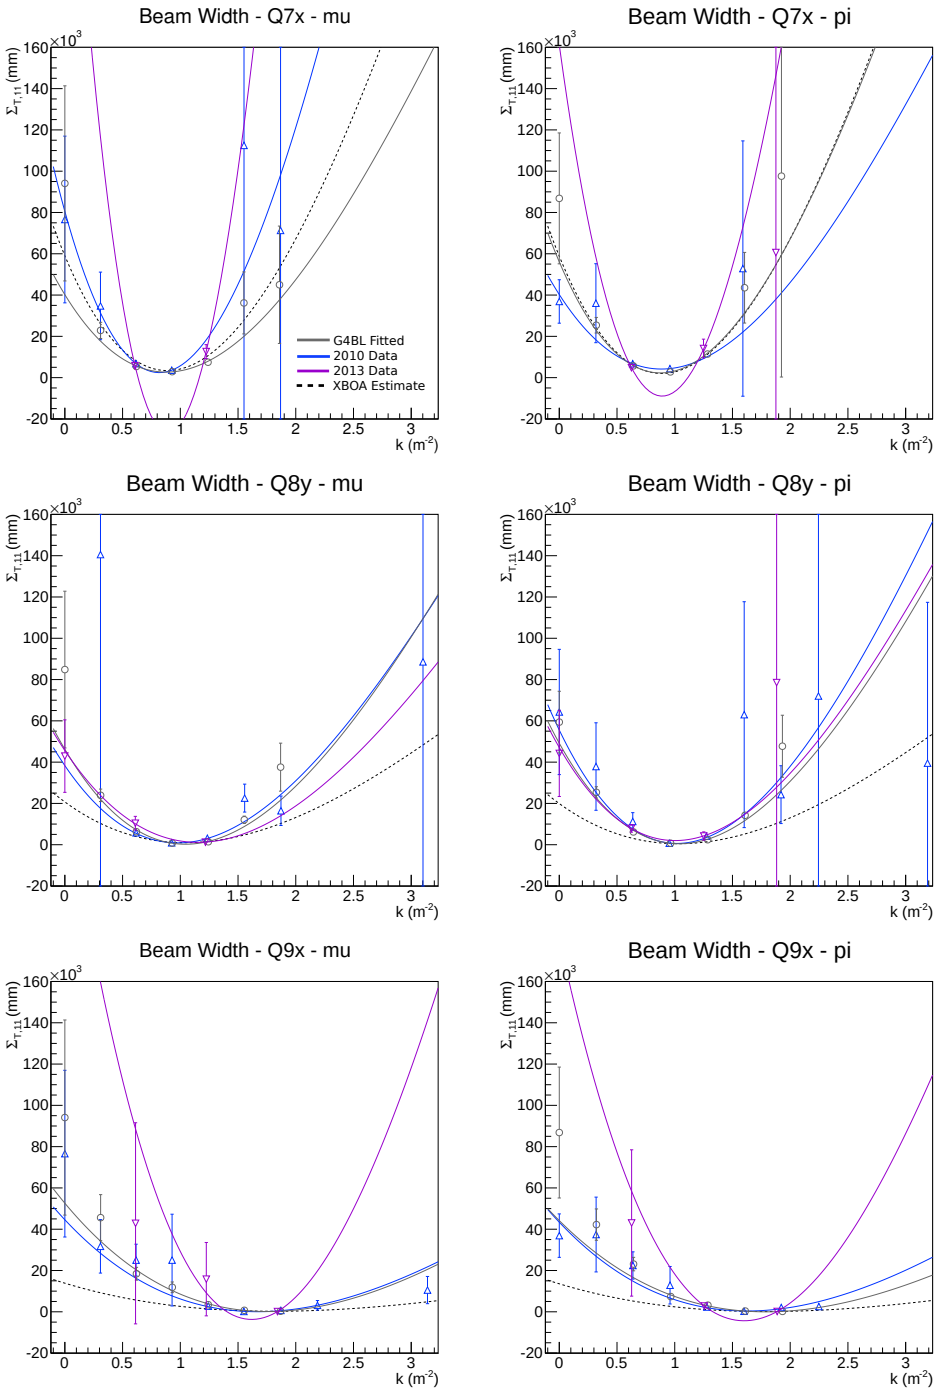

Figure 5.11: Fits to the beam width, in order to determine the covariance matrix just upstream from each quadrupole.

			Monte Carlo		Data	
Parameter			XBOA	Scan Fit	2010	2013
Q7x	Σ_{11}	μ	13181	9100 ± 1100	21800 ± 9800	80000 ± 25000
		π	13454	13100 ± 1100	8500 ± 3300	39000 ± 22000
	Σ_{12}	μ	3556	2390 ± 280	5000 ± 2000	23600 ± 7900
		π	3671	3460 ± 280	2240 ± 700	11300 ± 7200
	Σ_{22}	μ	1294	887 ± 73	1380 ± 310	4290 ± 120
		π	1199	1138 ± 69	990 ± 130	680 ± 130
	ϵ_N	μ	4.4	3.2 ± 1.9	5 ± 18	*
		π	2.6	2.6 ± 1.8	2.8 ± 3.0	*
	χ^2/n.d.f.	μ		5.0/3	0.2/2	0.3/1
		π		2.3/3	2.6/2	1.1/1
Q8y	Σ_{11}	μ	8123	20000 ± 1100	17400 ± 3300	14000 ± 11000
		π	8162	20000 ± 1100	23300 ± 5500	20000 ± 12000
	Σ_{12}	μ	1660	2310 ± 130	2990 ± 720	4000 ± 1500
		π	1646	4200 ± 270	4600 ± 1500	3700 ± 1900
	Σ_{22}	μ	542	581 ± 37	690 ± 160	1380 ± 600
		π	434	979 ± 61	1050 ± 350	1090 ± 360
	ϵ_N	μ	2.7	2.5 ± 1.1	4 ± 5	4 ± 13
		π	1.4	2.1 ± 2.4	3 ± 12	5 ± 7
	χ^2/n.d.f	μ		9.5/3	5.4/5	1.5/1
		π		9.8/3	4.7/5	0.6/1
Q9x	Σ_{11}	μ	8259	34200 ± 4800	31400 ± 5700	-7000 ± 10000
		π	8374	27600 ± 2200	32300 ± 5300	143000 ± 72000
	Σ_{12}	μ	1520	5460 ± 580	3970 ± 940	7700 ± 2100
		π	1475	4930 ± 300	3200 ± 1000	13900 ± 5800
	Σ_{22}	μ	485	960 ± 150	530 ± 140	4270 ± 160
		π	357	865 ± 89	600 ± 250	-1840 ± 160
	ϵ_N	μ	2.7	4 ± 7	2 ± 14	*
		π	1.4	*	5 ± 4	*
	χ^2/n.d.f	μ		3/3	6/5	0.4/1
		π		11/3	7.1/5	0.4/1

Table 5.5: The results of the covariance matrix fits shown in figure 5.11. Each Σ_{ij} is the i,jth element of the covariance matrix just upstream from the indicated quadrupole. A * indicates that the determinant of the covariance matrix was non-physical. The units are as follows: $\Sigma_{11}(\text{mm})^2$, $\Sigma_{12}(\text{mm})(\text{mrad})$, $\Sigma_{22}(\text{mrad})^2$

edges of which were scraped in the quadrupole. In the alignment study the goal was to observe the un-clipped beam and the increasing Σ_{11} element at each successive magnet is encouraging.

The scan showed reasonable agreement between the data points plotted from all three datasets, thus indicating consistency between the Monte Carlo and collected data. Both datasets from 2010 and 2013 also overlay well and indicate a reasonable level of beamline stability between the collection periods.

5.6 Conclusion

The aim of this chapter was to study both the alignment and covariance properties of the Step I beamline, utilising a quadrupole scan. Neither the detectors nor the beamline were designed to undertake this type of measurement. A substantial effort was placed into adapting the detectors to make a good parametrisation of the beam and selecting the optimal fit range to avoid scraping effects. The result of this was a large uncertainty on the beam parametrisation, which in turn passed the uncertainty into both the alignment and covariance studies.

The alignment study was able to achieve a resolution of $\approx$ 5 mm in some cases, which is a considerable achievement considering both the scale and uncertainty of the input data. The scan itself indicated a possible misalignment of Q9 in x, further analysis indicated this could be the result of a misalignment between the surveyed beam axis and particle beam axis.

The covariance study was able to demonstrate that the simulated beam matched both sets of data collected during the quadrupole scans. However substantial error on the covariance matrix terms prevented any meaningful measurement of the emittance from the scan.

Overall the scan enabled both a study of alignment and covariance, effectively using a single time of flight station. The scan itself is currently statistically limited and more data would increase the precision of the results. However, the data collection for each quadrupole requires a significant time($\approx$ 5 h), making substantial improvement implausible realistically. That said, the addition of the fibre tracker will allow a full reconstruction of each particle both upstream and downstream of the cooling cell. Therefore it may be possible to exploit this in a scan of the cooling cell to obtain a much higher precision measurement in substantially less time.

Chapter 6

Conclusions

The first chapter of this thesis has given an overview of the history of the neutrino and described the solar neutrino problem, which led to the discovery of neutrino oscillation. The current neutrino oscillation experiments and the current global best fit results have been shown, which identifies the tasks remaining for future neutrino experiments. The measured value of $\sin(\theta_{13})$ has shown that a number of next generation neutrino beams will be capable of measuring the CP violating phase δ. Although it is no longer a necessity, the Neutrino Factory still maintains a position of delivering the most precise measurement of δ. This precision comes at a substantial cost, and requires the construction of an advanced accelerator facility using a number of technologies, not all of which have been demonstrated. The muon cooling technology is one example, which is to be demonstrated by the Muon Ionisation Cooling Experiment.

Chapter two of this thesis introduced the concept of emittance and described the ionisation cooling process which is to be demonstrated by MICE. A brief overview of the MICE cooling cell was given, followed by the description of the MICE experiment, which measures each particle which traverses the cooling cell. The spectrometer system used in MICE aims to deliver an emittance measurement with a absolute precision of 0.1%.

Presently the first step of the MICE experiment has been completed, a number of preliminary measurements have also been made using the TOF system, KL and EMR detectors. MICE is eagerly awaiting the construction of step IV of the experiment, which will greatly improve the diagnostics available and begin the ionisation cooling study.

6.1 Target Diagnostics

In order to maintain a successful data taking campaign, MICE is dependent on the operation of a reliable target mechanism. Chapter three of this thesis described in detail the target system and how an actuation is accomplished. This was then followed by the development of three diagnostic techniques:

- The calculation of acceleration using the controller data has a resolution of $\approx 3\ \mathrm{m\,s^{-2}}$ (of $\approx 700\ \mathrm{m\,s^{-2}}$), however there is substantial variation with driving voltage, temperature and starting position. By placing cuts on the allowed starting position, and running at constant settings these effects can be neutralised. Ideally a temperature/position calibration could be implemented, which would decrease the variability substantially.

  The use of this parameter for monitoring changes in the mechanical performance of the mechanism is limited, since only a change in acceleration can be observed. This means that no absolute value for friction can be determined, which prevents comparison between stators. Finally, targets installed in ISIS operate intermittently and with changing dip depths which further confuse this parameter.

- The width of the starting position has been used to aid in monitoring friction in the target and is less susceptible to environmental effects than the acceleration calculation. However the restoring force at the end of the actuation is not precisely known, and can also vary between different controller configurations. This means that only a value which is proportional to the actual friction can be obtained. In addition this approach is limited by a maximum width (3.0 mm in two phase mode), which limits the maximum friction that can be observed.

  In practice the controller configuration is fairly consistent, which allows this parameter to add meaningful data to the mechanical diagnostics.

- Using the recorded trajectory of the target, an involved fitting technique has been developed which is capable of measuring the friction in each actuation. The fit showed a reasonable agreement with a numerical model, demonstrating that an absolute friction could be measured to $\approx 7\ \mathrm{m\,s^{-2}}$, or 0.3 N. When this technique was applied to real target data many changes observed in the friction measurement were also visible in the previous two techniques. This technique determines an absolute value, so it can be used to meaningfully compare friction

and other parameters between targets. It may also be possible to use this measurement to help calibrate the start position width.

These diagnostic techniques were then applied to a number of real targets, which had a number of step changes in acceleration due to temperature/voltage effects. On the whole the targets demonstrated an increase in friction with running time. This has allowed some general constraints to be placed on the performance of the mechanism after approximately one days running, which would have identified two targets that were unable to achieve one million actuations.

The next step for the QA process is to assure that a target will attain five million actuations on ISIS, which is not yet guaranteed. Unfortunately the R78 tests are often stopped after a target number of actuations is achieved, which makes the running lifetime ill-determined.

6.2 Tracker Readiness

The MICE scintillating fibre trackers are an essential part of the ionisation cooling experiment and have been described in chapter four. In this thesis the ISIS-1RF reference signal has been studied as a potential method to synchronise the fibre trackers to muon arrival. The particles observed in MICE were well correlated with the ISIS-1RF signal and a synchronisation system was designed and built. This system was tested using one tenth of the tracker in the single station test, where it operated well. As part of the test a number of diagnostics were described, which verify the timing of the various gates required for successful operation. A number of minor issues were revealed during the test, which are addressed by the addition of a dedicated buffer circuit. The remaining component of the synchronisation system is the EPICS interface, which will bring the delay units fully into the MICE ecosystem.

In addition the single station test demonstrated the need for an effective light injection system. This need was addressed by the construction of an LED calibration system, which was also detailed in chapter four of this thesis. In order to operate within the 4 T field of the tracker all electronic components are located outside of the high field region and light is routed using polymer optical fibre. The light injection system operates a minimum of 20 LED pulsers, each capable of adjustable brightness and a fixed < 3.5 ns RMS pulse width. At present the individual components of the system have been tested, however the combined system remains untested. Finally the control software to set the bias voltages must be written, which should be trivial

using the VME bus.

6.3 Beam Alignment

The final chapter of this thesis has investigated the use of a beam based alignment measurement of the final quadrupole triplet and time of flight detector. Data for the scan was collected in 2010 and again in 2013, with simulations made to mimic the data taking runs. The total number of particles collected by the scan at each point restricted the resolution to ≈ 5 mm. Since this limit is statistical, extended running could aid in improving the resolution, however a single magnet scan required ≈ 5 h, implying that a further improvement would require considerable beam time and is not recommended.

The addition of the trackers to the beamline adds an interesting possibility, which is that each particle could be reconstructed before a scan of a cooling channel magnet using the first tracker. The second tracker could then provide a second measurement of the particles. By using a suitable fitting technique, it may be possible to attain an alignment measurement of the cooling channel from a scan. In order to accomplish this the single particles would require tracking through the numerous fields, with multiple scattering also present. Further studies are needed to assess if this would be feasible.

Appendix A

Estimation of coil and magnetic efficiency during actuation

During the actuation of the target mechanism the control system switches through different phases in order to maintain maximum acceleration. The switching of the coils is not instantaneous; the coil stack, capacitor bank and interconnecting cables each contribute to the rise time of the circuit. While the precise source of this is not of concern, the general effect is. A further consideration is the effect of coil-magnet positioning during the switching of the coils.

The goal of this appendix is to estimate the average effect of both the coil rise time and the coil-magnet positioning, resulting in a function which describes $\overline{C(v)}$, to be used in chapter 3.

A.0.1 Coil Rise (t_{S}) time

The coils form a series resistor inductor circuit, of which we are interested in the impulse response as the circuit is switched on and then off. Response of RL circuits is an exponential rise towards an equilibrium current. In this case it can be expressed as the following:

$$I(t) = \begin{cases} 0, & \text{if } t < 0 \\ 1 - \exp\left(\frac{-t}{t_{\mathrm{S}}}\right), & \text{if } 0 \leq t \leq t_{\mathrm{on}} \\ I(t_{\mathrm{on}}) \exp\left(\frac{-t+t_{\mathrm{on}}}{t_{\mathrm{S}}}\right), & \text{otherwise} \end{cases} \tag{A.1}$$

where t_{S} is the coil rise time and the circuit is switched on at time $t = 0$, remains on for t_{on} and is then switched off at $t = t_{\mathrm{on}}$.

A more convenient parameterisation to use is dimensionless, where the coils are switched on at $x = 0$ and off at $x = 1$. In this case:

$$x = t/t_{\text{on}} \tag{A.2}$$

which allows $I(t)$ to be rewritten $I(x)$:

$$I(x, a) = \begin{cases} 0, & \text{if } x < 0 \\ 1 - \exp(\frac{-x}{a}), & \text{if } 0 \leq x \leq 1 \\ I(1)\exp(\frac{-(x-1)}{a}), & \text{otherwise} \end{cases} \tag{A.3}$$

and by comparison to equation A.1:

$$a = t_{\text{S}}/t_{\text{on}} \tag{A.4}$$

Next, the value of t_{on} can be calculated from the coil switching pitch ($p = 3$ mm) and shaft velocity (v):

$$t_{\text{on}} = p/v \tag{A.5}$$

$$a(v) = \frac{vt_S}{p} \tag{A.6}$$

Note that a has been constructed such that it collects all the physical parameters and increases with increasing velocity.

A.0.2 Coil-Magnet Positioning

In addition to rise time effects, the force from the coils moves though the peak value of a cosine curve, as the shaft changes position. This effect is also notable and worth calculating.

In total there are 6 states which the are switched in sequence during actuation. In an optimal configuration the coils switch such that the phase(ϕ) of the cosine function ($\cos(\phi)$) moves from $\phi = -\pi/6$ to $\phi = +\pi/6$. This can be expressed consistently with the rise time equation, where at $x = 0, \phi = -\pi/6$ and $x = 1, \phi = +\pi/6$, specifically:

$$J(x) = \cos\left(\frac{(2x-1)\pi}{6}\right) \tag{A.7}$$

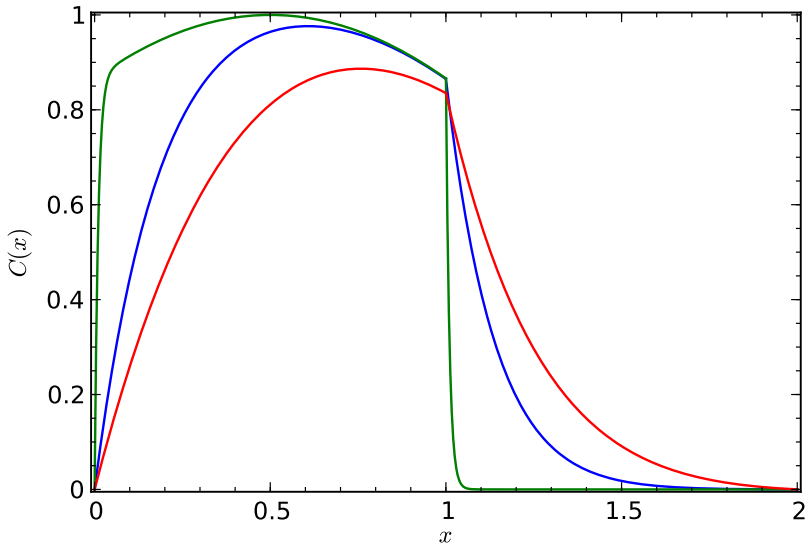

Figure A.1: Effect of coil rise time and sinusoidal response to position, for a number of different values of a. Grn: $a = 0.01$; Blu: $a = 0.15$; Red: $a = 0.30$

A.0.3 Overall fractional force

The overall fractional force delivered to the shaft is a product of the current and coil-magnet positioning:

$$C(x, a) = I(x, a)J(x) \tag{A.8}$$

and is visualised in figure A.1 which demonstrates the effect of increased coil switching times. The green line indicates when the coil rise time(t_{S}) is much shorter than the switching time (t_{on}). Conversely, the red line indicates the case when the coil rise time(t_{S}) has become a significant fraction of the switching time (t_{on}); note that this curve contains less area, indicating on average a reduced force.

To fully explore the effect of increasing velocity it is necessary to measure the average value of $C(x, a)$ for different a. This is trivially accomplished with the following integral:

$$\overline{C(a)} = \int_0^\infty C(x, a)dx \tag{A.9}$$

The result of the integral is shown in figure A.2. Note that when the velocity is zero, the target is capable of delivering on average 95% of the peak acceleration. As the target is accelerated, the velocity increases to a peak value around 5 $\mathrm{m\,s^{-1}}$, at which point $a \approx 0.25$ and only 89% of the peak force can be delivered.

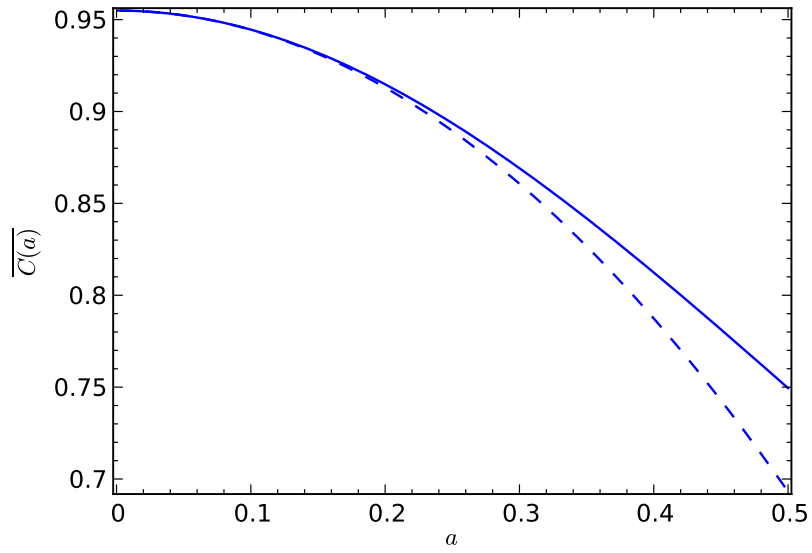

Figure A.2: The effect of increasing a on the average force delivered to the target shaft. The dashed line indicates the result of a Taylor expansion to second order about $a = 0$.

The exact expression for the integral was evaluated using the code given in figure A.3. This then yielded the result:

$$\begin{aligned}\overline{C(a)} =& \frac{3}{2}\left(\frac{\pi a^2}{\pi^2 a^2+9} - \frac{3\sqrt{3}a}{\pi^2 a^2+9}\right)\left(e^{\left(-\frac{1}{a}\right)}-1\right) \\ & - \frac{3\left(\pi^2 a^2 - 3\sqrt{3}\pi a - (\pi^2 a^2+9)e^{\frac{1}{a}}\right)e^{\left(-\frac{1}{a}\right)}}{2\left(9\,\pi+\pi^3 a^2\right)} \\ & - \frac{9\left(\sqrt{3}\pi a - 3\right)}{2\left(9\,\pi+\pi^3 a^2\right)} \qquad \text{(A.10)} \\ \overline{C(a)} \approx& \frac{3}{\pi} - \frac{1}{3}\,\pi a^2 \qquad \text{(A.11)}\end{aligned}$$

The approximation is the result of a Taylor expansion to second order about $a = 0$ and is used in the target analysis fitting routines to estimate loss of acceleration due to increased velocity.

```
a = var('a') # The fraction of risetime/switchtime
x = var('x') # Time axis
Ia(x) = 1 - exp(-x/a)
Ib(x) = Ia(x=1)*exp((-x+1)/a)
F(x) = cos(pi/6*(2*x - 1))
Ca(x) = F*Ia
Cb(x) = F*Ib
assume(a>0)
C_avg = Ca.integral(x,0,1) + Cb.integral(x,1,+Infinity)
```

Figure A.3: Source code to evaluate integral in the SAGE[88] mathematics library.

This page is intentionally left blank.

References

[1] Wolfgang Pauli. Letter to Tübingen conference participants, December 1930.

[2] C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, and A.D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.

[3] Raymond Davis. Attempt to detect the antineutrinos from a nuclear reactor by the $Cl^{37} + \overline{\nu} \rightarrow e^{-} + A^{37}$ reaction. *Phys. Rev.*, 97:766–769, Feb 1955.

[4] G. Danby, J-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, Jul 1962.

[5] DONUT Collaboration. Observation of tau neutrino interactions. *Physics Letters B*, 504:218–224, April 2001, hep-ex/0012035.

[6] J. Beringer et al. Review of particle physics. *Phys. Rev. D*, 86:010001, Jul 2012.

[7] T. D. Lee and C. N. Yang. Question of parity conservation in weak interactions. *Phys. Rev.*, 104:254–258, Oct 1956.

[8] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1415, Feb 1957.

[9] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. *Phys. Rev.*, 109:1015–1017, Feb 1958.

[10] Bruce T. Cleveland et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *The Astrophysical Journal*, 496(1):505526, Mar 1998.

[11] K. S. Hirata et al. Results from one thousand days of real-time, directional solar-neutrino data. *Phys. Rev. Lett.*, 65:1297–1300, Sep 1990.

[12] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, Jun 2002.

[13] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, Oxford, 2007.

[14] Y. Ashie et al. Evidence for an oscillatory signature in atmospheric neutrino oscillations. *Phys. Rev. Lett.*, 93:101801, Sep 2004.

[15] F. P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, Apr 2012.

[16] Y. Abe et al. Reactor electron antineutrino disappearance in the Double Chooz experiment. *Phys.Rev.*, D86:052008, 2012, 1207.6632.

[17] J. K. Ahn et al. Observation of reactor electron antineutrinos disappearance in the RENO experiment. *Phys. Rev. Lett.*, 108:191802, May 2012.

[18] K. Abe et al. Observation of electron neutrino appearance in a muon neutrino beam. *Phys. Rev. Lett.*, 112:061802, Feb 2014.

[19] P. Adamson et al. A Study of Muon Neutrino Disappearance Using the Fermilab Main Injector Neutrino Beam. *Phys.Rev.*, D77:072002, 2008, 0711.0769.

[20] R.B. Patterson. The NOνA experiment: status and outlook. *Nuclear Physics B - Proceedings Supplements*, 235236(0):151 – 157, 2013. The XXV International Conference on Neutrino Physics and Astrophysics.

[21] N. Agafonova et al. Evidence for $\nu_\mu \to \nu_\tau$ appearance in the CNGS neutrino beam with the OPERA experiment. *Phys. Rev. D*, 89:051102, Mar 2014.

[22] S. Choubey et al. International Design Study for the Neutrino Factory, Reference Design Report, Draft 1. `https://www.ids-nf.org/wiki/FrontPage/Documentation/RDR?action=AttachFile&do=get&target=IDS-NF-RDR-v1.0.pdf`, 2013.

[23] S. Choubey et al. International Design Study for the Neutrino Factory, Interim Design Report. 2011, 1112.2853.

[24] K.T. McDonald, H.G. Kirk, H. Park, T. Tsang, I. Efthymiopoulos, et al. The MERIT High-Power Target Experiment at the CERN PS. In *11th European Particle Accelerator Conference (EPAC06)*, 2009.

[25] S. Geer. Muon Colliders and Neutrino Factories. In *25th International Linear Accelerator Conference (LINAC10)*, September 2010, 1202.2140.

[26] Klaus Wille. *The Physics of Particle Accelerators: An Introduction*. Oxford University Press, USA, 2001.

[27] Helmut Wiedemann. *Particle Accelerator Physics*. Springer, 2007.

[28] G. Penn. Beam envelope equations in a solenoidal field. *Muon Collider Note*, 71, 2000, `http://nfmcc-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=71;filename=muc0071.pdf;version=1`.

[29] D. Adams et al. Characterisation of the muon beams for the Muon Ionisation Cooling Experiment. *Eur.Phys.J.*, C73(10):2582, 2013, 1306.1509.

[30] David Neuffer. Principles and Applications of Muon Cooling. *Part.Accel.*, 14:75, 1983.

[31] David Neuffer. Multi-TeV muon colliders. *AIP Conference Proceedings*, 156(1):201–208, 1987.

[32] S. Ozaki, R. Palmer, M. Zisman, and J. Gallardo. Feasibility Study-II of a Muon-Based Neutrino Source. *BNL-52623*, 2001, `www.cap.bnl.gov/mumu/studyii/FS2-report.html`.

[33] Mice technical reference document - draft version. `http://www.isis.rl.ac.uk/accelerator/mice/TR/MICE_Tech_Ref.html`, 2005.

[34] T. Carlisle and J. Cobb. Ionization Cooling in MICE Step IV. *IPAC2011*, C110904:877–879, 2011.

[35] D. Li et al. 201MHz Cavity R&D for MUCOOL and MICE. *EPAC 2006 Proceedings*, pages 1367–1369, 2006.

[36] A. Moretti et al. Effect of High Solenoidal Magnetic Fields on Breakdown Voltages of High Vacuum 805 MHz Cavities. *LINAC 2004 proceedings*, page 272, 2004.

[37] A. Moss. MICE RF power system specifications. *MICE Note*, 372, 2012, `http://mice.iit.edu/mnp/MICE0372.pdf`.

[38] K. Ronald. Muon Transit RF Phase Determination. *MICE Note*, 433, 2013, `http://mice.iit.edu/micenotes/restricted/pdf/MICE0433/MICE0433.pdf`.

[39] M. Hills and A. Nichols. MICE Essential Beamline Geometry. *MICE Note*, 242, 2009, `http://mice.iit.edu/mnp/MICE0242.pdf`.

[40] M. Bogomilov et al. The MICE Muon Beam on ISIS and the beam-line instrumentation of the Muon Ionization Cooling Experiment. *JINST*, 7(05):P05009, 2012.

[41] L. Cremaldi, D. A. Sanders, P. Sonnek, D. J. Summers, and J. Reidy. A Cherenkov Radiation Detector With High Density Aerogels. *IEEE Transactions on Nuclear Science*, 56:1475–1478, 2009, 0905.3411.

[42] Y. Karadzhov, M. Bonesini, J. S. Graulich, and R. Tsenov. TOF Detectors Time Calibration. *MICE Note*, 251, 2009, `http://mice.iit.edu/mnp/MICE0251.pdf`.

[43] M. Bonesini, R. Bertoni, A. de Bari, and G. Cecchet. The Refurbishing of MICE TOF0 and TOF1 detectors. *MICE Note*, 363, 2012, `http://mice.iit.edu/mnp/MICE0363.pdf`.

[44] R. Bertoni, M. Bonesini, A. deBari, G. Cecchet, Y. Karadzhov, and R. Mazza. The construction of the MICE TOF2 detector. *MICE Note*, 286, 2012, `http://mice.iit.edu/mnp/MICE0286.pdf`.

[45] R. Bertoni et al. Analysis of PID detectors (TOF and KL) performances in the MICE 2010 run. *MICE Note*, 337, 2010, `http://mice.iit.edu/mnp/MICE0337.pdf`.

[46] M. Rayner. *The development of a novel technique for characterizing the MICE muon beam and demonstrating its suitability for a muon cooling measurement.* PhD thesis, University of Oxford, 2011.

[47] A. Khan et al. MICE Scintillating Fibre Tracker Prototype. *MICE Note*, 90, 2005, `http://mice.iit.edu/micenotes/public/pdf/MICE0090/MICE0090.pdf`.

[48] M. Adinolfi, F. Ambrosino, A. Antonelli, M. Antonelli, F. Anulli, et al. The KLOE electromagnetic calorimeter. *Nucl.Instrum.Meth.*, A482:364–386, 2002.

[49] M. Bogomilov, M. Bonesini, O. M. Hansen, Y. Karadzhov, D. Orestano, and L. Tortora. Measurement of the pion contamination in the MICE beam. *MICE Note*, 416, 2013, `http://mice.iit.edu/mnp/MICE0416.pdf`.

[50] R. Asfandiyarov. A totally active scintillator calorimeter for the Muon Ionization Cooling Experiment (MICE). Design and construction. *Nucl.Instrum.Meth.*, A732:451–456, December 2013.

[51] P. Barrillon et al. Maroc: Multi-anode readout chip for mapmts. In *Nuclear Science Symposium Conference Record, 2006. IEEE*, volume 2, pages 809–814, Oct 2006.

[52] D. Bolognini et al. Frontend and readout electronics of the MICE Electron Muon Ranger detector. In *Real Time Conference (RT), 2010 17th IEEE-NPSS*, pages 1–5, May 2010.

[53] A. Dobbs, D. Forrest, and F. J. P. Soler. The MICE luminosity monitor. *Journal of Physics: Conference Series*, 408(1):012084, 2013.

[54] Matthias Clausen and Leo Dalesio. EPICS: Experimental physics and industrial control system. *ICFA Beam Dyn.Newslett.*, 47:56–66, 2008.

[55] W. Grandegger. DATE V1 - Data Aquisition Test Environment Users Manual. 1996. CERN-ALICE-INT-1996-16.

[56] C. D. Tunnell and C. T. Rogers. MAUS: MICE Analysis User Software. In *IPAC'11*, pages 850–852, 2011.

[57] T. J. Roberts and D. M. Kaplan. G4beamline simulation program for matter-dominated beamlines. In *Particle Accelerator Conference, 2007. PAC. IEEE*, pages 3468–3470, June 2007.

[58] C. Booth et al. The design, construction and performance of the MICE target. *JINST*, 9:P03006, 2013, 1211.6343.

[59] P. Smith, C. Booth, and P. Hodgson. Magnetic Modelling of the MICE Target Mechanism. *MICE Note*, 309, 2010, `http://mice.iit.edu/micenotes/public/pdf/MICE0309/MICE0309.pdf`.

[60] DuPont. *Vespel ®SCP-5000 polyimide isostatic shapes typical ISO properties.*

[61] The SPiDeR collaboration webpage. `https://heplnm061.pp.rl.ac.uk/display/spider/Home/`.

[62] Pieter Samyn, Jan Quintelier, Gustaaf Schoukens, and Patrick De Baets. Sliding properties of polyimide against various steel and dlc-coated counterfaces. In B Ivkovic, editor, *Proceedings of BALKANTRIB 05*, pages 506–515, 2005.

[63] B. Wang et al. The Design and Construction of the MICE Spectrometer Solenoids. *IEEE Transactions on Applied Superconductivity*, 19(3):1348–1351, Jun 2009.

[64] S. Virostek, D. Li, H. Pan, S. Prestemon, and R. Preece. Assembely and Test of a Modified Spectrometer Solenoid for MICE. *Proceedings of IPAC13*, 2013.

[65] The MICE Tracker Group. Specification of the scintillating and clear fibre for the MICE scintillating fibre trackers. *MICE Note*, 135, 2006, `http://mice.iit.edu/mnp/MICE0135.pdf`.

[66] M. Ellis et al. The design, construction and performance of the MICE scintillating fibre trackers. *Nucl.Instrum.Meth.*, A659:136–153, December 2011, 1005.3491.

[67] A. Bross, E. Flattum, D. Lincoln, S. Grnendahl, 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.*, A477(13):172 – 178, 2002.

[68] P. Rubinov. *AFEIIt Readout Format*, November 2007. DØ Note 5520.

[69] D. Adey. *Beam instrumentation and investigations into muon cooling at MICE.* PhD thesis, University of Warwick, 2012.

[70] T560 4-channel compact digital delay and pulse generator. Highland Technology data sheet `http://www.highlandtechnology.com/DSS/T560DS.shtml`.

[71] M. Takahashi et al. Illumination System for the MICE Tracker Station Assembly QA. *MICE Note*, 167, 2007, `http://mice.iit.edu/mnp/MICE0167.pdf`.

[72] Gigahertz-Optik. ODP97 - Barium Sulfate Coating. Technical Data Sheet.

[73] J. E. McMillan. Kraton paint mix ratios. Personal communication.

[74] Fibre Data. *Polymer Fibre Optic Cable, FDPF 4001 EH.* Technical Data Sheet.

[75] J.S. Kapustinsky, R.M. DeVries, N.J. DiGiacomo, W.E. Sondheim, J.W. Sunier, and H. Coombes. A fast timing light pulser for scintillation detectors. *Nucl.Instrum.Meth.*, A241(23):612 – 613, 1985.

[76] J. E. McMillan. Using the Sheffield Pulser. Personal communication.

[77] Bivar. *UV5TZ-405-15.* Technical Data Sheet.

[78] Hytek. *IP-DAC-8402.* Technical Data Sheet.

[79] Acromag. *AVME9630.* Technical Data Sheet.

[80] K. T. McDonald and D. P. Russell. Methods of emittance measurement. In M. Month and S. Turner, editors, *Frontiers of Particle Beams; Observation, Diagnosis and Correction*, volume 343 of *Lecture Notes in Physics, Berlin Springer Verlag*, pages 122–132, 1989.

[81] M. A. Rayner. Presentation of latest results. Presentation at the 30th collaboration meeting of the MICE experiment, July 2011, Oxford., 2011.

[82] O. M. Hansen, A. Blondel, and I. Efthymiopoulos. Towards a Symmetric Momentum Distribution in the Muon Ionisation Cooling Experiment. *Conf. Proc.*, C130512(CERN-ACC-2013-0281):TUPFI020, May 2013.

[83] C. Rogers. XBOA. Cross Platform Beam Optics Analysis `https://launchpad.net/xboa`.

[84] C. Croxford, M. Hills, A. Mariani, J. Giles, and A. Austin. Survey of MICE Beamline. *MICE Note*, 216, 2013, `http://mice.iit.edu/mnp/MICE0418.pdf`.

[85] T. Hayler and J. Palin. Survey of the MICE Hall 8th January 2013. *MICE Note*, 418, 2013, `http://mice.iit.edu/mnp/MICE0418.pdf`.

[86] M. Apollonio. MICE Materials Budget. Spreadhseet `http://www.mice.iit.edu/bl/Documentation/PION-BL-MAGIC-TABLE-v170809.xls`.

[87] D. A. Forrest. *The Muon Ionisation Cooling Experiment.* PhD thesis, University of Glasgow, 2011.

[88] W. A. Stein et al. *Sage Mathematics Software (Version 6.1.1).* The Sage Development Team, 2014. `http://www.sagemath.org`.