

Towards a Muon Collider

Studies from KURNS and MICE



A thesis submitted for the degree of
Doctor of Philosophy by

Craig Brown

Department of Electronic and Computer Engineering
Brunel University London

17 November 2023

Abstract

The proposed Muon Collider offers two distinct advantages in comparison to lepton and hadron colliders. The larger mass of the muon in comparison to the electron means synchrotron radiation is less of a concern and higher centre of mass collision energies can be reached. Secondly, the muon collider would collide point-like particles in comparison to hadron colliders, resulting in cleaner collision processes.

Challenges faced by the Muon Collider are investigated in this thesis. A charge exchange experiment was performed at the Kyoto University Institute for Integrated Radiation and Nuclear Science to investigate electron detachment cross-sections as a function of projectile energy. Results are presented for 11 MeV hydrogen anions striking a carbon foil.

Ionisation cooling, an increase in the position and momentum phase-space density of a beam, was investigated by the Muon Ionisation Cooling Experiment (MICE) for various absorber materials. The systematic uncertainties associated with the liquid hydrogen absorber are presented. The ionisation cooling result can be affected by various biases. The MICE momentum reconstruction was biased by the non-uniformity of the magnetic field in the tracker regions of the MICE experiment. It was also biased by misalignments of the solenoid, tracker and magnetic axes.

The ionisation cooling result can also be biased by transmission losses. Normalizing the phase space densities by their sample sizes as MICE had done was found to be incorrect. When transmission losses are missing not at random, the changing covariance matrix of the remaining distribution needs to be accounted for as well. Using a transfer matrix approach, a correction procedure was outlined to impute missing data points for the downstream distribution affected by transmission losses. The correct downstream covariance matrix could then be found, meaning the remaining downstream sample could be compared to the full upstream sample unaffected by any biases due to transmission losses.

Contents

Preamble	3
Abstract	3
Table of Contents	4
Acknowledgements	7
1 Introduction	9
1.1 Towards a muon collider	11
1.2 Basic description of the muon collider	12
1.3 Basic description of MICE	19
1.4 Transfer matrices	21
1.5 MICE reconstruction	25
2 11 MeV charge exchange at KURNS	31
2.1 Electron detachment cross-sections	34
2.2 Foil stripping	39
2.3 KURNS facility	40
2.4 Experimental design	41
2.5 Foil preparation	43
2.6 Experimental run plan	46
2.7 Experimental results	54
2.8 Calculating the partial stripping factor F_{PS1}	58
2.9 Discussion	62
3 Systematic effects of the MICE IH_2 absorber	67
3.1 Absorber vessel shape	68
3.2 Absorber vessel window thicknesses	69
3.3 Variation of the density of liquid hydrogen	70
3.4 Total systematic uncertainty on energy loss	77
4 Uncertainty of the MICE magnetic field	79
4.1 Dimensions of the spectrometer solenoid coils	79
4.2 Centre position of the match coils	81

4.3	Centre of the magnetic axis	84
4.4	Deformation of the solenoid coils	85
4.5	Tuning the magnetic field	86
4.6	Misalignments	90
4.7	Discussion	93
5	Momentum reconstruction at MICE	95
5.1	Linear solenoid transfer matrix assumptions	98
5.2	Biases of the MICE reconstruction	99
5.3	Quantifying the momentum bias	105
5.4	Time reconstruction	125
6	Transfer matrices	129
6.1	Transfer matrix between two stations	131
6.2	Transfer matrix between two trackers	136
7	Emittance, amplitude and density	143
7.1	Three types of cooling measurement	144
7.2	Missing data	150
7.3	Emittance, amplitude and density plots	155
7.4	Unbiased cooling	169
8	Conclusions	177
	Bibliography	181

Acknowledgements

I was in Holland Park when I received the email offering me the opportunity to study for a PhD at Brunel. The offer took me so much by surprise that I literally started skipping. How often in life do you literally start skipping?

I first and foremost would like to thank my supervisor Paul Kyberd for giving me the opportunity to study for a PhD at Brunel. Whenever I did ask Paul for help, he has been a valuable guiding force. This is especially true for the Muon Ionization Cooling Experiment (MICE) where it can be easy to be caught up in frivolous pedantry. Paul helped me look at the bigger picture. Why am I performing this particular piece of analysis or experiment, and what do I expect it to show? My co-supervisor for the project was Chris Rogers. If we ever agreed on anything, I am not sure, but in the end I still managed to produce this thesis. While I do acknowledge the help of the MICE collaboration, I do want to thank Paolo Franchini in particular. I hope his car survives to this day.

I would also like to acknowledge the help and guidance I received from the Intense Beams Group at the Rutherford Appleton Laboratory. This is the place where I spent most of my PhD, though I wouldn't recommend living in Didcot. The Intense Beam Group meetings can offer valuable constructive criticism on the work that is being presented, though due to the length of some of the meetings, one should be prepared to skip lunch. I would like to thank Shinji Machida in particular for not only funding the two trips to the Kyoto University Institute for Integrated Radiation and Nuclear Science (KURNS), but also for his insight and guidance into the charge exchange experiments performed at KURNS. This of course means I would also like to thank David Kelliher and Jean-Baptiste Lagrange from the Intense Beams Group, as well as Yoshihiro Ishi, Tomonori Uesugi and Hidefumi Okita from KURNS for helping me perform the charge exchange experiments at KURNS. No experiment is as easy as initially envisioned.

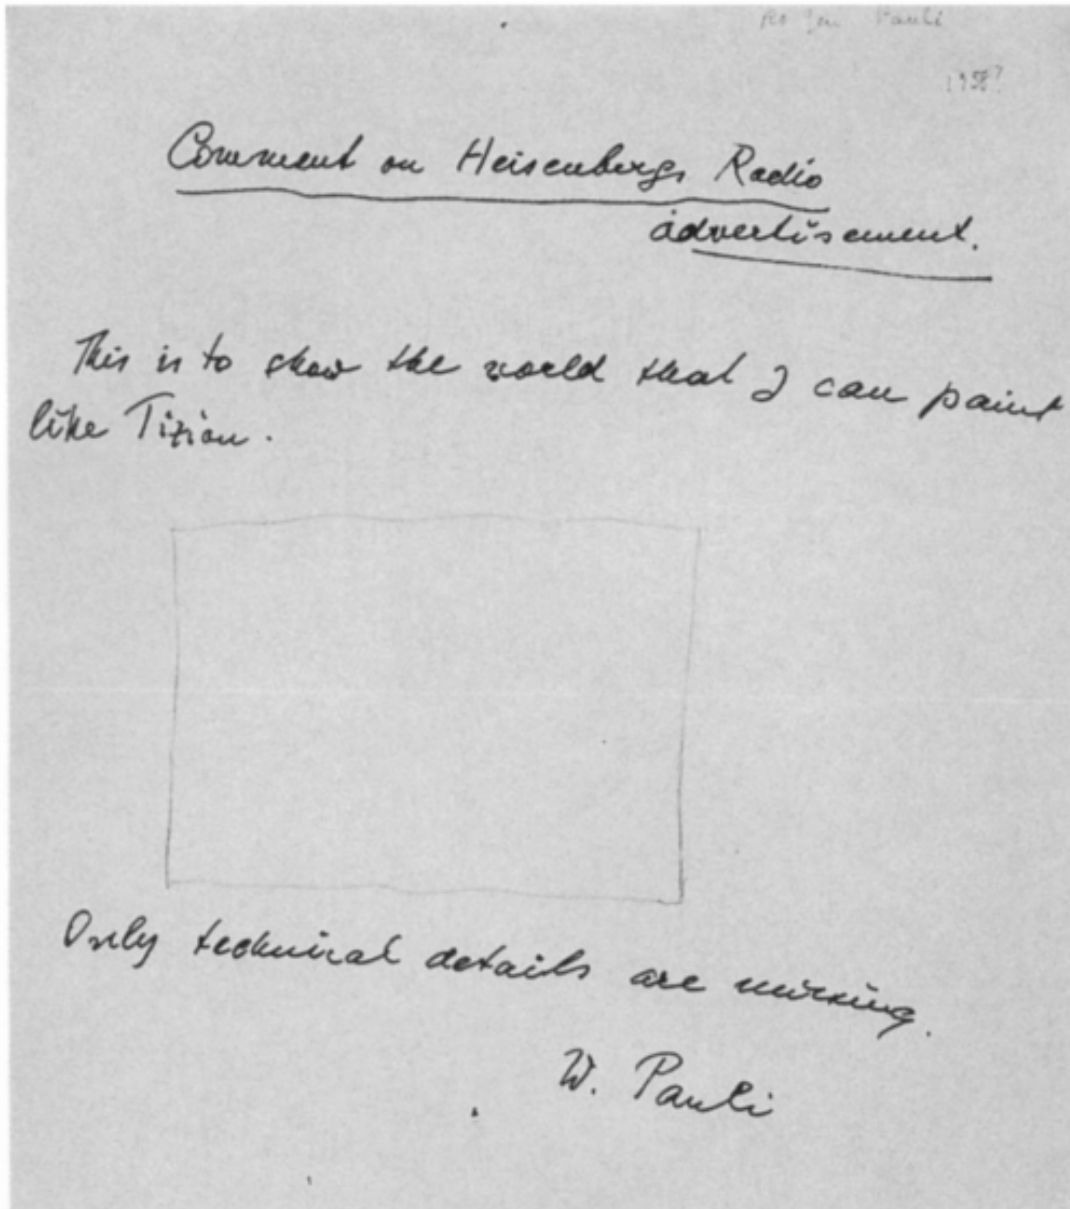


Figure 1: This is to show the world that I can paint like Titian. Only technical details are missing. Wolfgang Pauli¹

Introduction

The Large Electron-Positron (LEP) collider led to the precision measurements of the Z boson and W boson, while the Large Hadron Collider (LHC) led to the discovery of the Higgs boson. The LHC is expected to continue to run until the 2040s. Already plans are being proposed for the next generation of lepton and hadron colliders, dubbed 'Higgs Factories', such as the Future Circular Collider (FCC), the Compact Linear Collider (CLIC), the International Linear Collider (ILC), the Circular Electron Positron Collider (CEPC) and the Muon Collider, which will have greater centre of mass energies and luminosities, that would allow for greater precision measurements of the Higgs mass and width, as well as the Higgs' decay processes. Unlike previous colliders, the proposed colliders would not be searching for any new particles in particular, but rather looking at what happens when we do collide particles at higher energies and if there is new physics to be found, the main motivation being that the standard model is an incomplete description of particle physics.

The FCC could work in three different configurations, colliding protons with protons, electrons with electrons (or positrons) or protons with electrons (or positrons). The FCC does have some natural drawbacks. The collision of electrons and/or positrons is limited by synchrotron radiation. Protons, not being point-like particles, result in messy collision processes due to the quarks and gluons within them. The muon collider provides an alternative, colliding heavier point-like particles far less affected by synchrotron radiation. The muon collider has its own challenges of course, especially due to the muons short lifetime.

This thesis will examine some of the problems the muon collider currently faces and the progress that is being made towards making a muon collider a reality. This opening chapter will begin by giving a brief outline of the muon collider and the Muon Ionization Cooling Experiment (MICE).

The proposed muon collider is expected to use an 8 TeV proton beam. The proton beam can be created from a H^- beam during charge exchange injection. When the H^- beam passes through a carbon foil, it is stripped of its electrons, creating a proton beam. The efficiency of this process depends on the injection energy and the foil thickness used. The required carbon foil thickness is expected to scale with injection energy. However, experimental results have only been taken over a narrow energy range (181 – 800 MeV). To verify the relationship between the required foil thickness and the injection energy across a larger energy range, low energy (11 MeV) experimental results were taken at the Kyoto University Institute for Integrated Radiation and Nuclear Science (KURNS). The results are presented in chapter 2.

The remaining chapters will consider the MICE experiment. Chapter 3 will consider some of the systematic uncertainties associated with the liquid hydrogen absorber at MICE. It will also determine the density of the liquid hydrogen used during the MICE experiment.

The position of each particle can be measured at the ten tracker stations within the MICE cooling channel. The MICE reconstruction software uses these position measurements to reconstruct the momentum of each particle at the ten tracker stations. However, the MICE reconstruction software depends on the magnetic field within the MICE cooling channel. The uncertainty on the magnetic field is investigated in chapter 4.

How these uncertainties affect the momentum reconstruction will be considered in chapter 5. It will find that the MICE reconstruction software introduces small biases for the transverse momentum reconstruction and large biases for the longitudinal momentum reconstruction. The magnitude of the biases will be different for each particle, as the biases are also dependent on the trajectory of the particle in the MICE cooling channel.

Ionization cooling and emittance exchange can be demonstrated by changes in the emittance, amplitude distribution or phase-space density distribution of the beam. However, these measurements can be biased by transmission losses, as each type of measurement depends on the covariance matrix of the distribution. To overcome this issue, a transfer matrix approach is introduced in chapter 6. It considers how the phase-space volume of the initial upstream sample of particles would have changed if there were no transmission losses. The resulting covariance matrix can then be used to calculate the phase-space density distribution of the remaining downstream particles. As the resulting covariance matrix has accounted for transmission losses, the upstream sample of particles and the remaining downstream sample of particles can be directly compared. This will result in an unbiased demonstration of the ionizing cooling process for a limited fraction of the beam, before transmission losses begin to obscure the result. The results are presented in chapter 7.

Some of the results have been published elsewhere. The analysis on the liquid hydrogen absorber was presented in a MICE note² and a Journal of Instrumentation article³. Progress on the study of emittance exchange at MICE was presented at the Neutrino 2020 and ICHEP 2020 conferences^{4,5}. Previous MICE collaboration results⁶ presented at IPAC 2019 should be ignored, due to flaws in the analysis technique not discovered until later. Finally, the MICE collaboration has recently published papers on the emittance measurement and the cooling demonstration at MICE in the European Physical Journal C⁷ and Nature⁸ respectively.

1.1 Towards a muon collider

The main challenge for a muon collider is to achieve the required luminosity. The luminosity of a beam can be defined as:

$$\mathcal{L} = f_{coll} \frac{n_1 n_2}{4\pi\sigma_x\sigma_y} \mathcal{F} \quad (1.1)$$

where n_1 and n_2 describe the number of particles in each bunch colliding with frequency f_{coll} , σ_x and σ_y are the transverse RMS beam sizes⁹, while the factor \mathcal{F} accounts for geometric effects such as the crossing angle of the beams and the bunch length (though $\mathcal{F} = 1$ will be assumed). For Gaussian position and momentum distributions, the transverse RMS beam size σ can be explained in terms of emittance, a measure of the spread of the beam in position and momentum space through

$$\sigma^2 = \epsilon\beta \quad \epsilon = \frac{\sqrt[n]{|V|}}{mc} \quad (1.2)$$

where ϵ is the emittance of the beam, β is the beta function describing how focused the beam is at the interaction point, V is the covariance matrix of the n -dimensional position and momentum space considered, m is the mass of the beam particle and c is the speed of light. Taking $\mathcal{F} = 1$, the luminosity can then be described by:

$$\mathcal{L} = f_{coll} \frac{n_1 n_2}{4\pi\sqrt{\epsilon_x\beta_x\epsilon_y\beta_y}} \quad (1.3)$$

An experiment's luminosity can thus be increased by increasing the number of particles colliding, increasing the frequency of collisions, reducing the emittance of the beam or by reducing the beta function at the interaction point e.g. by focusing the beam with higher strength magnetic fields.

For the proton driver based muon collider (Fig. 1.1), the muons are produced as tertiary particles (protons striking a target producing pions which decay to muons). This production process results in the muons having a large spread in position and

momentum space (or phase-space) and thus this muon collider's main challenge will be to reduce the emittance of the muon beam within the muon's lifetime. Alternative muon collider schemes^{10–12} exist, though they typically suffer from low muon production cross-sections.

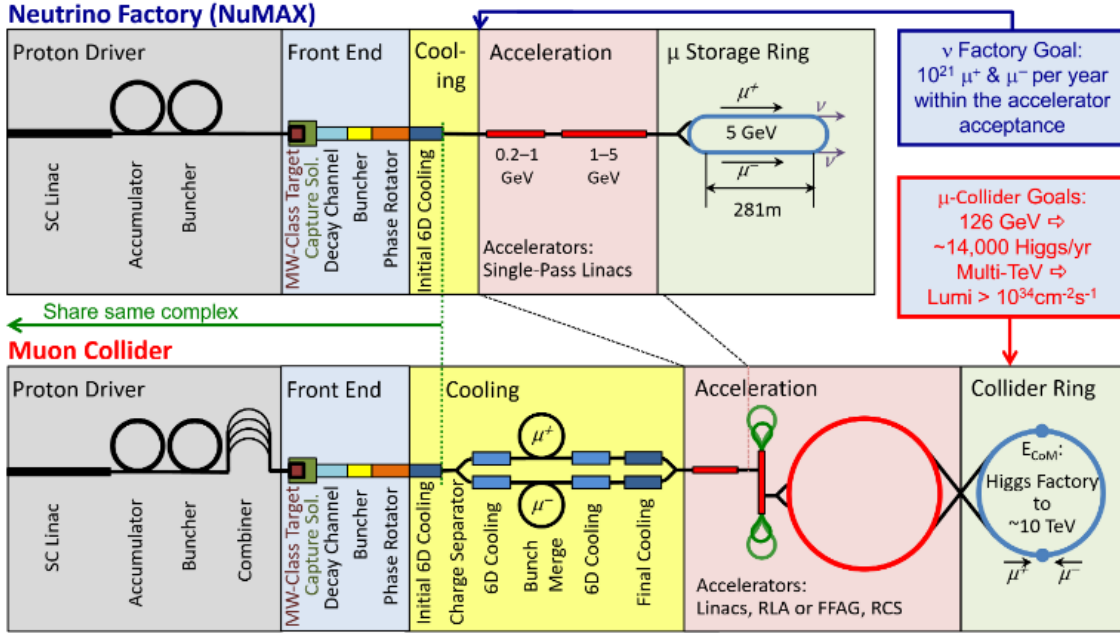


Figure 1.1: A basic description of the neutrino factory and muon collider. The neutrino factory, as well as nuSTORM, could be considered as stepping stones towards a muon collider, incrementally implementing more complex physics designs and technologies¹³.

The proton driver based muon collider^{14–17} can be implemented in incremental steps (nuSTORM → Neutrino Factory → Higgs Factory → Multi-Tev Muon Collider). This structure would provide both a neutrino and a Higgs physics program where each additional step could be added as the relevant physics processes and technologies are demonstrated, culminating in a multi-TeV muon collider. A brief outline of the muon collider is given next.

1.2 Basic description of the muon collider

1.2.1 Proton driver

To deliver enough pions and muons, the proton driver^{13,18} needs to provide a proton beam of sufficient power (in the megawatt range). Providing proton beams of sufficient energy or beam current is already possible. However, the main challenge is to deliver both at the same time¹⁹.

While the proton driver could directly produce protons and inject them into an

accumulator ring, the proton beam is typically derived from a H^- beam. This makes injection into an accumulator ring easier. As the H^- particles and the protons will follow different trajectories in a magnetic field, this can be exploited during charge exchange injection (see chapter 2). The H^- particles in the beam are stripped of both of their electrons via a thin stripping foil, with bump magnets used to join the new protons with the existing proton beam in the accumulator ring. The proton beam current thus depends on how many H^- particles can be produced, delivered and stripped into the accumulator ring.

The H^- beam is typically produced via a surface plasma source where neutral hydrogen atoms impact a surface and can pick up electrons²⁰⁻²³. The binding energy (or affinity) of the electron to the hydrogen atom is low at 0.7542 eV, and thus a low work function surface is used (typically caesium) to increase the likelihood of the hydrogen atom picking up an electron during impact. Magnets are then used to guide the H^- beam to a superconducting linear accelerator, where a series of oscillating electric potentials are used to accelerate and chop the beam to match the predefined frequency of the RF cavities in the accumulator ring.

When the H^- beam is injected into the accumulator ring, it is stripped of its electrons, with the protons accumulating in the ring. Once a sufficient number of protons have accumulated in the ring, the protons are collected into bunches and accelerated via RF cavities. This results in a high energy and high intensity proton beam that can be extracted to a target area.

1.2.2 Front end

In the target area, the proton beam strikes a target, creating pions. The number of pions produced depends on the target material and the proton beam's energy. For proton beams of megawatt power, the main consideration is the survival of the target itself due to the high heat loads imparted. Experiments have been completed of proton beams impacting on a liquid mercury jet target by MERIT²⁴ and on solid targets such as beryllium, carbon, aluminium, copper, tin, tantalum and lead by HARP²⁵⁻²⁷. Due to mercury's hazardous nature²⁸, solid graphite and fluidised tungsten powder are considered more likely options for a muon collider²⁹⁻³¹.

When the proton beam strikes a target material, the pions produced will scatter at large angles and at various momenta. To capture the pions, it has been proposed to use a high field strength solenoid³² (Fig. 1.2). The captured pions are then led to a decay channel consisting of a series of magnets and drift spaces, where the pions decay to muons and any remaining pions are separated from the beam. Finally, the muons are phase rotated in a series of RF cavities to decelerate high energy particles and accelerate low energy particles until a homogenous energy beam is obtained³³.




Figure 1.2: When the proton beam strikes a target (in this case mercury) the resulting pions will scatter and must be captured in a solenoid³⁴.

1.2.3 Cooling

Going into the cooling section, the beam has a low energy spread but a large transverse emittance. This can be reduced by cooling the beam, a reduction in the spread of the beam in position and momentum space. The most common cooling techniques include radiation^{35–37}, laser^{36–38}, electron^{39–43}, stochastic^{38,44,45}, ionisation^{8,46} and frictional cooling^{47–51}. However, most of these techniques do not sufficiently cool the muon beam within the muon’s lifetime. For the muon collider, it has been proposed to use ionisation cooling.

In ionisation cooling^{8,46}, a beam is passed through an absorber material, where the beam loses energy and momentum. The beam is then accelerated by passing it through an RF cavity, which recovers the lost longitudinal momentum of the beam but not the transverse momentum. This reduces the overall emittance of the beam.



Figure 1.3: Ionisation cooling⁵². The overall momentum of the beam is reduced by an absorber material before the longitudinal momentum of the beam is restored by a RF cavity, reducing the overall emittance of the beam.

Scattering and energy loss can be reduced by using a low atomic number absorber, as can be seen from the Bethe-Bloch equation⁵³, which describes the mean energy loss of moderately relativistic heavy charged particles,

$$-\left\langle \frac{dE}{dx} \right\rangle = 4\pi N_A r_e^2 m_e c^2 z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (1.4)$$

where N_A is Avogadro's constant, r_e and m_e are the classical electron radius and mass, z refers to the charge of the incident particle, Z and A are the atomic and mass numbers of the material, β and γ are the Twiss parameters, T_{\max} is the maximum kinetic energy that can be imparted to a free electron in a single collision, while I is the mean excitation energy. The $\delta(\beta\gamma)$ term, is a density effect correction term due to ionisation energy loss.

The emittance of a beam changes when passing through an absorber material and is described by the ionisation cooling equation⁸:

$$\frac{d\varepsilon_{\perp}}{ds} \cong -\frac{\varepsilon_{\perp}}{\beta^2 E_{\mu}} \left\langle \frac{dE}{ds} \right\rangle + \frac{\beta_{\perp} (13.6 \text{ MeV})^2}{2\beta^3 E_{\mu} m_{\mu} X_0}. \quad (1.5)$$

where $\frac{d\varepsilon_{\perp}}{ds}$ is the rate of change of normalized transverse emittance (ε_{\perp}). β , E_{μ} and m_{μ} are the muon velocity, energy and mass, $\frac{dE}{ds}$ is the magnitude of the mean energy loss rate through ionisation, X_0 is the absorber radiation length and β_{\perp} is the transverse beta function at the absorber⁵⁴.

In Eq. 1.5, the cooling effect is described by the negative term which reduces the emittance of the beam, while the positive term is the heating effect. When the cooling term is larger than the heating term, the emittance of a beam can be reduced. If that differential is large enough, then the beam can be cooled quickly. As the cooling term has a dependence on the atomic number of the absorber material, it is advantageous to use a low atomic number material such as liquid hydrogen to cool the beam quickly.

While ionisation cooling reduces the transverse components of the beam, Emittance exchange^{5,55,56} can be used to additionally reduce the longitudinal components of the beam. In Emittance exchange, a muon beam is passed through a dipole magnet to create both a position spread and a position-energy correlation in the beam. The beam is then passed through a wedge shaped absorber of specific thickness to eliminate the momentum dispersion. As a result, the beam has a lower longitudinal emittance but a greater transverse emittance. The transverse emittance of the beam can then be reduced through ionisation cooling (Fig. 1.4).

In reverse emittance exchange, the beam is first passed through a wedged shaped absorber and then through a dipole magnet. This reduces the transverse emittance of the beam at the expense of a larger longitudinal emittance.

A muon collider will require significant 6D cooling to achieve the necessary luminosity, i.e. a reduction in the spread of both the transverse and the longitudinal components of the beam (Fig. 1.5).




Figure 1.4: *Left:* Emittance exchange reduces the longitudinal emittance at the expense of an increased transverse emittance. The transverse emittance can then be reduced via ionisation cooling. *Right:* Reverse emittance exchange reduces the transverse emittance at the expense of an increased longitudinal emittance.

1.2.4 Acceleration

Muon beams are most efficiently cooled at low momenta in the MAP proposal. When the muon beam emerges from the cooling section, it needs to be rapidly accelerated to overcome the muon's short lifetime^{33,57,58}. This can be done by using high-gradient RF cavities, though they carry a significant financial cost. It is therefore advantageous for the beam to pass multiple times through the same RF cavities. To increase the muon's energy as rapidly and efficiently as possible, the accelerator facility design plans to use several different accelerators. These include the use of linacs, Recirculating Linear Accelerators (RLA), Fixed Field Alternating Gradient machines (FFA) and Rapid Cycling Synchrotrons (RCS)³³.

The initial linac rapidly accelerates the cooled muon beam, before the muon beam makes a number of passes through the RLA (Fig. 1.6). The rapid acceleration of the linac makes the beam more relativistic and extends the muon's lifetime. It also reduces the beam's phase-space volume, making the RLA viable. Due to the muon's short lifetime, the RF cavities will need to operate with as high a gradient as possible. This will increase the energy spread of the beam. As the arcs of the RLA only have a limited energy acceptance, this will lead to beam losses³³.

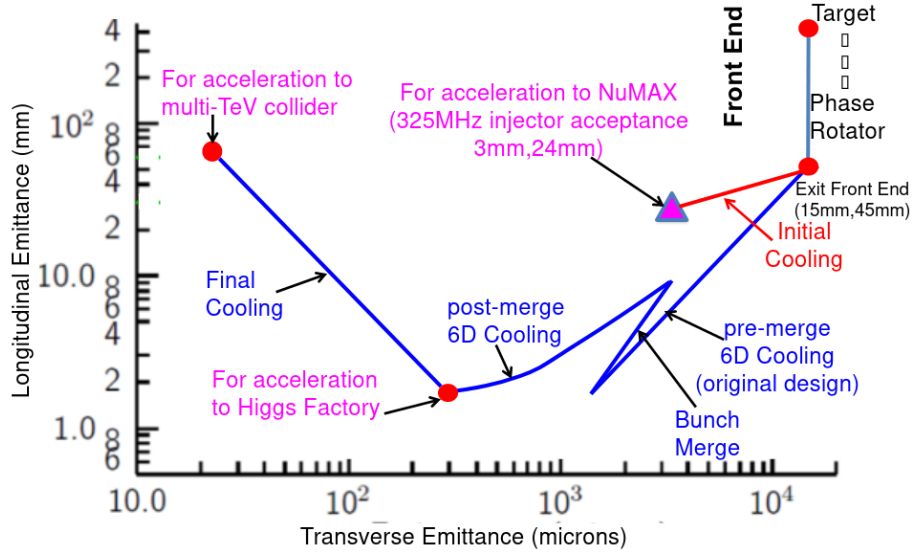


Figure 1.5: The muon beam on production occupies a large emittance. This must be reduced to achieve the luminosity required for a muon collider. The graph shows the emittance of the beam after each stage, with cooling used to achieve the desired transverse and longitudinal emittances¹³.

The design and size of the RLA limit its feasible energy reach. To continue accelerating a muon beam to TeV energies, RCSs and FFA machines have been proposed^{34,59}. The RCS has a larger energy acceptance compared to the RLA, but it faces its own difficulties, such as requiring rapidly ramping magnets to follow the increase in beam energy. The RF cavity gradients also need to be far larger than for typical synchrotrons, leading to a limited longitudinal acceptance of the beam.

At higher beam energies⁶⁰, FFA's have also been proposed as an alternative to RLA's. At higher beam energies, the beam makes more regular passes through the RF cavity of the FFA and becomes more efficient than the RLA. In an FFA, the magnetic field becomes stronger as the beam increases in energy⁶¹. This is

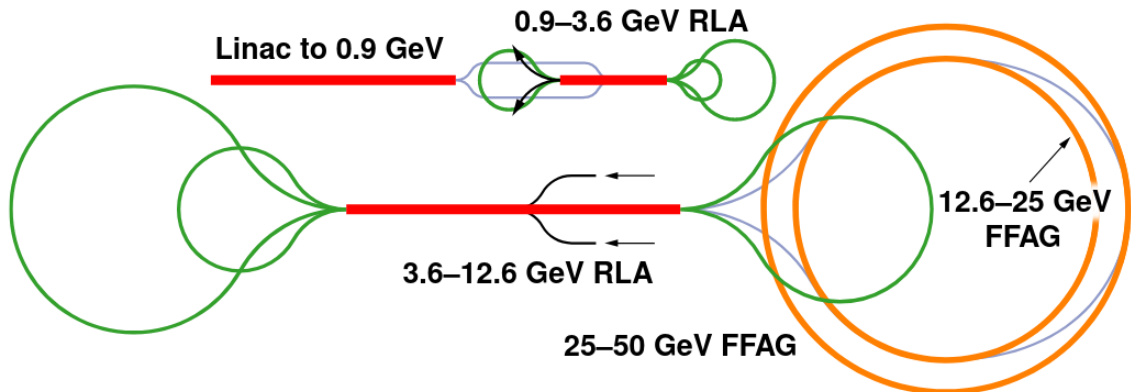


Figure 1.6: Various stages of an acceleration system design that could be used in a neutrino factory or muon collider³³.

achieved by the beam moving to a higher radius orbit as it accelerates, with the field getting continuously stronger at higher radii. EMMA has demonstrated a variation of this design, called the linear non-scaling FFA, where the orbit doesn't scale with momentum⁶². This allows for a more compact design, though the magnets become increasingly more complex.

A further improvement called the vertical FFA has been proposed to eliminate time-of-flight variations of the beam⁶³⁻⁶⁵. Instead of going to a higher radius orbit as the beam energy increases, the beam orbit would change with height as the beam energy increases. As the orbital path length remains fixed, relativistic beams would then travel at a fixed revolution frequency, i.e. isochronously.

The acceleration section of the muon collider requires significant optimisation⁵⁷. Different accelerator types will be optimum in terms of acceptance and muon survival at different energies. This means that each centre-of-mass collision energy (e.g. 0.125 TeV, 3 TeV, 14 TeV) will have a different optimal acceleration scenario in terms of maximising the number of muons in the collider ring. For example, at lower collider energies it may be more efficient to accelerate the beam more slowly to increase the number of accepted muons, even at the expense of more decays e.g. by using a lower RF frequency in a RCS/FFA.

1.2.5 Collider ring

The muon collider proposes to use a traditional ring. At higher energies it becomes more difficult to focus the beam at the interaction point as the beta function is inversely proportional to the Lorentz factor. This increases the required strength of the final focusing magnets at the interaction point to achieve the desired luminosity⁶⁶.

While radiation considerations from muon decays play a significant factor throughout the muon collider complex, they will be most severe in the collider ring due to the high energies involved and the muons nearing the end of their lifetime. This will require the use of significant shielding and perhaps the use of open mid plane magnets to prevent the decay products putting excessive power loads into the surrounding magnets. This applies to the diagnostic systems near the interaction point as well, where background reactions and decays need to be minimised⁶⁷.

1.3 Basic description of MICE

The Muon Ionization Cooling Experiment was a proof of concept experiment to demonstrate ionisation cooling. The experiment recreated a section of the cooling channel design for a potential neutrino factory or muon collider. Based at the Rutherford Appleton Laboratory, it used the ISIS proton beam to strike a titanium target. The pions produced were captured using a quadrupole triplet and led to a dipole-solenoid-dipole combination where pions decayed to muons. The strength of the dipole field was used to select muons of the desired momentum. Another two quadrupole triplets focused the beam into the cooling section of the experiment. Here the beam passed through the upstream spectrometer solenoid, an absorber material and the downstream spectrometer solenoid (Fig. 1.8).

MICE used a series of detectors to distinguish and track the particles passing through the cooling channel^{3,68,69}. Two Time-of-flight (TOF) detectors were placed before the cooling section, while another TOF was placed after the cooling section. The TOF system consisted of scintillator slabs and photomultiplier tubes which could distinguish muons from electrons and pions, based on their differing velocities⁷⁰. Before entering the TOF system, the dipole bend had selected particles of equal momentum. This meant lighter particles would travel faster between the first two TOF stations than the heavier particles (Fig. 1.7). At higher momenta, it becomes more difficult to distinguish between the flight time differences of the particles, and therefore Cherenkov counters can be used instead. Finally, the Kloe-Light and the Electron-Muon-Ranger form a combined electromagnetic calorimeter after the ionisation cooling section. This can be used to separate decay electrons from muons⁷¹.




Figure 1.7: The time-of-flight differences between TOF0 and TOF1 can be used to distinguish muons (middle peak) from electrons (left peak) and pions (right peak)³.




Figure 1.8: A schematic of the MICE experiment⁸. Muons could be distinguished from pions and electrons based on their flight time differences between two TOF stations. The particles could also be tracked through ten tracker stations. The position measurements at the ten tracker stations were used to reconstruct the momentum of the particles in the MICE cooling channel.

The MICE experiment operated at a low rate, allowing for each individual particle passing through the cooling channel to be measured. The MICE beam was then assembled from all the individual particle measurements. The position and momentum of each particle was measured using two scintillating fibre trackers housed within multi-Tesla spectrometers solenoids either side of an absorber material, with each tracker containing five stations of three planes of scintillating fibres⁷². The stations were able to identify the position of each particle, while the longitudinal spacing and the helical trajectory of each particle was used to deduce the momentum of each particle. The reconstruction software will be further explored in section 1.5.

The position and momentum measurements of the beam particles were used to calculate the overall emittance of the beam, before and after an absorber material. In MICE, a number of different absorber materials were investigated. A lithium hydride disk and an aluminium vessel filled with liquid hydrogen were used to investigate the change in transverse emittance of a muon beam, while a polyethylene wedge was used to investigate emittance exchange, a change in both the transverse and longitudinal emittances of the beam.

In Eq. 1.5, the change in transverse emittance depended on the beam's initial energy and input emittance. For this reason MICE recorded data at various beam energies (140, 170, 200 and 240 MeV/c) and multiple input emittances (3, 4, 6 and 10 mm). For the MICE cooling channel, the equilibrium emittance was expected to be approximately 4 mm, i.e. when the heating and cooling terms of Eq. 1.5 balanced. Ionisation cooling would therefore be expected to be seen at input emittances greater than 4 mm (for the liquid hydrogen or lithium hydride absorbers).

The various nominal input emittances were created using a diffuser containing two brass and two tungsten irises. When one or more of the irises were closed, the beam passing through the diffuser was heated. The heating of the beam can be similarly described by the ionisation cooling equation (Eq. 1.5). Therefore, depending on which combination of irises were closed, the input beam could be heated to a desired input emittance, before the beam was passed through the cooling channel.

In addition to taking data at the MICE experiment, the MICE collaboration has written the MICE Analysis User Software (MAUS) to simulate and analyse the MICE experiment⁷³. The software contains a simulated replica of each physical structure in the MICE experiment and uses Geant4 to simulate the particles travelling through the MICE cooling channel. The MAUS software was used as a tool to help understand and verify the physics processes seen in the MICE experiment. Further information on the MICE experiment can be found in some of the recent MICE theses^{74–77}.

The transfer matrix for a particle travelling through a solenoid is considered in the following section. A transfer matrix can be used to describe the motion of a particle between two longitudinal positions. The predictable behaviour of a particle travelling through a magnetic field is exploited for the momentum reconstruction software (section 1.5), which only uses the physical measurements of the MICE experiment, i.e. the hits in the tracker and TOF stations.




Figure 1.9: The MICE diffuser containing two brass and two tungsten irises. The irises can be used in various configurations to heat the input beam to various input emittances in the cooling channel.

1.4 Transfer matrices

The transformation of a particle from some initial point, $u_{in}(\vec{x}, \vec{p})$ to some later point, $u_{fin}(\vec{x} + d\vec{x}, \vec{p} + d\vec{p})$ of variation ds can be written in terms of a transfer matrix M , such that⁷⁸:

$$u_{fin} = Mu_{in}, \quad du_{fin} = Mdu_{in} \quad (1.6)$$

This transfer matrix M is a combination of all the effects which affect a particle's motion. For example, a particle travelling through a quadrupole triplet can be described by some transfer matrix M which is a combination of the three transfer matrices of each individual quadrupole. If one knows the transfer matrix for an arbitrary magnet element, then one can derive a transfer matrix for a combined lens system consisting of quadrupoles, dipoles, solenoids, etc.

The transfer matrix for a magnet can be derived from the Hamiltonian^{79–85}. For an accelerator system in the Frenet-Serret curvilinear coordinate system⁸⁶, the Hamiltonian, K , is given by:

$$K = - \left(1 + \frac{x}{\rho} \right) \left[p^2 - (p_x - eA_x)^2 - (p_y - eA_y)^2 \right]^{\frac{1}{2}} - eA_s \quad (1.7)$$

where x , y and s denote the transverse and longitudinal coordinates along the centre of the accelerator system with a radius of curvature defined by $\rho(s)$, while the momentum, charge and vector potential are described by p , e and A .

A number of assumptions and approximations can then be made for a particular accelerator system. For MICE, this meant the assumption of a linear and uniform solenoid field within the tracker regions of the MICE experiment. For such a solenoid, the magnetic fields and vector potentials can be given by^{86,87}:

$$B_x = x \sum_{k=0}^{\infty} b_{2k+1} (x^2 + y^2)^k \quad (1.8)$$

$$B_y = y \sum_{k=0}^{\infty} b_{2k+1} (x^2 + y^2)^k \quad (1.9)$$

$$B_s = \sum_{k=0}^{\infty} b_{2k} (x^2 + y^2)^k \quad (1.10)$$

$$A_x = y \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \quad (1.11)$$

$$A_y = -x \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \quad (1.12)$$

$$A_s = 0 \quad (1.13)$$

with

$$b_{2k+1} = \frac{-1}{2(k+1)} b'_{2k} \quad (1.14)$$

$$b_{2k+2} = \frac{1}{2(k+1)} b'_{2k+1} \quad (1.15)$$

Substituting into Eq. 1.7 the above vector potentials, then

$$K = - \left(1 + \frac{x}{\rho} \right) \left[p^2 - \left(p_x - ey \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right)^2 - \left(p_y + ex \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right)^2 \right]^{\frac{1}{2}} - eA_s \quad (1.16)$$

The root can be expanded using

$$\begin{aligned}
 (a+b)^n &= a + na^{n-1}b + \frac{n(n-1)a^{n-2}b^2}{2} + \dots = \sum_{l=0}^n \frac{n!}{l!(n-l)!} a^{n-l}b^l = \\
 &= \frac{n(n-1)\dots(n-l+1)}{l(l-1)\dots 1} = \sum_{l=0}^n \binom{n}{l} a^{n-l}b^l
 \end{aligned} \tag{1.17}$$

Making the substitutions $a = p^2$, $n = 1/2$,

$$b = - \left(p_x - ey \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right)^2 - \left(p_y + ex \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right)^2 \tag{1.18}$$

and dividing the Hamiltonian K , by a reference momentum p_0 to remain canonical, then

$$\begin{aligned}
 \tilde{H} = \frac{K}{p_0} &= - \left(1 + \frac{x}{\rho} \right) \frac{1}{p_0} \sum_{l=0}^{\frac{1}{2}} \frac{\frac{1}{2}!}{l! \left(\frac{1}{2} - l\right)!} (p^2)^{\frac{1}{2}-l} \\
 &\times \left[- \left(p_x - ey \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right)^2 \right. \\
 &\quad \left. - \left(p_y + ex \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right)^2 \right]^l
 \end{aligned} \tag{1.19}$$

For a solenoid with a straight geometry and approximating p_0 by p , then

$$\begin{aligned}
 \tilde{H} &= \sum_{l=0}^{\frac{1}{2}} \frac{\frac{1}{2}!}{l! \left(\frac{1}{2} - l\right)!} p^{-2l} \left[- \left(p_x - ey \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right)^2 \right. \\
 &\quad \left. - \left(p_y + ex \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right)^2 \right]^l
 \end{aligned} \tag{1.20}$$

$$\begin{aligned}
 \tilde{H} &= \sum_{l=0}^{\frac{1}{2}} \frac{\frac{1}{2}!}{l! \left(\frac{1}{2} - l\right)!} p^{-2l} \left[-p_x^2 - p_y^2 + 2e(p_xy - p_yx) \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right. \\
 &\quad \left. - \left(\sum_{k=0}^{\infty} \frac{eb_{2k}}{2(k+1)} (x^2 + y^2)^{k+\frac{1}{2}} \right)^2 \right]^l
 \end{aligned} \tag{1.21}$$

If only linear optic terms up to $l = 1$ are considered (i.e. ignoring higher order terms), then

$$\begin{aligned} \tilde{H} = 1 - \frac{1}{2p^2} \left(-p_x^2 - p_y^2 + 2e(p_xy - p_yx) \sum_{k=0}^{\infty} \frac{b_{2k}}{2(k+1)} (x^2 + y^2)^k \right. \\ \left. - \left(\sum_{k=0}^{\infty} \frac{eb_{2k}}{2(k+1)} (x^2 + y^2)^{k+\frac{1}{2}} \right)^2 \right) \end{aligned} \quad (1.22)$$

Similarly, there are linear optic terms only when $k = 0$, then

$$\tilde{H} = 1 + \frac{1}{2} \left(\left(\frac{p_x}{p} \right)^2 + \left(\frac{p_y}{p} \right)^2 \right) - \frac{eb_0}{2p} \left(\left(\frac{p_x}{p} \right) y - \left(\frac{p_y}{p} \right) x \right) + \frac{1}{2} \left(\frac{eb_0}{2p} \right)^2 (x^2 + y^2) \quad (1.23)$$

Rewriting in terms of our original Hamiltonian K , using $K = \tilde{H}p$, then

$$K = p + \frac{1}{2p} (p_x^2 + p_y^2) - \frac{eb_0}{2p} (p_xy - p_yx) + \frac{e^2 b_0^2}{8p} (x^2 + y^2) \quad (1.24)$$

Let $g = \frac{eb_0}{2}$, then

$$K = p + \frac{1}{2p} (p_x^2 + p_y^2) - \frac{g}{p} (p_xy - p_yx) + \frac{g^2}{2p} (x^2 + y^2) \quad (1.25)$$

Using the Hamiltonian equations of motion

$$x' = \frac{\partial K}{\partial p_x} \quad x'' = \dot{p}'_x = -\frac{\partial K}{\partial x} \quad y' = \frac{\partial K}{\partial p_y} \quad y'' = \dot{p}'_y = -\frac{\partial K}{\partial y} \quad (1.26)$$

then

$$x' = \frac{p_x - gy}{p} \quad x'' = \dot{p}'_x = \frac{-p_y g - xg^2}{p} \quad (1.27)$$

$$y' = \frac{p_y + gx}{p} \quad y'' = \dot{p}'_y = \frac{p_x g - yg^2}{p} \quad (1.28)$$

The transformation from an initial point, $u_{in}(x, y, p_x, p_y)$ to some later point $u_{fin}(x + dx, y + dy, p_x + dp_x, p_y + dp_y)$ of variation ds can be given in terms of a transfer matrix \mathbf{M} , such that⁷⁸:

$$u_{fin} = \mathbf{M}u_{in} \quad du_{fin} = \mathbf{M}du_{in} \quad (1.29)$$

with the condition that⁸⁸:

$$\mathbf{M}^T \mathbf{J} \mathbf{M} = \mathbf{J} \quad (1.30)$$

Then

$$\begin{pmatrix} x + x'ds \\ y + y'ds \\ p_x + p'_x ds \\ p_y + p'_y ds \end{pmatrix} = I_4 \begin{pmatrix} x \\ y \\ p_x \\ p_y \end{pmatrix} + \begin{pmatrix} 0 & -g/p & 1/p & 0 \\ g/p & 0 & 0 & 1/p \\ -g^2/p & 0 & 0 & -g/p \\ 0 & -g^2/p & g/p & 0 \end{pmatrix} ds \begin{pmatrix} x \\ y \\ p_x \\ p_y \end{pmatrix} \quad (1.31)$$

The linear transfer matrix for a particle travelling through a solenoid, described by a variation ds , is then given by:

$$\mathbf{M}(ds) = I_4 + \begin{pmatrix} 0 & -g/p & 1/p & 0 \\ g/p & 0 & 0 & 1/p \\ -g^2/p & 0 & 0 & -g/p \\ 0 & -g^2/p & g/p & 0 \end{pmatrix} \quad (1.32)$$

The predictable nature of a particle travelling through a solenoid field is exploited for the MICE momentum reconstruction software.

1.5 MICE reconstruction

The MICE momentum reconstruction^{74,75,89} can be broadly divided into five stages: digitisation, clustering, space point production, pattern recognition and Kalman filtering. Each stage is considered separately in the following subsections.

1.5.1 Digitisation

The MICE cooling channel contains two trackers, with each tracker containing five stations of three planes. Each plane consists of 212 to 214 channels, with each channel containing seven $350 \mu\text{m}$ thick scintillating fibres arranged as in Fig. 1.10, with the fibres producing scintillating light when ionizing radiation passes through them⁸⁹. The scintillating light is led to visible photon counters, where the signal is digitised using analogue-to-digital converters and converted into a number of photo-electrons using a simple conversion factor. After accounting for noise, if the number of photo-electrons is beyond a threshold value (at least 2 photo-electrons⁷⁵), then the channel number and number of photo-electrons produced are recorded as a digit.




Figure 1.10: *Left:* Each of the MICE tracker stations consists of three planes of scintillating fibres rotated at 120 degrees with respect to one another. *Right:* Each plane consists of 212 to 214 channels, with each channel consisting of seven $350\ \mu\text{m}$ thick scintillating fibres.

1.5.2 Clustering

For every channel that contains a digit, the neighbouring channel is also checked for another digit. If digits are found in neighbouring channels, a two-digit cluster is formed, with the unweighted average channel value used to define the plane co-ordinate. For all other digits, a single-digit cluster is formed.

1.5.3 Space point production

If all three planes of a station have a cluster coinciding in transverse position space, then a “Triplet” space point is formed. After all “Triplet” space points have been found, “Doublet” space points are formed from clusters where only two planes coincided. As each channel is uniquely numbered and positioned, the channel numbers that were used to form a “Triplet” or a “Doublet” can be used to determine the transverse position of the space point at the intersection and centre of those channels.

1.5.4 Pattern recognition

For a uniform solenoid field, the particle’s helical trajectory projected on to the transverse position plane traces out a circle (neglecting energy loss and scattering). A potential candidate track can then be found by performing a circle fit of the spacepoints at the five different tracker stations in the transverse position plane (using a linear least-squares technique). The helical pattern recognition process checks through all the different combinations of spacepoints at the five different stations of a tracker.

For each potential candidate track, the circle centre (x_c, y_c) and the radius, R , can be found (Fig. 1.11). From these, the turning angle, ϕ , and the distance, s , the distance the particle travels along the orbital path in the transverse position plane (the arc-length), can be found between every successive pair of stations of a tracker. Using the distance between the trackers along the longitudinal z -axis, a straight line fit can be made in the (z, s) plane, as the ratio between the arc-length distance, s , and the longitudinal distance between stations should remain constant, for a constant transverse momentum particle in a uniform solenoid field.

For each potential track, a χ^2 fit can be made on the circle fit in the (R, ϕ) plane. Similarly, another χ^2 fit can be made on the straight line fit in the (z, s) plane. Potential tracks are then ranked according to their combined χ^2 fit. The lowest combined χ^2 fit potential track is then accepted as the candidate track. The above procedure is then repeated using spacepoints from only 4 of the 5 stations.

The transverse momentum of the candidate track can then be obtained from the radius of the circle via:

$$R = \frac{p_t}{qB_z} \quad (1.33)$$

where R is the radius, p_t is the transverse momentum, q is the charge of the particle and B_z is the strength of the longitudinal solenoid field. The longitudinal momentum can be determined using the spacing between the five stations and the phase advance (the turning angle) the particle makes along the circular orbit via:

$$\frac{ds}{dz} = \sqrt{1 + (p_t/p_z)^2} \quad (1.34)$$

where s is the arc-length, z is the distance between stations, p_t is the transverse momentum and p_z is the longitudinal momentum.

1.5.5 Kalman filtering

As the scintillating fibres have a finite resolution due to their physical size, a Kalman filter is used to smooth the components of the candidate track. From Hunt, the Kalman Filter^{75,89} “is an optimal, linear estimator that correctly takes into account all errors and correlations of measured data . . . For non-linear systems, the Kalman filter provides only the optimal linear approximation”. The Kalman Filter⁷⁵ contains three stages: a prediction stage, a filtering stage and a smoothing stage.

In the prediction stage, the system is predicted at a future stage (i.e. a subsequent station of the tracker) using the current state of the system.

In the filtering stage, use is made of the actual measurement at the future state of the system to adjust the predicted state based on the weight of the error between the measured and the predicted state.




Figure 1.11: *Left:* The projection of a particle's helical trajectory on to the transverse plane in a uniform solenoid field. *Right:* The vertical and horizontal position of the particle as a function of longitudinal position. Displayed are also the five stations of the tracker. These are used to reconstruct the momentum of a particle⁷⁴.

In the smoothing stage, using the predicted and measured states at the five tracker stations, the Kalman filter is propagated in reverse from the optimal estimate of the final state to adjust each preceding stage. Each stage (the stations of the tracker) now contains an optimal linear estimate for the system.

For MICE, an extended Kalman filter was used. This could account for energy loss and multiple Coulomb scattering. At the first station of the tracker, the initial state vector obtained from the pattern recognition stage is given by $u = (x, p_x, y, p_y, q/p_z)$, the position and momentum of the particle in phase-space co-ordinates, where q is the charge of the particle. The predicted state at a future stage for each phase-space co-ordinate as implemented in MAUS is then given by⁷⁵:

$$x' = x + \frac{p_x}{p_t} R \sin \phi - \frac{p_y}{p_t} R (1 - \cos \phi), \quad p'_x = p_x \cos \phi - p_y \sin \phi, \quad (1.35)$$

$$y' = y + \frac{p_y}{p_t} R \sin \phi + \frac{p_x}{p_t} R (1 - \cos \phi), \quad p'_y = p_y \cos \phi + p_x \sin \phi, \quad (1.36)$$

$$z' = z + \Delta z, \quad p'_z = p_z \quad (1.37)$$

The equations can be rewritten in terms of a propagator matrix F_k (see below), that the initial state vector can be multiplied by. The propagator matrix can then be multiplied by matrices describing the energy loss and scattering of the particle⁷⁵.

In MAUS, the energy loss is described by the mean energy loss of the Bethe-Bloch equation (Eq. 1.4). The momentum loss the particle experiences is then described

by the fractional momentum loss, μ_k , so that

$$\mu_k = \frac{\sqrt{(E_0 - \Delta E)^2 - m^2}}{p_0} \quad (1.38)$$

where E_0 , p_0 and m are the initial energy, initial momentum and mass of the particle. The propagator matrix is then simply updated by multiplying the propagator matrix by the energy loss matrix, P_k , so that $F_k \rightarrow P_k F_k$, where

$$F_k = \begin{pmatrix} 1 & \frac{R}{p_t} \sin \phi & 0 & \frac{R}{p_t} (\cos \phi - 1) & 0 \\ 0 & \cos \phi & 0 & -\sin \phi & 0 \\ 0 & \frac{R}{p_t} (1 - \cos \phi) & 1 & \frac{R}{p_t} \sin \phi & 0 \\ 0 & \sin \phi & 0 & \cos \phi & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad P_k = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \mu_k & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & \mu_k & 0 \\ 0 & 0 & 0 & 0 & 1/\mu_k \end{pmatrix} \quad (1.39)$$

Scattering causes the emittance of a beam to grow. In multiple Coulomb scattering, a particle will undergo many small-angle deflections when passing through a medium. The RMS width for the projected angular distribution can be given by⁵³:

$$\theta_0 = \frac{13.6 \text{ MeV}}{\beta c p} z \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \frac{x z^2}{X_0 \beta^2} \right] \quad (1.40)$$

where p , β , c , and z are the momentum, relativistic beta, speed of light, and charge number of the incident particle, while x/X_0 is the thickness of the scattering medium in radiation lengths. The scattering matrix⁷⁵, Q_k is then given by

$$Q_k = \theta_0^2 J_k S_k J_k^T \quad (1.41)$$

where θ_0 is the RMS scattering angle, S_k is a matrix describing the deflection of the transverse position and momentum components of the particle, and J_k is a Jacobian matrix required to transform the S_k matrix to helical phase-space. The matrices are given by,

$$S_k = \begin{pmatrix} d^2/3 & d/2 & 0 & 0 \\ d/2 & 1 & 0 & 0 \\ 0 & 0 & d^2/3 & d/2 \\ 0 & 0 & d/2 & 1 \end{pmatrix} \quad J_k = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & p_z & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & p_z \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (1.42)$$

where d is the distance between planes. The matrix J_k uses the paraxial approximation, which assumes that the longitudinal momentum is far greater than the transverse momentum.

Some assumptions were made by the MICE momentum reconstruction software. It assumed a constant solenoid field. However, chapter 4 will show that the magnetic field is non-uniform in the tracker regions. It will also show that the magnitude of the non-uniformity is different for the upstream and downstream trackers.

The MICE momentum reconstruction software attempts to make a circle fit using the five spacepoints of a tracker. However, chapter 5 will show that the MICE reconstruction software doesn't account for the changing trajectory of the particle in the spectrometer solenoids whenever it encounters energy loss. Each time the particle hits a tracker station, the particle will deviate from its initial helical trajectory path onto a new lower radius helical trajectory path. This similarly applies to the projected helical path in the transverse position plane. The particle appears to spiral inwards each time the particle hits a tracker station.


How these effects can bias the momentum reconstruction are investigated in chapter 5. It will find small biases for the transverse momentum reconstruction and large biases for the longitudinal momentum reconstruction. These biases can naturally affect the ionisation cooling result.

Chapters 6 and 7 will show that the ionisation cooling result can be severely affected by transmission losses. How to overcome the bias due to transmission losses is investigated in chapter 6 with the help of transfer matrices. Chapter 7 will then present an ionisation cooling result unaffected by transmission losses.


The next chapter however begins with the low projectile energy charge exchange experiment performed at KURNS.

11 MeV charge exchange at KURNS

In charge exchange, two interacting particles transfer one or more electrons between each other, resulting in a change of charge of the two interacting particles, e.g. the collision of a proton with a Helium atom⁹⁰:



The probability of the proton capturing an electron from the helium atom can be given in terms of the cross-section, σ_{10} , where the subscript denotes the change in charge of the proton, going from +1 to 0, after gaining the electron. It can similarly denote the electron detachment cross-section from an ion e.g. for a hydrogen anion, the cross-section σ_{-11} denotes the double electron detachment process:



While both electrons can be stripped simultaneously from a hydrogen anion, it is far more likely to occur as a two-step process, i.e. σ_{-10} followed by σ_{01} :



An electron can be stripped in a number of ways, including by black-body radiation^{91,92}, Lorentz stripping^{93,94}, laser stripping⁹⁵⁻⁹⁷, intrabeam stripping^{98,99}, gas stripping and foil stripping¹⁰⁰⁻¹⁰². Most of these processes degrade the beam quality or can even lead to beam losses. However, some of these stripping processes can also be exploited for charge exchange injection.

Charge exchange injection was first proposed by Luis Alvarez¹⁰³ in 1951, where a H^- beam injected into a synchrotron is passed through a thin foil (e.g. carbon) and stripped of its electrons, creating a proton beam. The differing trajectories of the negative H^- ions and the positive protons in a magnetic field can be exploited to create a high intensity proton beam. This is achieved by stripping the H^- beam in the same path as the circulating proton beam (often with the help of an electrostatic septum and magnetic kickers). Freshly stripped protons are continuously added to the same transverse phase-space as the circulating proton beam, and therefore the circulating proton beam increases in intensity.

The foil thickness, typically measured as a surface density in $\mu g/cm^2$, dictates what fraction of the beam is stripped of its electrons when passing through the foil (Fig. 2.1). The fraction of the H^- beam remaining intact or turning to H^0 or H^+ for a foil with a surface density x , is given by^{104,105}:

$$N^- = e^{-(\sigma_{-10} + \sigma_{-11})x}, \quad (2.5)$$

$$N^0 = \frac{\sigma_{-10}}{\sigma_{-10} + \sigma_{-11} - \sigma_{01}} [e^{-\sigma_{01}x} - e^{-(\sigma_{-10} + \sigma_{-11})x}], \quad (2.6)$$

$$N^+ = 1 - N^- - N^0. \quad (2.7)$$

The ability to progressively increase the intensity of the proton beam, as well as the low beam losses when the stripping efficiency is high, make H^- injection favourable to direct proton injection. Additionally, any unstripped or partially stripped H^- and H^0 can be easily removed from the beam via a beam dump, while the stripped electrons can be removed via an electric or magnetic field.

There are drawbacks to using stripping foils during H^- injection^{106–109}. The emittance of the beam will grow due to the scattering of the beam off of the foil. The thin foils can also be difficult to handle, very fragile, and have a tendency to curl up. This limits the maximal transverse area of foil that can be used and may limit the transverse acceptance of the beam. Energy is also deposited into the foil by the beam as it passes through. This can degrade the surface of the foil and even create holes. If the beam intensity is high enough, the deposited energy will cause the foil to heat and potentially even sublimate^{110–114}

An alternative to using foil stripping during H^- injection is to use laser stripping^{96,97,115}. The photons in the laser beam can detach loosely bound electrons from the H^- ions. They can also bring electrons into an excited state for H^0 atoms, that can subsequently be stripped using a magnetic field. Work is currently underway to make laser stripping viable for higher intensity beams of different pulse durations.

The foil stripping cross-section depends on the beam energy, with the stripping




Figure 2.1: The charge fraction remaining of a 200 MeV H^- beam as it passes through various carbon foil thicknesses. The lines represent a best fit of Equations 2.5, 2.6 and 2.7 to the experimental results of Webber and Hojvat¹⁰⁵.

cross-section generally decreasing as the beam energy increases. This is due to the progressively shorter interaction times of the beam with the foil. Many authors therefore scale^{114,116-118} the stripping cross-section by a factor of $1/\beta^2$ (where β is the relativistic beta), to determine the optimum thickness of stripping foil to use.

This chapter will investigate the validity of such an approach in the low and high energy limit. For the high energy limit, the consideration will be the injection of a H^- beam into a Muon Collider, based on the 8 GeV Proton Driver from Fermilab's Project X. For the low energy limit, the consideration will be the injection of a 3 MeV beam from the Front End Test Stand¹¹⁹ (FETS) at the Rutherford Appleton Laboratory (RAL) into a vertical fixed field alternating gradient (FFA) machine^{64,65}. The considerations will be based on the current theoretical and experimental results. Additionally, foil stripping measurements were taken at the Kyoto University Institute for Integrated Radiation and Nuclear Science (KURNS) using an 11 MeV H^- beam.

2.1 Electron detachment cross-sections

In gas stripping, electrons can be stripped from the beam as it passes through a beam pipe if the beam interacts with any residual gas particles. This source of beam loss can be reduced by maintaining an appropriate vacuum and reducing the distance the beam travels. Figures 2.2 and 2.3 show some experimental electron loss cross-sections as a function of projectile energy in various gases for H^- and H^0 beams, respectively. As expected, the electron-loss cross-section is larger for H^- beams than for H^0 beams at the same energy. Most experimental results are however at relatively low energies. A theoretical description of the electron detachment cross-section as it extends into higher projectile energies in terms of the Born approximation was given by Gillespie.

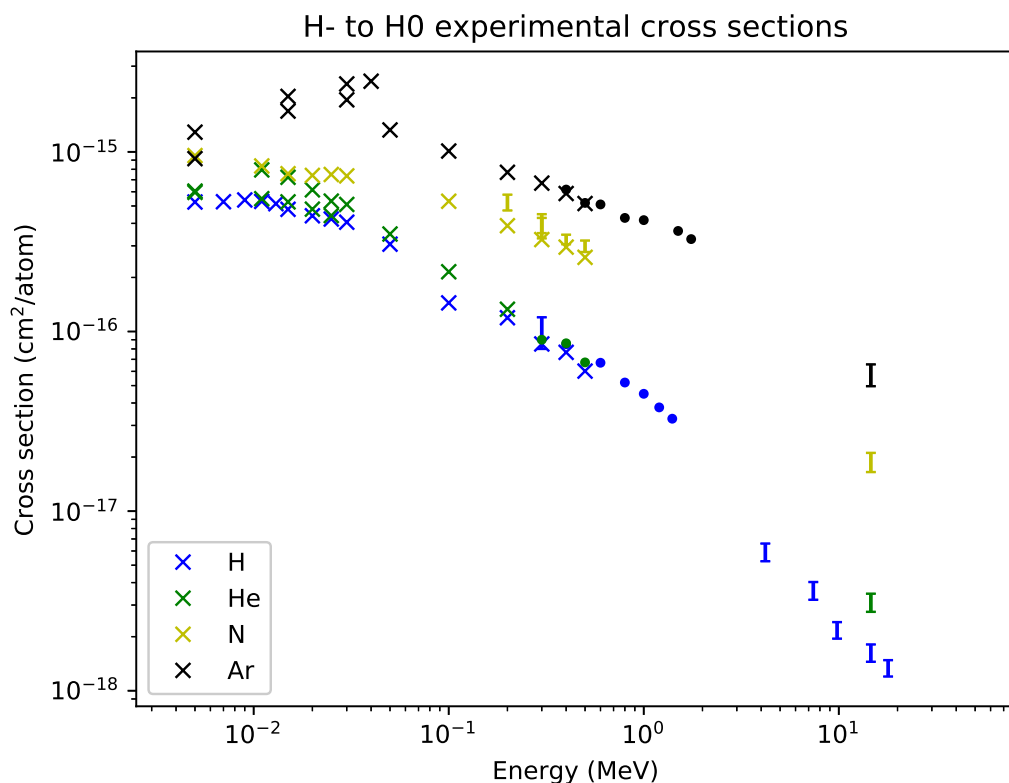


Figure 2.2: Experimental electron loss cross-sections of a H^- beam passing through various gases (error bars where given). The cross-sections are given on a per-atom basis, which ignores any differences in the cross-section of the H^- ion being stripped by an atom or a molecule. Dot markers are the σ_{-10} cross-section, while 'x' markers and values with error bars additionally include the double electron loss cross-section, i.e. $\sigma_{-10} + \sigma_{-11}$. Values are taken from Saha et al.¹²⁰, Heinemeier et al.¹²¹, Alisson¹²², Kovacs¹²³, Dmitriev¹²⁴, and Smythe and Toevs¹²⁵.

For hydrogen, Gillespie^{100–102} derived the following total electron loss cross-section σ_{nm} , the sum of the cross-sections resulting in one electron being lost, $\sigma_{-1,0}$ and both




Figure 2.3: Experimental electron loss cross-sections of a H^0 beam passing through various gases. The σ_{01} cross-sections are given on a per-atom basis with values taken from Alisson¹²², Dmitriev¹²⁴, Acerbi et al.¹²⁶, Welsh et al.¹²⁷ and Smythe and Toevs¹²⁵.

electrons being lost, $\sigma_{-1,1}$:

$$(\sigma_{-1,0} + \sigma_{-1,1}) = 8\pi a_0^2 \frac{\alpha^2}{\beta^2} \sum_{n \neq 0} \sum_m [I_{nm} - J_{nm}(\beta^2) - K_{nm}(\beta^2)] \quad (2.8)$$

where

$$I_{nm} = \int_0^\infty |F_n^{(1)}(K)|^2 |F_m^{(2)}(K)|^2 \frac{d(a_0 K)}{(a_0 K)^3}, \quad (2.9)$$

$$J_{nm}(\beta^2) = \int_0^{a_0 K_{min}} |F_n^{(1)}(K)|^2 |F_m^{(2)}(K)|^2 \frac{d(a_0 K)}{(a_0 K)^3}, \quad (2.10)$$

$$K_{nm}(\beta^2) = \int_{a_0 K_{max}}^\infty |F_n^{(1)}(K)|^2 |F_m^{(2)}(K)|^2 \frac{d(a_0 K)}{(a_0 K)^3}, \quad (2.11)$$

where a_0 , α , β and K are the Bohr radius, the fine structure constant, the beta velocity and the momentum transfer respectively, while $F_n^{(j)}(\vec{K})$ is the elastic form factor given by

$$F_n^{(j)}(\vec{K}) = {}_j \langle n | Z_N^{(j)} - \sum_{l=1}^{Z_e^{(j)}} \exp(i\vec{K} \cdot \vec{r}_l^{(j)}) | 0 \rangle_j, \quad (2.12)$$

where $Z_N^{(j)}$ and $Z_e^{(j)}$ are the nuclear and electronic charges respectively of the j^{th} atomic system.

The summed values calculated by Gillespie for the I_{nm} and J_{nm} integrals of a hydrogen anion striking various target atoms have been tabulated in Table 2.1. For hydrogen like atoms, Gillespie used a 39-term Weiss wave function¹²⁸, while for non-hydrogen like target atoms, the Hartree-Fock (HF) and Configuration-Interaction (CI) models were used. The Weiss, HF and CI models are correlated wave functions describing many electron systems¹²⁹.

The $K_{nm}(\beta^2)$ integral differs from the $J_{nm}(\beta^2)$ integral only by its integration limits and is a factor of $(m_e/m_p)^2$ smaller than the $J_{nm}(\beta^2)$ integral. The $K_{nm}(\beta^2)$ integral therefore has a negligible contribution.

Table 2.1: Tabulated I_{nm} and $J_{nm}(\beta^2)$ for incident H^- on various target atoms using the Configuration-Interaction (CI) and Hartree-Fock (HF) models.

Atom	Model	$S^{(2)}(-1)$	$S^{(2)}(0)$	$S^{(2)}(1)$	$\sum_{n \neq 0} \sum_m I_{nm}$	$\frac{\beta^2}{\alpha^2} \sum_{n \neq 0} \sum_m J_{nm}(\beta^2)$
H^-	-	7.484	2	1.495	-	-
H	-	1	1	4/3	2.42	2.434
H_2	-	1.55	2	3.34	-	2 x 2.71
He	-	0.7525	2	8.167	2.81	10.28
C	CI	2.9	6	100	16.85	110
N	HF	2.6	7	150	18.58	160
N	CI	2.6	7	150	17.76	160
O	HF	2.3	8	200	19.38	210
O	CI	2.3	8	200	18.75	210
Ne	HF	2.02	10	303	20.3	316
Ne	CI	2.02	10	303	19.9	316
Ar	HF	5.502	18	1149	65.2	1170
Ar	CI	5.502	18	1149	62.4	1170
Kr	HF	7.86	36	5340	152.8	5350
Xe	HF	10.3	54	13200	315	13100
Rn	HF	14.4	86	36300	567.8	35900

The cross-sections $(\sigma_{-1,0} + \sigma_{-1,1}) = 8\pi a_0^2 \frac{\alpha^2}{\beta^2} \sum_{n \neq 0} \sum_m [I_{nm} - J_{nm}(\beta^2)]$ of incident H^- ions on various target atoms have been plotted in Fig. 2.4 using the values calculated by Gillespie. There are however some limitations to Gillespie's method, especially in comparison to experimental measurements (Fig. 2.5).

The first is that the first Born approximation is not valid at low projectile energies. It assumes a weak scattering potential where the wave function is only slightly different from the incident plane wave. The region of validity given by Inokuti¹³⁰ is

$$T = \frac{1}{2}mv^2 \gg RZ_N^2, \quad (2.13)$$

where R is the Rydberg energy and Z_N^2 is the nuclear charge. The purpose of Eq. 2.13 is to remove the non-physical low energy behaviour. Therefore, in Fig. 2.5, the electron loss cross-section of H^- impacting on H and He has been plotted for energies where $T > 2RZ_N^2$, and for H^- impacting on other elements at energies where $T > RZ_N^2$.

The electron loss cross-sections are often measured from the impact of H^- ions against target molecules rather than against target atoms, and subsequently reported on a per-atom basis. This limits the accuracy to which experimental results can be compared to the theoretical description by Gillespie. For example, the electron loss cross-section of H^- impacting on a H_2 molecule may not necessarily be twice the electron loss cross-section of H^- impacting on a H atom, but only approximately that. This difference can be seen in the summation of the $J_{nm}(\beta^2)$ integral in Table 2.1, where on a per-atom basis, the $J_{nm}(\beta^2)$ integral of H^- impacting on a H_2 molecule is approximately 11% larger than that of H^- impacting on a sole hydrogen atom. One should also note that the cross-sections represented by dots in Fig. 2.5 do not include the $\sigma_{-1,1}$ cross-sections, however they are expected to be only a few percent of the $\sigma_{-1,0}$ cross-section¹⁰⁰.

The cross-sections are a sum of all possible electron detachment cross-sections from the ground state up. Thus, if our particles are already in an excited state (e.g. for a H^0 ion impacting a target atom), the measured electron loss cross-sections will differ from those predicted by Gillespie¹⁰¹.

Gillespie's method uses the first Born approximation and thus neglects higher order effects. Gillespie himself supposed that the second and higher order Born approximations¹⁰⁰ may introduce corrections of the order $(\frac{\alpha^3}{\beta^3})$ and higher. These contributions could improve the low energy results. Figures 2.4 and 2.5 already give an indication that the outer shell structure is of importance to the cross-section, with higher order Born approximations likely to give rise to sub-shell structure.

Gillespie has described the cross-section in Eq. 2.8 in terms of the expansion parameter α^2/β^2 , with the approach taken "kinematical in origin and is correct, relativistically." This explains the curvature seen in the cross-sections in Figures 2.4 and 2.5, when the incident H^- beam approaches relativistic velocities^{87,131}. As β approaches 1, the interaction time between the incident H^- ion and target atom becomes constant, with the cross-section no longer reducing in size, though this assumes there are no other relativistic effects.




Figure 2.4: Electron detachment cross-sections for H^- beam incident on various target atoms at various incident projectile energies.




Figure 2.5: Experimental electron detachment cross-sections compared to the theoretical values by Gillespie. 'x' markers do not include $\sigma_{-1,1}$, whereas 'I' do. Errors are given where available.

2.2 Foil stripping

The theory developed by Gillespie in the previous section was for the impact of an incident ion upon a target atom, resulting in the loss of one or more electrons by the incident ion. The treatment doesn't differentiate between target atoms in a gaseous state or in a solid state. This means that the same equations derived in the previous section apply to both gas stripping and foil stripping. The only difference is that for foil stripping, the H^- beam is striking a much denser medium. This also means that if the foil is thick enough, the H^- beam can be completely stripped to H^+ over a very short distance. This is exploited for charge exchange injection. The foils used for charge exchange injection are typically made of carbon due to its low atomic number and high sublimation point.

The stripping efficiency is given by the ratio of the number of H^- ions stripped to H^+ , compared to the initial H^- population. For a given surface density x , the fraction of the H^- beam remaining intact or turning to H^0 or H^+ were given by Equations 2.5, 2.6 and 2.7. Experimental results of the measured H^- , H^0 and H^+ populations for a given foil surface density have been given by Saha et al.¹²⁰, Webber and Hojvat¹⁰⁵, and Gulley et al.¹³² for 181, 200 and 800 MeV H^- beams respectively in Figures 2.1 and 2.6. All three graphs follow the exact same shape and have the same peak in the H^0 population occurring at just over 50%, with the only difference being that higher energy beams require higher foil surface densities to achieve the same charge fractions. This has led many authors to scale^{114,116-118} the stripping cross-section by a factor of $1/\beta^2$, to determine the optimum thickness of stripping foil to use at a given beam energy.

The validity of the scaling can be explained in terms of the I_{nm} and $J_{nm}(\beta^2)$ integrals used to calculate the electron loss cross-sections (Table 2.1). For a given interaction, for example



the I_{nm} and $(\beta^2/\alpha^2)J_{nm}(\beta^2)$ integrals of the interaction remain constant, with the curvatures in the cross-sections at higher energies seen in Fig. 2.5 occurring when the projectile energy becomes relativistic. As the H^- and H^0 particles share similar masses, they will approach relativistic velocities at nearly identical energies, and given that the I_{nm} and $(\beta^2/\alpha^2)J_{nm}(\beta^2)$ integrals remain constant for a given interaction, the ratio between the σ_{-10} and σ_{01} electron loss cross-sections also remains near constant as a function of projectile energy, even at relativistic velocities. This also explains the identical shapes of the graphs in Figures 2.1 and 2.6.

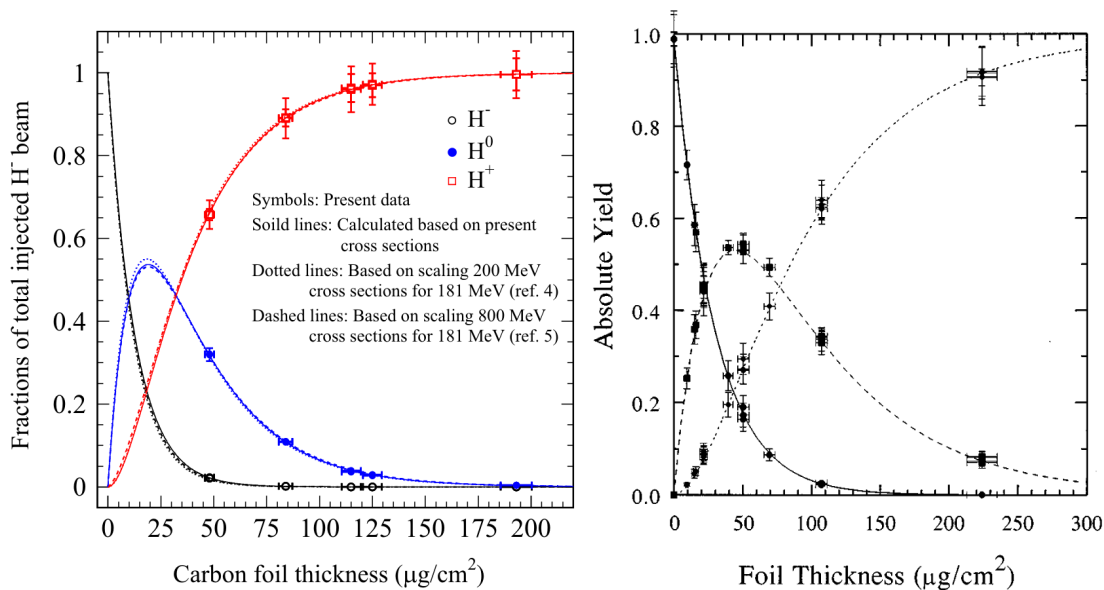


Figure 2.6: Survival Probability of (*Left:*) a 181 MeV, and (*Right:*) a 800 MeV H^- beam passing through various foil thicknesses of carbon by Saha et al.¹²⁰ and Gulley et al.¹³².

Recall that the description by Gillespie was only valid when the projectile energy was sufficiently high (Eq. 2.13). This similarly means that scaling the cross-section by $1/\beta^2$ is also only valid when the projectile energy is sufficiently high. There is some arbitrariness to Eq. 2.13 as to when the projectile energy is large enough. In Fig. 2.5 it has been drawn for H^- ions impacting on carbon when $T > RZ_N^2$, as the purpose of Eq. 2.13 is to remove the non-physical low energy behaviour of the Gillespie treatment i.e. when the cross-section begins to plateau at low energy. For valid projectile energies, the experimental cross-sections tend to be slightly lower than the cross-sections predicted by Gillespie. Additionally, experimental cross-sections of a H^- beam striking a carbon foil have only been taken at projectile energies that are at least an order of magnitude greater than a H^- beam striking any gaseous particle. To bridge this gap in projectile energies, a charge exchange experiment was performed at KURNS using an 11 MeV H^- beam and is described in the following sections. The main purpose of the experiment was to validate the $1/\beta^2$ scaling of the cross-section with projectile energy. This would aid in determining what the optimum foil thickness would be for the 3 MeV H^- beam of the FETS-FFA.

2.3 KURNS facility

The Kyoto University Institute for Integrated Radiation and Nuclear Science (KURNS) facility (Fig. 2.7) houses a 150 MeV Fixed Field Alternating gradient main ring (FFA) and the Multiplex Energy Recovery Internal Target ring (MERIT). The

facility wasn't designed to perform a foil stripping experiment, however alterations to the beam transport line could be made to make such an experiment possible. The design used for the experiment is not an optimal design, but rather what was possible at the KURNS facility with the equipment available.




Figure 2.7: Layout of the KURNS facility. An 11 MeV H^- beam is generated from the linac and transported to the MERIT ring and the FFA Main Ring via a series of dipole and quadrupole magnets.

2.4 Experimental design

The main practical consideration for a foil stripping experiment is to separate and measure the H^- , H^0 , and H^+ charge fractions after the H^- beam had passed through a carbon foil. At KURNS this was only possible near the D2 magnet (also called BM2, Fig. 2.8). If the dipole magnet was switched off, all of the particles in the beam would continue straight down the beamline, however if the magnet was switched on, only neutral particles could continue straight down the beamline, with negative particles travelling around the dipole bend. If the polarity of the dipole magnet was reversed, then positive particles could be selected to travel around the dipole bend instead. The beampipe near the D2 magnet also allowed for beam measurement devices such as Faraday cups to be inserted into the beamline.

Carbon foils could be inserted in two locations, either upstream or downstream of the dipole magnet. Upstream of the dipole magnet, at the Foil 1 location (F1), a movable Faraday cup (FC1) could be inserted behind the foil. Downstream of the dipole magnet, at the Foil 2 location (F2), a stationary Faraday cup (FC2) was placed behind the foil. A third Faraday cup (FC3) was placed after the dipole bend to measure the charged components of the beam when the dipole magnet was switched on.




Figure 2.8: The beam transport line near the D2 magnet at KURNS, where D refers to a dipole magnet and Q refers to a quadrupole magnet.

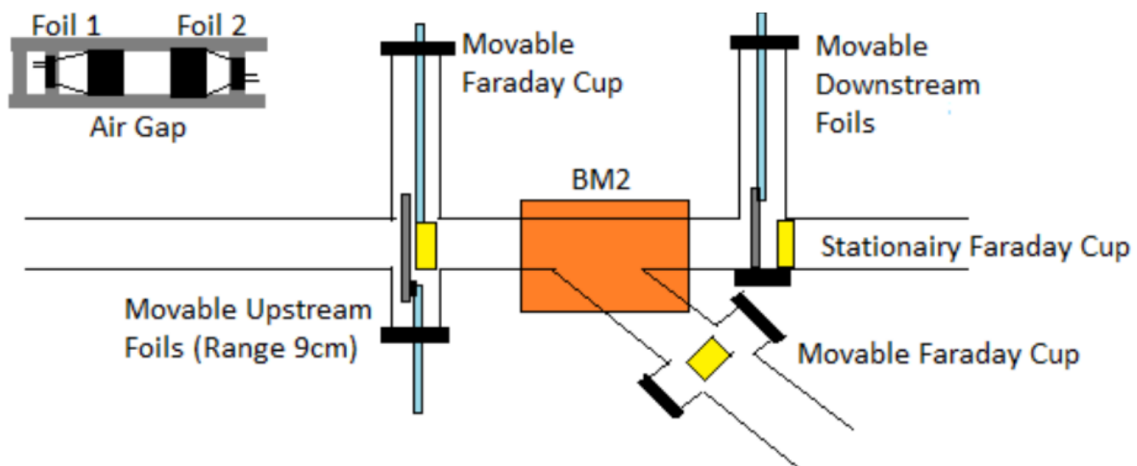


Figure 2.9: The foil stripping measurement set up around the dipole D2 magnet (also called BM2). Behind the upstream foil, at the Foil 1 location (F1), a movable Faraday cup (FC1) could be inserted. Behind the downstream foil, at the Foil 2 location (F2), a stationary Faraday cup (FC2) was placed. A third Faraday cup (FC3) was placed after the dipole bend of the magnet to measure any charged particles in the beam. Up to two foils could be mounted at each foil location at the same time (see top left), with the appropriate foil then slid into the beamline.

The H^- beam will naturally contain some H^0 contamination and some H^+ contamination. These can arise from production processes in the H^- source, or through effects such as intrabeam stripping and residual gas stripping when the beam is travelling through the beam pipe. These effects were reduced at KURNS through the use of a vertical chicane in the transport line between Q3 and Q4. Additionally, when a foil was inserted at the F1 location, neutral particles in the beam were stripped to H^+ and deflected away by the dipole magnet.

The carbon foil at the F1 location will also strip the H^- beam. If the foil is thin enough, then only a fraction of the beam will be fully stripped to H^+ , with the partially stripped neutral component continuing on to FC2. To measure this neutral component, a further thicker foil was placed at the F2 location just before FC2, which completely strips the neutral component to H^+ . The electron loss cross-section due to foil stripping can then be determined by what percentage of the beam was only partially stripped by the first foil. This will depend on the foil thicknesses used and the projectile energy of the beam. During the beam transport line, the H^- beam had a nominal projectile energy of 11 MeV.

The required foil thicknesses for partial stripping can be estimated using the theory by Gillespie or by scaling the experimental cross-sections. Unfortunately, Gillespie only provided values of the I_{nm} and $(\beta^2/\alpha^2)J_{nm}(\beta^2)$ integrals for H^- impacting on carbon and not for H^0 impacting on carbon. This can be circumvented with the help of some experimental results (Table 2.2). The ratio of the σ_{01} to σ_{-10} cross-sections should remain constant with projectile energy, meaning the σ_{-10} cross-section at 11 MeV can be scaled by that ratio to obtain the σ_{01} cross-section at 11 MeV. Using Eqs. 2.5, 2.6 and 2.7, the charge fractions of a 11 MeV H^- beam passing through various surface densities of carbon have been plotted in the top plot of Fig. 2.10. Noting that the experimental cross-sections tend to be a fraction lower than the values predicted by Gillespie, the cross-sections used for the bottom plot of Fig. 2.10 have additionally been scaled by that ratio.

2.5 Foil preparation

The lowest foil thicknesses available at KURNS were approximately $5 \mu\text{g}/\text{cm}^2$. These were made by the Arizona Carbon Foil Company and had an error of $\pm 10\%$ on the foil thickness¹³³. The foils had been arc evaporated onto a glass substrate (a glass slide) covering a transverse area of $50 \text{ mm} \times 70 \text{ mm}$. To remove the foils from the glass substrate, the foils were floated on to the surface of a liquid, in this case deionised water. This was achieved by placing the glass slide onto a tilted holder in a water tank. The tank was then slowly filled with water through a hose pipe to meet the surface of the foil.

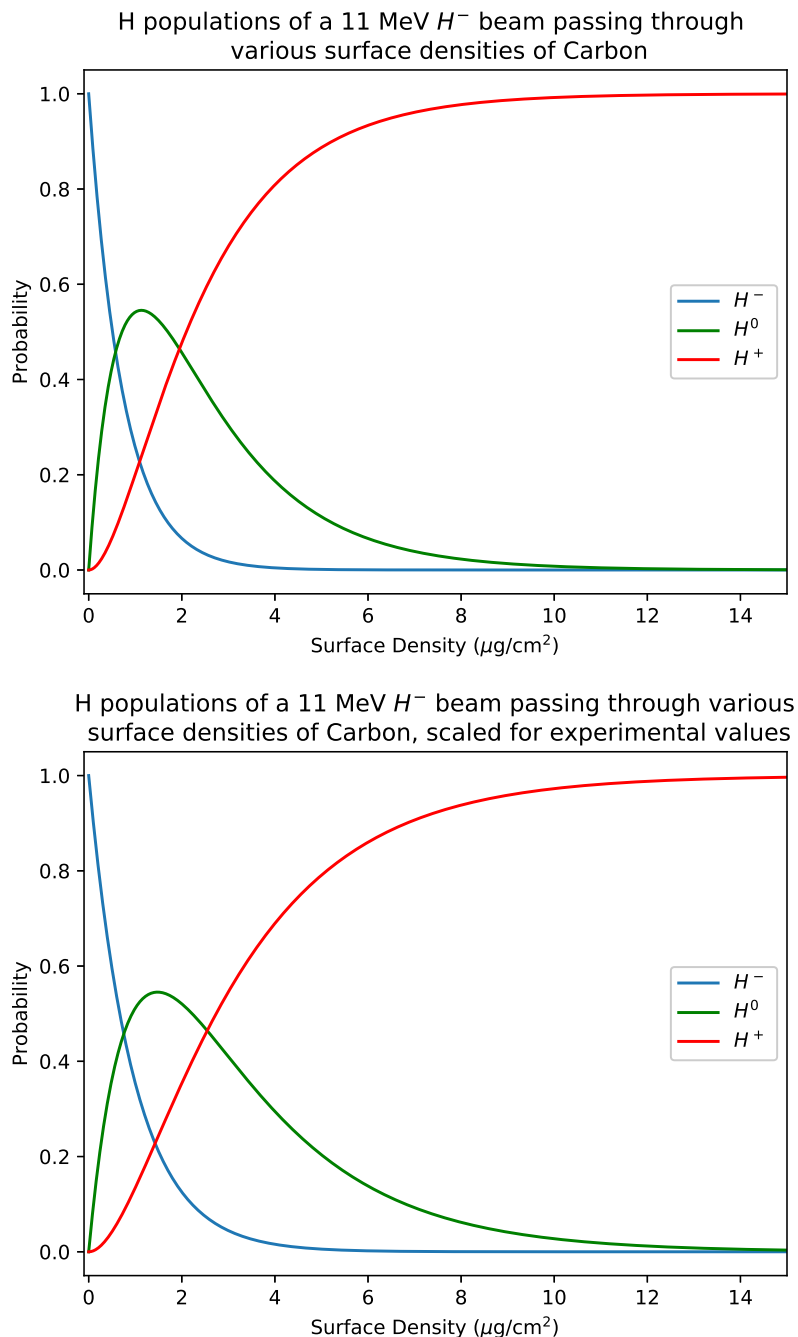


Figure 2.10: Predicted Survival Probability of an 11 MeV H^- beam passing through a given surface density of carbon. The $\sigma_{-1,0}$ cross-section has been obtained using the equations by Gillespie, while the $\sigma_{0,1}$ cross-section has been obtained by scaling the $\sigma_{-1,0}$ cross-section by the ratio between the $\sigma_{-1,0}$ and $\sigma_{0,1}$ cross-sections at higher energies. The cross-sections used for the bottom figure have additionally been scaled by the ratio between the experimental and predicted values.

Energy (MeV)	$\sigma_{-1,0(the)}$	$\sigma_{-1,0(exp)}$	$\sigma_{-1,1(exp)}$	$\sigma_{0,1(exp)}$
181	2.1213	1.58	0.002	0.648
200	1.9677	1.56	0.08	0.6
800	0.8905	0.676	0.012	0.264
Energy (MeV)		$\sigma_{0,1(exp)}/\sigma_{-1,0(exp)}$	$\sigma_{-1,0(exp)}/\sigma_{-1,0(the)}$	$(\sigma_{0,1(exp)} + \sigma_{1,1(exp)})/\sigma_{-1,0(the)}$
181		0.4101	0.7448	0.7458
200		0.3846	0.7928	0.8335
800		0.3905	0.7591	0.7726

Table 2.2: *Top:* Tabulated experimental cross-sections (in units of 10^{-18} cm^2 per atom) of H^- incident on C from Webber and Hojvat, Saha et al. and Gulley et al. as well as Gillespie's theoretical cross-sections. *Bottom:* Ratio between the experimental σ_{01} to σ_{-10} electron loss cross-sections and ratio of the experimental to theoretical σ_{-10} electron loss cross-sections as a function of energy.

The foil then slowly peeled from the glass as the water level was raised further. The floated foil was then guided over a foil holder by gently blowing air on the surface of the water near the foil. The foil holder had two pencil inserts attached using aluminium tape, one near the bottom and one near the top of the foil holder. When the floated foil partially overlapped with the submerged foil holder in the water tank, the water was slowly removed from the tank via a hose pipe (by suction), causing the foil to slowly hang off of the top pencil insert of the foil holder as if on a washing line. Either side of the pencil insert, the overhanging foil adhered to itself, causing the effective foil thickness to double. The bottom pencil insert relieved the surface tension on the foil, and as the rest of the foil was removed from the water, the small remaining leftover foil was torn by the surface tension along this edge. This method allows for thin foils with a large transverse area to be made.

For KURNS, a foil with a surface density of $9 \mu\text{g}/\text{cm}^2$ and transverse area of $4 \text{ cm} \times 3 \text{ cm}$ was made using this method. This foil was then installed at the foil 1 location of the experiment (Fig. 2.11). The foil was however very fragile and by the end of the experiment, the foil had developed a small tear near one of its edges. Using the electron loss cross-sections from Fig. 2.10, an 11 MeV H^- beam passing through a foil with a surface density of $9 \mu\text{g}/\text{cm}^2$ would be expected to produce measurable charge fractions of 0% H^- , 1.33% H^0 and 98.67% H^+ (or $<0.01\%$ H^- , 4.12% H^0 and 95.88% H^+ when scaled for experimental results). A foil of $17.6 \mu\text{g}/\text{cm}^2$ total thickness was also installed at the foil 2 location, and would be expected to nearly completely strip a H^- beam to H^+ .




Figure 2.11: The upstream $9 \mu\text{g}/\text{cm}^2$ thick carbon foil inserted into the beamline at the Foil 1 location. The foil has been rotated by 90 degrees to allow a top-down view. The foil was very fragile and shows a slight tear along one of its edges.

2.6 Experimental run plan

Even with the thinnest foils available at KURNS, the unstripped H^- component is expected to be tiny and not measurable using the Faraday cups available at KURNS. This means the determination of the electron detachment cross-sections will depend on the measured H^0 and H^+ populations. These measured populations can however be biased by any contamination present in the beam. The beam travelling through the transport line at KURNS has three main components:

- the H^- beam,
- the H_{cont}^0 beam contamination,
- the H_{cont}^+ beam contamination.

The measured populations at FC2 and FC3 will depend on the polarity of the dipole magnet. Additionally, the measured populations will change when carbon foils are present in the beamline. To measure the partial stripping of the H^- beam by the upstream foil, the populations of the H_{cont}^0 and the H_{cont}^+ beam contamination need to be accounted for.

The experimental run plan that was used is detailed in Table 2.3. Runs 2–14 cover all 12 combinations of Dipole settings and foil insertion positions ($3 \times 2 \times 2$). Additional runs were taken with the upstream foil moved from its initial central

position to determine the width of the beam, and runs with the first Faraday cup inserted into the beamline to measure changes in the intensity of the beam.

For the H^- beam, the following components can arise:

- $H^- \rightarrow H_{uns}^-$, the H^- beam remains unstripped (expected to be minute).
- $H^- \rightarrow H^+$, the beam is fully stripped to H^+ by either foil.
- $H^- \rightarrow H^0$ at F1, the beam is only partially stripped by the upstream foil to H^0 and then either gas stripped or foil stripped by the downstream foil.

The partial stripping of the H^- beam to H^0 can be described by the partial stripping factor F_{PS1} . The measurement of this factor allows for the measurement of the electron detachment cross-sections of an 11 MeV H^- beam. For the H_{cont}^0 beam contamination, the following components can arise:

- $H_{cont,GS}^0$, the H_{cont}^0 beam contamination is gas stripped to H^+ between F1 and Faraday cup 2.
- $H_{cont,F1S}^0$, the H_{cont}^0 beam contamination is foil stripped to H^+ by F1.
- $H_{cont,F2S}^0$, the H_{cont}^0 beam contamination is foil stripped to H^+ by F2.

Similarly, the H_{cont}^0 beam contamination may not be fully stripped by the upstream foil to H^+ and can be described by the partial stripping factor F_{PSZ} . Additionally, a factor F_{GS} can be used to describe the fraction of the H_{cont}^0 contamination that is gas stripped between the dipole magnet and FC2. The H^+ beam contamination can not be stripped further and thus will remain H^+ .

If any beam component passes through a foil, it will be affected by scattering and thus a reduction in the measured beam intensity. The change in the physical size of the beam as it is affected by scattering couldn't be measured directly, and thus the effect on the change in the beam intensity as it passes through either foil is approximated by the factors F_1 and F_2 , the ratio in the beam intensities when a foil was present compared to when it was not.

Table 2.4 indicates the measurable hydrogen components at FC2 and FC3 in each of the runs that are due to the beam contaminants, stripped particles or various combinations. The 12 run combinations may feel superfluous, however, most measurements are in some way biased by the H^0 or H^+ beam contamination and must be accounted for. The following subsections will deal with the measurements of the beam signal and its contaminants for each run.

Table 2.3: The experimental run plan, going through all 12 combinations of dipole and foil insertion settings (runs 2-14), while minimising any foil movements between runs. Extra runs were taken to investigate the effect of the beam only partially striking the foil. The amplifier used for Faraday cup 1 had an issue, as the polarity of the signal never changed when impacted on by positive or negative beams, and therefore those measurements should be treated with care. The expected signal strength measured at a Faraday cup is given as a percentage of the typical H^- beam intensity (Values in brackets are scaled for experimental results). Values below 100%, which would indicate partial stripping, assume a perfect experimental environment. Measurements affected by beam impurities and residual gas stripping are shown in Table 2.4. Unfortunately, the linac producing the H^- beam underwent a failure when files 081-090 were recorded. This resulted in a substantially lower beam intensity for measurements taken after that point. The right-hand column corresponds to the filenames for each run¹³⁴.

Run	Dip	F1	F2	FC1	FC2	FC3	Signal Strength			Filename tek0xxxALL.csv
							@ FC1 (%)	@ FC2 (%)	@ FC3 (%)	
1	Off	No	No	Yes	Yes	Yes	100	0	0	001 – 010
2	Off	No	No	No	Yes	Yes	0	100	0	011 – 020
3	On	No	No	No	Yes	Yes	0	0	100	021 – 030
4	On	No	Yes	No	Yes	Yes	0	0	100	031 – 040
5	Off	No	Yes	No	Yes	Yes	0	100	0	041 – 050
6	Off	Yes	Yes	No	Yes	Yes	0	100	0	051 – 060
7	On	Yes	Yes	No	Yes	Yes	0	1.33 (4.12)	0	061 – 070
8	On	Yes	No	No	Yes	Yes	0	0	0	071 – 080
9	Off	Yes	No	No	Yes	Yes	0	98.67 (95.88)	0	081 – 090
10	Rev	Yes	No	No	Yes	Yes	0	0	98.67 (95.88)	091 – 100
11	Rev	Yes	Yes	No	Yes	Yes	0	1.33 (4.12)	98.67 (95.88)	101 – 110
12	Rev	No	Yes	No	Yes	Yes	0	0	0	111 – 120
13	Rev	No	No	No	Yes	Yes	0	0	0	121 – 130
14	Off	No	No	No	Yes	Yes	0	100	0	131 – 140
15	Off	No	No	No	Yes	Yes	0	100	0	141 – 150
16	On	Yes	Yes	No	Yes	Yes	0	1.33 (4.12)	0	151 – 160
17	On	Yes	Yes	No	Yes	Yes	0	1.33 (4.12)	0	161 – 170
18	Off	No	No	Yes	Yes	Yes	100	0	0	171 – 180

Table 2.4: The experiment designed in Fig. 2.9 is expected to measure the partial stripping of an 11 MeV H^- beam by a carbon foil. Residual gas stripping and H^0 and H^+ contaminant components can bias the foil stripping measurements. The table accounts for what beam components and contaminants can be measured by Faraday cup 2 (FC2) at the end of the straight section of the experiment and by Faraday cup 3 (FC3) after the curved section of the dipole magnet. The dipole magnet (Dip) was run in three settings: Off, where the beam travels straight to the end of the beamline, hitting Faraday cup 2 (FC2); On, where the H^- components travel through the dipole bend hitting Faraday cup 3 (FC3); and reverse polarity, where the polarity of the dipole magnet is reversed and the H^+ component travels to FC3. When a thin carbon foil is placed in the path of the H^- beam, it will primarily strip to H^+ and partially strip to H^0 . The H^0 component can be further stripped by a thicker carbon foil (F2) and therefore measure the partial stripping achieved by the first foil (F1). In the table, GS means residual gas stripping, while F1S and F2S is stripping by the first and second foils respectively.

Faraday cup				FC2	FC2	FC2	FC2	FC2	FC2	FC2	FC2	FC2	FC2	FC3	FC3	FC3
Component Measured				H-	H+	H+	H+	H+	H+	H+	H+	H+	H+	H-	H+	H+
Contaminant					H+	H0	H0	H0							H+	
Stripping										H-	H-	H- → H0	H- → H0			H-
										F1S	F2S	@F1 GS	@F1 F2S			F1S
Run	Dip	F1	F2													
2, 14	Off	No	No	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No
5	Off	No	Yes	No	Yes	Yes	No	Yes	No	Yes	No	No	No	No	No	No
9	Off	Yes	No	No	Yes	No	Yes	No	Yes	No	Yes	No	No	No	No	No
6	Off	Yes	Yes	No	Yes	No	Yes	No	Yes	No	Yes	Yes	Yes	No	No	No
3	On	No	No	No	No	Yes	No	No	No	No	No	No	No	Yes	No	No
4	On	No	Yes	No	No	No	No	Yes	No	No	No	No	No	Yes	No	No
8	On	Yes	No	No	No	No	No	No	No	No	Yes	No	No	No	No	No
7	On	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes	Yes	No	No	No
13	Rev	No	No	No	No	Yes	No	No	No	No	No	No	No	No	Yes	No
12	Rev	No	Yes	No	No	No	No	Yes	No	No	No	No	No	No	Yes	No
10	Rev	Yes	No	No	No	No	No	No	No	No	Yes	No	No	No	Yes	Yes
11	Rev	Yes	Yes	No	No	No	No	No	No	No	Yes	Yes	Yes	No	Yes	Yes

2.6.1 Measurement of the partial stripping factor F_{PS1} - Foil (Measurements at FC2 for runs 7 and 11)

The signals measured at FC2 for runs 7 and 11 respectively can be used to measure the partial stripping of the H^- beam and therefore the partial stripping factor F_{PS1} . Using Table 2.4, the measured signal at FC2 can be described by the following equation:

$$H_{measured} = \left((H^- \rightarrow H^0)_{GS} + (H^- \rightarrow H^0)_{F2S} + H_{cont,uns,GS}^0 + H_{cont,uns,F2S}^0 \right) \times F_1 \times F_2 \quad (2.15)$$

For both runs, the dipole magnet was on, with the only difference being the polarity of the dipole magnet. Therefore, only neutral particles can travel through the dipole magnet towards FC2. The only neutral particles arising from the beam come from the H^- beam that was only partially stripped to H^0 or from the H_{cont}^0 that was left unstripped by the upstream foil. Each component can then be either gas stripped or foil stripped by the thick downstream foil to H^+ . Finally, the measured signal will be affected by scattering in both foils.

The H^- population that is left unstripped can be described by $1 - F_{PS1}$, while the H_{cont}^0 population that is left unstripped can be described by $1 - F_{PSZ}$. If the H^- beam is stripped to H^+ , the polarity of the measured signal changes, introducing a factor of -1. Finally, as the downstream foil is so thick, no partial stripping should occur, meaning the downstream component is solely made up of the partially stripped H^- beam and the unstripped H_{cont}^0 beam contamination, that have both been either gas or foil stripped. Eq. 2.15 then simplifies to:

$$H_{measured} = \left(-(1 - F_{PS1}) \times H^- + (1 - F_{PSZ}) \times H_{cont}^0 \right) \times F_1 \times F_2 \quad (2.16)$$

The partial stripping factor F_{PS1} is then given by

$$F_{PS1} = \frac{H_{measured}}{F_1 F_2 H^-} - (1 - F_{PSZ}) \frac{H_{cont}^0}{H^-} + 1 \quad (2.17)$$

where H^- and H_{cont}^0 are the initial populations of each component within the beam travelling through the transport line. To calculate the factor F_{PS1} , these populations, as well as the factors F_1 , F_2 and F_{PSZ} , need to be found and will be determined across the 12 different runs. The following subsections will describe what each run will measure at FC2 and FC3.

2.6.2 Measurement of the partial stripping factor F_{PS1} - Gas (Measurements at FC2 for runs 8 and 10)

Similarly to runs 7 and 11 in the previous section, runs 8 and 10 can also consider the partial stripping of the upstream foil. However, as the downstream foil is not present in the beamline, only the neutral particles that were gas stripped are measured at FC2. The measured signal at FC2 is then given by:

$$H_{measured} = \left((H^- \rightarrow H^0)_{GS} + H_{cont,uns,GS}^0 \right) \times F_1 \quad (2.18)$$

Residual gas stripping of the neutral component of the beam between the dipole magnet and FC2 is described by the factor F_{GS} . Eq. 2.18 then similarly simplifies to:

$$H_{measured} = \left(- (1 - F_{PS1}) \times H^- \times F_{GS} + (1 - F_{PSZ}) \times H_{cont}^0 \times F_{GS} \right) \times F_1 \quad (2.19)$$

The partial stripping factor F_{PS1} is then given by

$$F_{PS1} = \frac{H_{measured}}{F_1 F_{GS} H^-} - (1 - F_{PSZ}) \frac{H_{cont}^0}{H^-} + 1 \quad (2.20)$$

2.6.3 Measurement of the H_{cont}^0 beam contamination (Measurements at FC2 for runs 3, 4, 12 and 13)

When the upstream foil is removed from the beamline and the dipole magnet is on, only the H_{cont}^0 beam contamination can pass through the dipole magnet. When the downstream foil is removed from the beamline, then FC2 measures the neutral component of the beam that was gas stripped. Runs 3 and 13 then measure:

$$H_{measured} = H_{cont}^0 \times F_{GS} = H_{cont,GS}^0 \quad (2.21)$$

If the downstream foil was present, however, then the remaining neutral component is completely stripped by the thick downstream foil. For runs 4 and 12, the completely stripped H_{cont}^0 beam contamination is measured, though it is affected by scattering in the downstream foil.

$$H_{measured} = (H_{cont,GS}^0 + H_{cont,F2S}^0) \times F_2 = H_{cont}^0 \times F_2 \quad (2.22)$$

2.6.4 Measurement of the variance of the beam intensity (Measurements at FC2 for runs 2 and 14)

Runs 2 and 14 had the exact same conditions and were used to measure the variance of the beam intensity during the experiment. When no foils were present in the beamline and the dipole magnet was switched off, FC2 measured the whole beam intensity, including the beam contamination:

$$H_{measured} = H^- + H_{cont}^+ + H_{cont,GS}^0 \quad (2.23)$$

The H^- beam component is then given by:

$$H^- = H_{measured} - H_{cont}^+ - H_{cont,GS}^0 \quad (2.24)$$

2.6.5 Measurement of the scattering factors F_1 and F_2 (Measurements at FC2 for runs 5, 6 and 9)

When the dipole magnet was switched off, the effect of scattering by either foil could be determined using the measured signal at FC2. For run 5, when the downstream foil was present in the beamline, the measured signal is given by:

$$H_{measured} = (H_{cont}^+ + H_{cont,GS}^0 + H_{cont,F2S}^0 + (H^- \rightarrow H_{F2S}^+) + H_{uns,F2}^-) \times F_2 \quad (2.25)$$

The thick downstream foil should result in no unstripped H^- component being present, with a partial stripping factor of 1 assumed ($F_{PS2} = 1$), while a change of charge states from H^- to H^+ introduces a factor of -1. The thickness of the downstream foil also means that the whole of the H_{cont}^0 beam contamination is either gas or foil stripped. Eq. 2.25 then simplifies to:

$$H_{measured} = (H_{cont}^+ + H_{cont}^0 - H^-) \times F_2 \quad (2.26)$$

The scattering factor F_2 is then given by

$$F_2 = \frac{H_{measured}}{H_{cont}^+ + H_{cont}^0 - H^-} \quad (2.27)$$

A similar description can be made when only the upstream foil is present in the beamline for the scattering factor F_1 . The measured signal for run 9 is given by:

$$H_{measured} = \left[H_{cont}^+ + H_{cont,F1S}^0 + H_{cont,uns,GS}^0 + (H^- \rightarrow H_{F1S}^+) + (H^- \rightarrow H^0)_{GS} + (H^- \rightarrow H_{uns}^-) \right] \times F_1 \quad (2.28)$$

$$H_{measured} = \left[H_{cont}^+ + H_{cont}^0 (F_{PSZ} + (1 - F_{PSZ})F_{GS}) - H^- (F_{PS1} + (1 - F_{PS1})F_{GS} + H_{uns}^-) \right] \times F_1 \quad (2.29)$$

However, as the partial stripping factors F_{PS1} and F_{PSZ} are unknown, it is difficult to determine what the scattering factor F_1 should be. Instead, run 6 will be considered where both foils are present in the beamline and every part of the beam is completely stripped. The measured signal for run 6 is given by:

$$H_{measured} = \left[H_{cont}^+ + H_{cont,F1S}^0 + H_{cont,uns,GS}^0 + H_{cont,uns,F2S}^0 + (H^- \rightarrow H_{F1S}^+) + (H^- \rightarrow H^0)_{GS} + (H^- \rightarrow H^0)_{F2S} + (H^- \rightarrow H_{uns}^-)_{F2S} \right] \times F_1 \times F_2 \quad (2.30)$$

The thick downstream foil completely strips the three H_{cont}^0 contamination contributions, meaning they can be combined to the original H_{cont}^0 beam contamination:

$$\begin{aligned} H_{cont,F1S}^0 + H_{cont,uns,GS}^0 + H_{cont,uns,F2S}^0 &= H_{cont}^0 \times F_{PSZ} \\ &+ (1 - F_{PSZ}) \times H_{cont}^0 \times F_{GS} \\ &+ (1 - F_{PSZ}) \times H_{cont}^0 \times (1 - F_{GS}) \\ &= H_{cont}^0 \end{aligned} \quad (2.31)$$

This can be similarly be done for three stripped H^- contributions:

$$\begin{aligned} (H^- \rightarrow H^+)_{F1S} + (H^- \rightarrow H^0)_{GS} + (H^- \rightarrow H^0)_{F2S} &= -F_{PS1} \times H^- \\ &- (1 - F_{PS1})(F_{GS}) \times H^- \\ &- (1 - F_{PS1})(1 - F_{GS}) \times H^- \\ &= -H^- \end{aligned} \quad (2.32)$$

Eq. 2.30 can then be rewritten as

$$H_{measured} = \left[H_{cont}^+ + H_{cont}^0 - H^- - H_{uns}^- \right] \times F_1 \times F_2 \quad (2.33)$$

F_1 is then given by

$$F_1 = \frac{H_{measured}}{(H_{cont}^+ + H_{cont}^0 - H^- - H_{uns}^-) \times F_2} \quad (2.34)$$

2.6.6 Measurement of the H^+ beam contamination (Measurements at Faraday cup 3)

The measurements for Faraday cup 3 are simpler to consider. For runs 2, 5, 6 and 9 the dipole magnet is switched off which should result in no measured signal:

$$H_{measured} = 0 \quad (2.35)$$

As the downstream foil cannot affect the measurements at FC3, each of the remaining measurements is effectively taken twice. When the dipole magnet was switched on, runs 3 and 4 measured the full H^- beam

$$H_{measured} = H^- \quad (2.36)$$

while runs 7 and 8 measured the H^- beam component that remained unstripped. The intensity of the measured beam component will however be affected by scattering off of the first foil.

$$H_{measured} = H_{uns}^- \times F_1 \quad (2.37)$$

When the polarity of the dipole magnet was reversed, runs 12 and 13 measured the H_{cont}^+ beam contamination

$$H_{measured} = H_{cont}^+ \quad (2.38)$$

while runs 10 and 11 measured the H_{cont}^+ beam contamination and the stripped H^- beam component.

$$H_{measured} = (H_{cont}^+ - H^- \times F_{PS1}) \times F_1 \quad (2.39)$$

As the H^- beam component changes polarity when stripped, the measured signal will also change polarity and is therefore multiplied by a factor of -1.

2.7 Experimental results

The pressure level recorded throughout the period of data-taking was $1-2 \times 10^{-5}$ Pa. There is some uncertainty on this value as the experiment was performed in the beam transport line, while the pressure level was recorded in the main ring. When the foils were installed in the beam transport line, the vacuum needed to be broken. This resulted in some gas leak and outgassing issues thereafter. By the time the experiment was performed, the pressure levels were still higher than during typical beam operation at KURNS. The measured signal for the gas stripped component

may therefore be affected by the vacuum improving throughout the experiment, though the vacuum improves by no more than a factor of two.

The H^- beam will naturally undergo fluctuations. Therefore, 10 measurements were taken for each run, with each measurement recording the average output of the amplified beam current at the Faraday cups from 512 beam pulses. The measurements were then recorded in files as detailed in Table 2.3. During run 9, the linac underwent a brief failure. Once the run was completed, the experiment was paused until that afternoon. However, when the experiment resumed with run 10, the beam pulse shape had changed and the beam was now operating at a significantly lower intensity. This makes comparison of runs before and after run 9 more difficult.

The current measured by each Faraday cup in the beam transport line at KURNS was amplified using a current to voltage amplifier, with the subsequent voltage recorded on an oscilloscope in the control room. The amplifier also inverted the polarity of the signal, resulting in a H^- beam that displayed a positive voltage. Faraday cup 1 appeared to have a problem, however, as it always showed a positive signal, no matter whether a positive or negative beam had impacted the Faraday cup. Therefore, only the results of Faraday cups 2 and 3 will be considered here.

When the upstream foil was placed into the centre of the beamline, the large transverse area of the foil meant none of the beam could reach FC2 without passing through the foil. This was not the case for FC3 due to the dipole magnet. A small number of slower or faster particles could pass either side of the foil and still find a path to be bent onto FC3 due to the dipole magnet. That path corresponded with a position-momentum correlation in the transport line. This was due to the presence of a vertical chicane in the transport line. The position-momentum correlation meant a time cut could be applied to the beam pulse, to select only particles that did pass through the upstream foil (0.015 to 0.075 ms).

To compare all of the runs, the integral of the measured signal during the interval of the time cut was taken. To ensure a fair comparison, a zero offset was removed from the measured signal before the interval was taken. This was due to the measured signal showing a zero offset even when the beam was off. Figures 2.12 and 2.13 show the integrated beam voltage in a time interval going from 0.015 to 0.075 ms for the various dipole magnet and foil insertion configurations for Faraday cups 2 and 3 respectively. The integrated values have also been tabulated in Table 2.5.

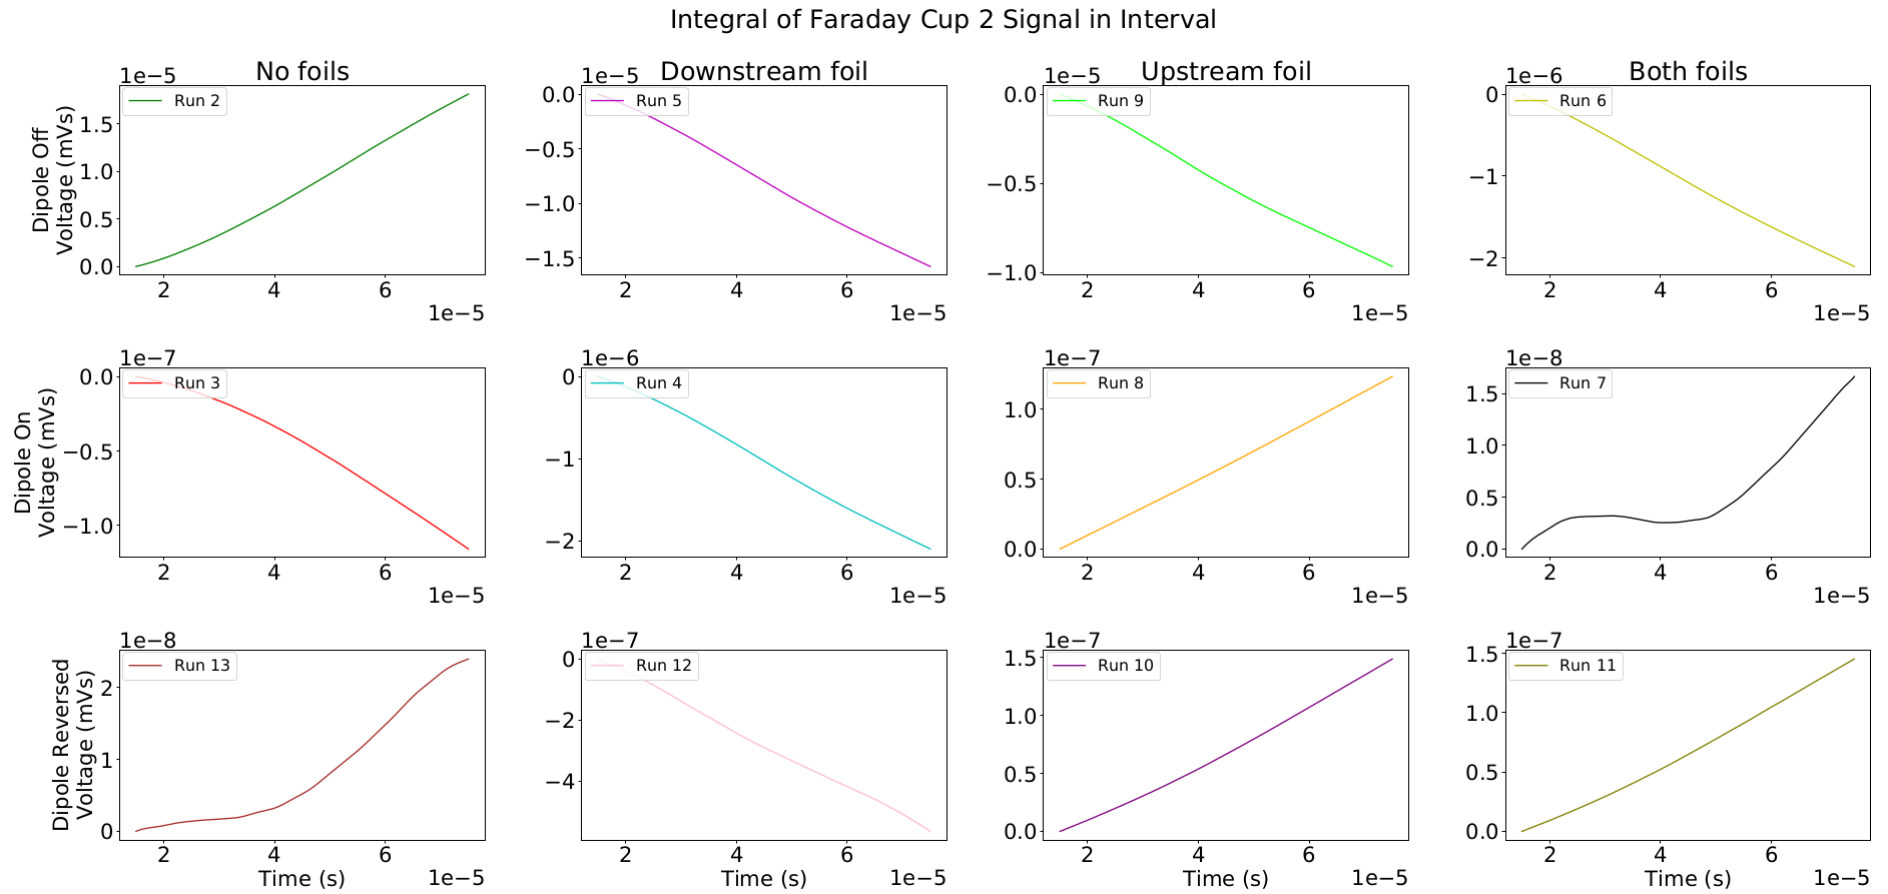


Figure 2.12: Data from the 5 December 2019. The integral of the corrected Faraday cup 2 signal for the twelve different run combinations within a time interval where the beam occupies an area covered by the foil. Provided the integrated voltage is large enough, the integrated voltage as a function of time is approximately linear in the time interval going from 0.015 to 0075 ms.




Figure 2.13: Data from the 5 December 2019. The integral of the corrected Faraday cup 3 signal for the twelve different run combinations within a time interval where the beam occupies an area covered by the foil. Provided the integrated voltage is large enough, the integrated voltage as a function of time is approximately linear in the time interval going from 0.015 to 0.075 ms.

Table 2.5: The integrated average beam signal measured in the time interval going from 0.015 to 0.075 ms at Faraday cup 2 and Faraday cup 3 after the zero offset has been removed from all runs in units of 10^{-9} mVs i.e. pVs. The errors given are only the error on the average beam signal measurement.

Run	Dip	F1	F2	FC2	Error	FC3	Error
2	Off	N	N	17994.805	35.693	-54.132	1.022
14	Off	N	N	7300.915	129.807	-18.971	1.186
5	Off	N	Y	-15926.172	89.362	41.865	1.796
9	Off	Y	N	-9656.308	143.836	20.668	1.496
6	Off	Y	Y	-2170.824	4.808	-1.000	2.334
3	On	N	N	-204.816	2.065	22096.392	177.080
4	On	N	Y	-2150.464	11.891	20678.501	67.198
8	On	Y	N	-6.875	0.928	12.782	1.278
7	On	Y	Y	-113.456	0.512	16.501	1.028
13	Rev	N	N	-107.132	0.562	-110.204	0.474
12	Rev	N	Y	-638.699	8.901	-105.924	0.697
10	Rev	Y	N	16.770	0.911	-12354.274	143.926
11	Rev	Y	Y	13.092	0.965	-12001.341	128.333

2.8 Calculating the partial stripping factor F_{PS1}

Sections 2.6.1 and 2.6.2 showed that the partial stripping factor, F_{PS1} , depended on the factors F_1 , F_2 , F_{PSZ} and F_{GS} . Additionally, F_{PS1} depended on the H^- beam and H_{cont}^0 beam contamination populations, which can vary throughout the experiment. The following subsections will determine the variance of the beam intensity throughout the experiment, as well as the various factors required to calculate F_{PS1} .

2.8.1 Error on the beam intensity

Due to the linac failure during run 9 and the subsequent reduction in the beam intensity thereafter, the specific runs taken to measure the variance of the beam intensity at the start and end of the experiment showed a great variance. When the dipole magnet was switched off and no foils were present in the beamline, Faraday cup 2 measured an integrated voltage (in units of picoVolt-seconds) of 17994.805 pVs and 7300.915 pVs for runs 2 and 14 respectively. The average integrated voltage for runs 2 and 14 is 12647.860 ± 5346.945 pVs, giving an uncertainty of $\pm 42.3\%$ on every measurement. This large error is unfortunate, with the variance on the beam signal likely to be much lower either side of the linac failure. It was attempted to take a second set of data, however the downstream foil holder failed during those runs. This means the large error will permeate throughout the rest of the results.

2.8.2 H^- beam and H^+ beam contamination at FC3

The signals measured by Faraday cups 2 and 3 cannot be directly compared, as the transverse areas of each cup will occupy slightly different transverse areas of the beam. The ratio of the beam populations is not expected to change, however. Therefore, the ratio of the H_{cont}^+ beam contamination compared to the H^- beam population is expected to remain constant for both Faraday cups and will be exploited as the H_{cont}^+ beam contamination cannot be measured directly by FC2.

Runs 3 and 4 measured the pure H^- beam at FC3, finding an average integrated voltage of 21387.447 ± 189.401 pVs. Similarly, the pure H^+ beam contamination was measured during runs 12 and 13 at FC3, finding an average integrated voltage of 108.064 ± 0.843 pVs. When including the variance of the beam intensity, the H^+ beam contamination as a fraction of the pure H^- beam signal is $0.5\% \pm 0.3\%$.

The H^- beam that was left unstripped by the upstream foil (less any scattering, i.e. constant factor F_1) measured an average integrated voltage of 14.642 ± 1.860 pVs for runs 7 and 8. The fraction of the H^- beam that is left unstripped by the upstream foil is then given by

$$\frac{H_{uns}^-}{H^-} = \frac{H_{measured}}{H^- \times F_1} = \frac{14.642}{21387.447 F_1} = 0.000685/F_1 \quad (2.40)$$

The fraction of the H^- beam that is left unstripped by the thicker downstream foil will be even smaller, and thus this contribution is considered negligible.

2.8.3 H^- beam and $H_{cont,GS}^0$ beam contamination at FC2

The H^- beam cannot be measured directly at FC2, however it can be inferred using Eq. 2.23 and runs 2 and 14. The H_{cont}^+ beam contamination is only 0.5 % as large as the H^- beam. Therefore, the value H_{cont}^+ can be substituted by $-0.005 H^-$ (the change of state introduces a factor of -1). Rearranging Eq. 2.23 for H^- , then

$$H^- = \frac{H_{measured} - H_{cont,GS}^0}{0.995} \quad (2.41)$$

Runs 3 and 13 measured $H_{cont,GS}^0$ directly, however, the measurements taken before and after the linac failure need to be treated separately. Table 2.6 shows the $H_{cont,GS}^0$ values along with the corresponding H^- values. Using runs 2 and 3, it can be seen that the H^0 beam contamination that is gas stripped between the dipole magnet and the Faraday cup is $1.197 \pm 0.468\%$ of the intensity of the pure H^- beam (Table 2.6). Similarly, for runs 14 and 13 the value is $1.439 \pm 0.600\%$. While the values are consistent with each other, they are still extremely large considering the loss rate due to residual gas stripping is expected to be minimal at the measured vacuum pressure levels ($1-2 \times 10^{-5}$ Pa).

Table 2.6: The first row compares the H^0 beam contamination that is gas stripped as a ratio of the H^- beam intensity before the linac failure, while the second row compares the ratio after the linac failure. The errors are large, as the error due to the variance of the beam intensity is large.

Run	$H_{measured}$	Run	H_{GS}^0	H^-	H_{GS}^0/H^- (%)	Error (%)
2	17994.805	3	-204.816	18292.045	1.197	± 0.468
14	7300.915	13	-107.132	7445.668	1.439	± 0.600

2.8.4 Scattering factors F_1 and F_2

The scattering factor F_2 can be determined using Eq. 2.25 and the value measured for run 5. It however depends on H_{cont}^0 which can be found using Eq. 2.22 and the value measured for run 4. Substituting Eq. 2.22 into Eq. 2.25 and rearranging for F_2 results in

$$F_2 = \frac{H_{measured,Run4} - H_{measured,Run5}}{H^- - H_{cont}^+} = 0.749 \pm 0.487 \quad (2.42)$$

After the beam has passed through the downstream foil, $25.1\% \pm 48.7\%$ of the beam is scattered or deflected beyond the aperture of the Faraday cup. The error is large due to the number of parameters used to calculate F_2 , that suffer from the large error on the variance of the beam intensity.

The scattering factor F_1 can be found in a similar manner using Eq. 2.34 and run 6. The H_{cont}^0 beam contamination is found in the exact same way as for F_2 . The H_{uns}^- component will only make a minute difference and cannot be found directly using FC2. However, a value for H_{uns}^- can be found by assuming its ratio to H^- remains constant for FC2 and FC3, i.e. using Eq. 2.40. Substituting in Eqs 2.22 and 2.40 into Eq. 2.34 and rearranging for F_1 , then

$$F_1 = \frac{H_{measured,Run6} + 0.000685H^- \times F_2}{(H_{cont}^+ - H^-) \times F_2 + H_{measured,Run4}} \quad (2.43)$$

If Eq. 2.42 is rearranged for $(H_{cont}^+ - H^-)$ and substituted into Eq. 2.43, then

$$F_1 = \frac{H_{measured,Run6} + 0.000685H^- \times F_2}{H_{measured,Run5}} = 0.136 \pm 0.081 \quad (2.44)$$

2.8.5 Gas stripping factor F_{GS}

The residual gas stripping factor is simply the ratio between the gas stripped H_{cont}^0 beam contamination and the total H_{cont}^0 beam contamination:

$$F_{GS} = \frac{H_{cont,GS}^0}{H_{cont}^0} = \frac{H_{measured,Run3} \times F_s}{H_{measured,Run4}} = 0.071 \pm 0.063 \quad (2.45)$$

2.8.6 H^0 beam contamination

Having calculated F_2 , the H^0 beam contamination as a function of the initial H^- beam intensity can be calculated. Using Eq. 2.22, then

$$\frac{H_{cont}^0}{H^-} = \frac{H_{measured}}{F_2 \times H^-}. \quad (2.46)$$

Treating the runs either side of the linac failure separately, then for run 4:

$$\frac{H_{cont}^0}{H^-} = -0.157 \pm 0.138 \quad (2.47)$$

The minus sign is due to the different charge states of H^- and the stripped H_{cont}^0 component. Similarly, for run 12:

$$\frac{H_{cont}^0}{H^-} = -0.114 \pm 0.101 \quad (2.48)$$

The two values are consistent with each other, though the reduction in the absolute value of the H_{cont}^0 component can also be explained by the vacuum improving for the latter runs of the experiment.

2.8.7 Calculating the stripping factor F_{PS1}

In Eqs. 2.17 and 2.20, F_{PS1} has a dependence on F_{PSZ} . The dependence is given by:

$$-(1 - F_{PSZ}) \frac{H_{cont}^0}{H^-} \quad (2.49)$$

In a similar way that the expected hydrogen populations were calculated for a H^- beam passing through varying surface densities of carbon in Fig. 2.10, the expected hydrogen populations for a H^0 beam passing through varying surface densities of carbon can also be calculated. For a $9 \mu\text{g}/\text{cm}^2$ foil, this results in a partial stripping factor of 0.992 (or a partial stripping factor of 0.975 when scaled for experimental results).

Using either Eqs. 2.47 or 2.48, and a value of 0.992 for F_{PSZ} , then Eq. 2.49 equates to approximately 0.001 for either scenario (or approximately 0.003 when scaled for experimental results). This means that F_{PSZ} has a negligible contribution on F_{PS1} . In fact, if the contribution were non-negligible, then this would in fact show the partial stripping of a H^0 beam. Eqs. 2.20 and 2.17, then reduce to

$$F_{PS1} = \frac{H_{measured}}{F_1 F_{GS} H^-} + 1 \quad (2.50)$$

$$F_{PS1} = \frac{H_{measured}}{F_1 F_2 H^-} + 1 \quad (2.51)$$

For runs 8 and 10, which measured the partial stripping of a H^- beam to H^0 , that was then subsequently gas stripped, this resulted in values of 0.961 ± 0.043 and 1.233 ± 0.284 respectively. For runs 7 and 11, which measured the partial stripping of a H^- beam to H^0 , that was then subsequently foil stripped, this resulted in values of 0.939 ± 0.057 and 1.017 ± 0.018 respectively. The partial stripping factors are summarised in Table 2.7.

Table 2.7: Summary of the calculated partial stripping factors (F_{PS1}).

Measurement Type	Run	Dipole	Foil 2	Integrated Value	F_{PS1}
Direct (FC2)	8	On	No	-6.875	0.961 ± 0.043
Direct (FC2)	10	Reversed	No	16.770	1.233 ± 0.284
Direct (FC2)	7	On	Yes	-113.456	0.939 ± 0.057
Direct (FC2)	11	Reversed	Yes	13.092	1.017 ± 0.018

2.9 Discussion

The partial stripping factors of Table 2.7 require some consideration. Negative integrated values in Table 2.7 are indicative of H^+ particles, while positive integrated values are indicative of H^- particles (or electrons), though no negative particles should be able to reach Faraday cup 2 in the time interval the integration takes place.

The positive values of runs 10 and 11 and the overall low integrated values of runs 8, 10 and 11 may therefore be due to a number of reasons. As runs 8 and 10 measured residual gas stripping, the integrated values for those runs were expected to be very low in any case. Additionally, runs 10 and 11 were taken after the linac failure and were affected by a lower beam intensity. The measured signals for each run are plotted in Fig. 2.14, with the central black line showing the average of the 10 beam samples.

A capacitance effect can be seen at the end of the beam signal for each of the four runs, with the capacitance effect dissipating between beam samples. A capacitance effect can build during a beam sample, however. If the measured beam signal is already very small, as is the case for runs 10 and 11, then the measured beam signal could turn from a negative to a positive. Additionally, for the integrated signal, only the time interval where the beam fully passed through the upstream foil was considered (0.015 to 0.075 ms).

A positive beam signal, due to negative particles striking FC2 and not due to some capacitance effect, should be impossible as the dipole magnet is either on or has

its polarity reversed during all four runs. The negative particles should be deflected away from FC2.

The Faraday cups could however be affected by noise, though this is considered unlikely. When the beam was switched off, FC2 and FC3 measured integrated voltages of less than 1 pVs. When the beam was switched on, however, FC3 was able to measure a small voltage even for runs where the dipole magnet was switched off. The average integrated signal for those runs (runs 2, 5, 6 and 9) was 7.401 ± 20.610 pVs. The value is small and may be due to stray particles striking the Faraday cup. It is however unlikely for stray particles to strike FC2 for runs 7, 8, 10 and 11, as the dipole magnet was switched on and should be able to deflect any stray particles. The small measured beam signals (<17 pVs), as well as the capacitance effect, mean the partial stripping factors for runs 8, 10 and 11 are considered untrustworthy. For run 7, the measured beam signal is large (>100 pVs), and will be considered for the rest of the analysis.

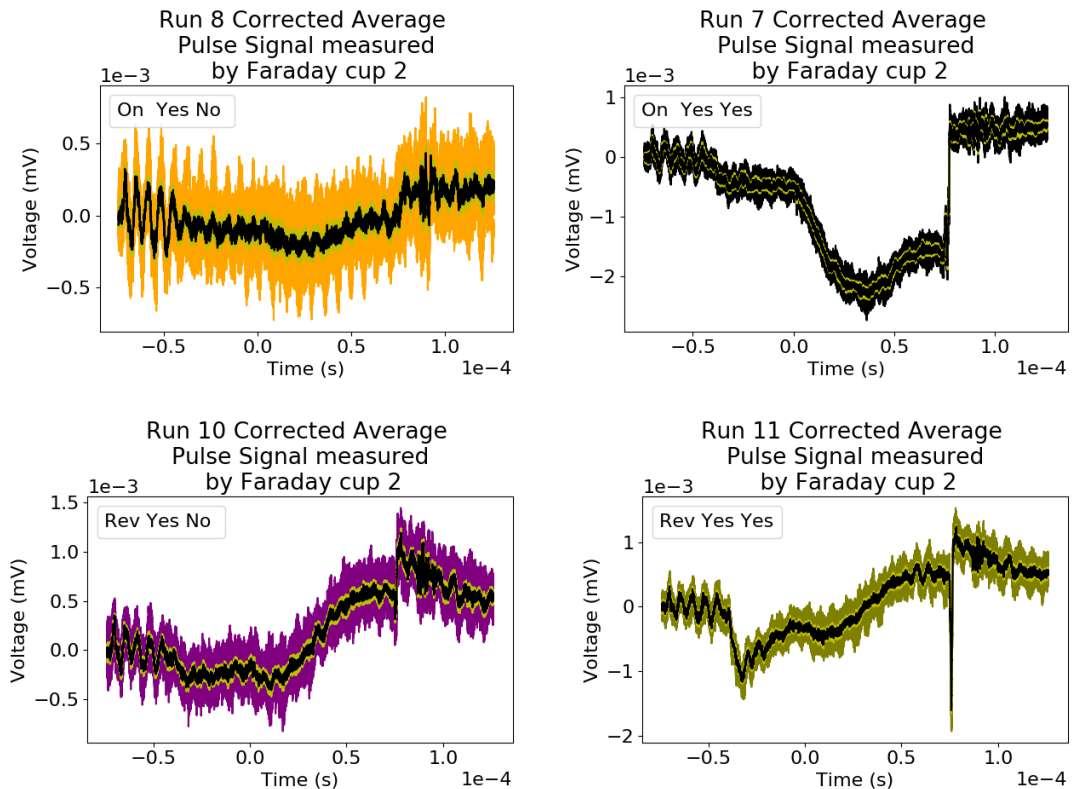


Figure 2.14: The voltage measured at Faraday cup 2. For runs 7 and 8 the dipole magnet was switched on, while for runs 10 and 11 the dipole magnet had its polarity reversed. All four runs had the upstream foil present, while only runs 7 and 11 had the downstream foil present (Dipole: On/Reversed, Foil 1: Yes, Foil 2: Yes/No).

2.9.1 Electron loss cross-sections at 11 MeV

The electron loss cross-sections can be obtained using Eqs. 2.5, 2.6 and 2.7. The double electron loss cross-section σ_{-11} is typically only a small fraction of the single electron loss cross-section σ_{-10} and will be considered negligible^{105,132}. The electron loss cross-sections can be calculated for two components, the unstripped H^- component and the partially stripped H^0 component.

Unstripped H^- component

Due to the thickness of the foil used, the measured unstripped component of the beam is very small (<15 pVs) and carries a significant uncertainty. The cross-section σ_{-10} can however still be calculated using Eqs. 2.5 and 2.40, finding that

$$e^{-\sigma_{-10}x} = N^- = \frac{0.000685}{F_1} = 0.005 \pm 0.004 \quad (2.52)$$

Taking the natural log of both sides and rearranging for σ_{-10} , then

$$\sigma_{-10} = \frac{-\ln(0.005)}{x} \quad (2.53)$$

For a foil with a surface density of $9 \mu\text{g}/\text{cm}^2$, the number of atoms per square centimetre, x , is given by

$$x = \frac{\text{Foil Surface Density} \times \text{Avogadro's number}}{\text{Carbon Molar Mass}} = 4.51258 \times 10^{17} \text{ atoms}/\text{cm}^2 \quad (2.54)$$

Using Eq. 2.53, the σ_{-10} electron loss cross-section calculated for the unstripped H^- component then becomes

$$\sigma_{-10} = 1.17_{+0.41}^{-0.13} \times 10^{-17} \text{ cm}^2 \quad (2.55)$$

Partially stripped H^0 component

The σ_{-10} and σ_{01} electron loss cross-sections can be calculated using Eq. 2.6. Assuming that σ_{-11} has a negligible contribution, then

$$N^0 = \frac{\sigma_{-10}}{\sigma_{-10} - \sigma_{01}} [e^{-\sigma_{01}x} - e^{-\sigma_{-10}x}] \quad (2.56)$$

Table 2.2 noted that for experimental values, the ratio between the σ_{-10} and σ_{01} cross-sections remained constant, provided the projectile energy is sufficiently large. For a hydrogen anion striking a carbon atom, a projectile energy of at least 1 MeV

should be sufficient based on Fig. 2.5. In Table 2.2, the average σ_{-10}/σ_{01} ratio is 0.395. Substituting this value into Eq. 2.6, then

$$N^0 = \frac{1}{0.605} [e^{-0.395\sigma_{-10}x} - e^{-\sigma_{-10}x}] \quad (2.57)$$

From Eq. 2.7, the H^0 population can be written in terms of the H^+ and H^- populations:

$$N^0 = 1 - N^- - N^+ \quad (2.58)$$

Eq. 2.57 can then be rewritten as

$$0 = \frac{1}{0.605} [e^{-0.395\sigma_{-10}x} - e^{-\sigma_{-10}x}] - (1 - N^- - N^+) \quad (2.59)$$

The stripped H^+ component corresponds to the partial stripping factor F_{PS1} (providing the unstripped H^- component is very small). The σ_{-10} cross-section in Eq. 2.59 can then be found using the Newton-Raphson-Simpson method¹³⁵⁻¹³⁸.

For Eq. 2.59, f and f' are given by

$$f(\sigma_{-10}) = \frac{1}{0.605} [e^{-0.395\sigma_{-10}x} - e^{-\sigma_{-10}x}] - (1 - N^- - N^+) \quad (2.60)$$

$$f'(\sigma_{-10}) = \frac{1}{0.605} [-0.3905xe^{-0.395\sigma_{-10}x} + xe^{-\sigma_{-10}x}] \quad (2.61)$$

Using an initial guess of $2.0 \times 10^{-17} \text{ cm}^2/\text{atom}$ for the σ_{-10} cross-section and a partial stripping factor of $F_{PS1} = 0.939 \pm 0.057$, the σ_{-10} cross-section calculated for the partially stripped H^0 component is given by $\sigma_{-10} = 1.896_{+1.483}^{-0.401} \times 10^{-17} \text{ cm}^2/\text{atom}$. Including the stated error¹³³ on the foil thickness of 10%, the σ_{-10} cross-section then becomes

$$\sigma_{-10} = 1.896_{+1.859}^{-0.536} \times 10^{-17} \text{ cm}^2/\text{atom}. \quad (2.62)$$

The corresponding σ_{01} cross-section then becomes

$$\sigma_{01} = 0.749_{+0.734}^{-0.212} \times 10^{-17} \text{ cm}^2/\text{atom}. \quad (2.63)$$

The σ_{-10} cross-section compares well with the value predicted by Gillespie ($\sigma_{-10} = 2.703 \times 10^{-17} \text{ cm}^2/\text{atom}$) and the cross-section scaled for experimental results ($\sigma_{-10} = 2.067 \times 10^{-17} \text{ cm}^2/\text{atom}$). The 11 MeV σ_{-10} cross-section measured during this experiment has been plotted in Fig. 2.15 and appears to be consistent with the other experimental results and the value predicted by Gillespie.




Figure 2.15: The σ_{-10} electron loss cross-section for a hydrogen anion striking a carbon atom as a function of projectile energy. The red line corresponds to the expected cross-sections as a function of projectile energy by Gillespie. The values at 181 MeV, 200 MeV and 800 MeV correspond to the experimental results by Saha¹²⁰ et al., Webber and Hojvat¹⁰⁵ and Gulley¹³² et al. The cross-section at 11 MeV corresponds to the result of this experiment.

Systematic effects of the MICE IH₂ absorber

In the Muon Ionization Cooling Experiment (MICE), a muon beam was passed through a liquid hydrogen absorber to demonstrate ionisation cooling. The emittance of a beam changes when passing through an absorber material and is described by the ionisation cooling equation

$$\frac{d\varepsilon_{\perp}}{ds} \cong -\frac{\varepsilon_{\perp}}{\beta^2 E_{\mu}} \left\langle \frac{dE}{ds} \right\rangle + \frac{\beta_{\perp} (13.6 \text{ MeV})^2}{2\beta^3 E_{\mu} m_{\mu} X_0}. \quad (3.1)$$

where $\frac{d\varepsilon_{\perp}}{ds}$ is the rate of change of normalized transverse emittance (ε_{\perp}), β , E_{μ} and m_{μ} are the muon velocity energy and mass, $\frac{dE}{ds}$ is the magnitude of the mean energy loss rate through ionization, X_0 is the absorber radiation length and β_{\perp} is the transverse beta function at the absorber⁵⁴. In Eq. 3.1, the negative term is the cooling effect which reduces the emittance of the beam, while the positive term is the heating effect which increases the emittance of the beam.

The systematic uncertainties associated with the liquid hydrogen absorber described in this chapter follow on from the work by Green and Yang¹³⁹. These include the systematic uncertainties on the density of liquid hydrogen at different temperatures and pressures, the contraction of the absorber vessel as it is cooled down to cryogenic temperatures and the variation of the absorber vessel window thicknesses.

The absorber vessel was set at the centre of the Focus Coil magnet (Fig. 3.1). This reduces the effect of scattering by focusing the beam on the absorber. The absorber vessel has an inner diameter of 300 mm and a length between its end flanges of 230 mm. The length between the two curved absorber vessel windows along the central axis is 350 mm.

Hydrogen was chosen as an absorber as it is a low Z material which suffers less scattering for a given energy loss. The liquid hydrogen is contained in an aluminium




Figure 3.1: The liquid hydrogen absorber vessel surrounded by the absorber focus coil^{140,141}.

absorber vessel and as a high Z material, the aluminium vessel worsens the cooling performance (Eq. 3.1). The aluminium windows are thinnest at the centre of the beamline where the beam density is greatest.

3.1 Absorber vessel shape

The aluminium absorber vessel is cooled from room temperature to the operating temperature of the experiment (≈ 20 K), which results in the vessel contracting. The linear contraction of Al-6061 as it is cooled from 293 K is given by

$$\alpha = -4.1277 \times 10^{-3} - 3.0389 \times 10^{-6}T + 8.7696 \times 10^{-8}T^2 - 9.9821 \times 10^{-11}T^3 \quad (3.2)$$

where T is the operating temperature¹⁴². The equation is a line of best fit of data collated by NIST (National Institute of Standards and Technology) and has an associated curve fit error of 4%.

At the MICE operating temperature, there is a linear contraction of the vessel along each plane of 0.415% (293 K \rightarrow 20 K), resulting in a warm bore length (350 mm) contraction of $1.4525 \text{ mm} \pm 4\%$. The vessel is held suspended in place, meaning the vessel is free to contract along each plane without restriction, ensuring there are no forces created to distort the shape of the vessel.

To minimise energy loss and Coulomb scattering by the absorber vessel, the windows are kept as thin as possible. However, they must not rupture when handling any internal pressure they are subjected to. For safety considerations^{140,141}, it is

necessary for the liquid hydrogen circuit to be pressurised above atmospheric pressure to prevent air ingress. The vessel must also be capable of handling up to 1.5 bar, the relief valve set pressure.

These pressures result in a deflection of the absorber windows and were modelled by Green and Yang using ANSYS¹³⁹. The uncertainty in the model's window deflection was 20%. It did show a linear expansion of the window deflection with pressure up to 2 bar when the windows began to yield. The pressure sensors were accurate to ± 5 mbar (0.25% of 2 bara full scale). At 1085 ± 5 mbar, the typical MICE operating pressure, this corresponds to a deflection of $0.5374 \text{ mm} \pm 0.1076 \text{ mm}$ (model uncertainty) $\pm 0.0022 \text{ mm}$ (sensor uncertainty) at the centre of the absorber window.

For the MICE experiment, the contraction due to cooling was 1.4525 mm ($\pm 4\%$), with each window deflecting $0.5374 \pm 0.004 \text{ mm}$ ($\pm 20\%$) at 1.085 bar. This results in a combined effect of:

$$1.4525 (\pm 0.0581) - 2(0.5374 (\pm 0.1098)) = 0.3777 \pm 0.1629 = 0.4 \text{ mm} \pm 0.2 \text{ mm} \quad (3.3)$$

3.2 Absorber vessel window thicknesses

The amount of energy loss and cooling experienced by a muon passing through the absorber depends on the amount of aluminium and liquid hydrogen traversed. There are four windows, two absorber wall windows of the vessel and two safety windows (Fig. 3.1).

The windows¹⁴⁰ show some variation in thickness. At the centre of the absorber, the total amount of aluminium the muon beam passes through is 785 ± 13 microns, a variance of 1.68%. However, as the windows are thin, the effects on energy loss are negligible. A 200 MeV muon passing along the central axis of an empty absorber vessel loses 0.345 MeV, which introduces a 0.006 MeV uncertainty on the energy loss.

Table 3.1: The measured absorber window thicknesses compared to their design at centre of absorber.

At centre of absorber	Measured (μm)	Design (μm)
Safety window 1	197 ± 8	210
Absorber window	174 ± 5	180
Absorber window	184 ± 2	180
Safety window 2	230 ± 9	210
Total	785 ± 13	780

3.3 Variation of the density of liquid hydrogen due to varying temperature and pressure

The energy lost by a muon travelling through the liquid hydrogen absorber depends on the path length the muon travelled through and on the density of the liquid hydrogen. The temperature was recorded by eight LakeShore Cernox 1050 SD sensors¹⁴³ (Fig. 3.2). Four of the sensors were used solely as temperature sensors, while the other four were used as both temperature and level sensors. The level sensors were used when the absorber vessel was being filled to know how much liquid hydrogen was in the vessel and during the experiment to ensure the liquid hydrogen reached the top of the vessel.

They were arranged in pairs with two mechanically clamped at the top of the vessel, two at a rotation of 45°, a further two at a further rotation of 90° and a final two at a further rotation of 45° to be at the bottom of the vessel. The temperature sensors were labelled TSA, TSB, TSD and TSE from top to bottom, while the level sensors were labelled LSA, LSB, LSD and LSE from top to bottom.

The sensors have a typical sensor accuracy of ± 9 mK and a long-term stability of ± 12 mK at 20 K¹⁴⁴⁻¹⁴⁶. The magnetic field dependent temperature error at 2.5 T is 0.04% $\Delta T/T$, equivalent to ± 8 mK at 20 K¹⁴⁴⁻¹⁴⁸. These are the quoted uncertainties given by the manufacturer of the sensors. The importance of magnetic fields on temperature measurements is that they cause reversible calibration shifts. When the magnetic field is removed, the sensors return to their original calibration.

To reduce the uncertainty in the liquid hydrogen density, a calibration procedure was devised using the boiling point. It is assumed that after cooldown and before venting, the liquid hydrogen absorber was kept in equilibrium conditions in steady-state. During this time, the pressure readings remained steady at 1.085 Bar.

Cooldown and liquefaction were completed slowly over eight days until the 25 September 2017 at a pressure of 1.15 Bar, after which the vessel's pressure was lowered to 1.085 Bar and stabilised during the early hours of the 26 September 2017^{140,143}. The vessel then remained in this steady-state until the 16 October 2017 when the venting process began. During this process, the coldhead was switched off and the heaters were switched on, delivering a nominal power of 50 Watt to the absorber vessel. This resulted in an increase in pressure up to the relief valve set pressure (1.505 Bar) and an increase in temperature until it stabilised at the boiling temperature. At this temperature, the liquid hydrogen turned to gas and began emptying from the vessel. A rapid increase in temperature followed once all the liquid hydrogen had boiled off.

During venting at that increased pressure, the temperature remained near constant

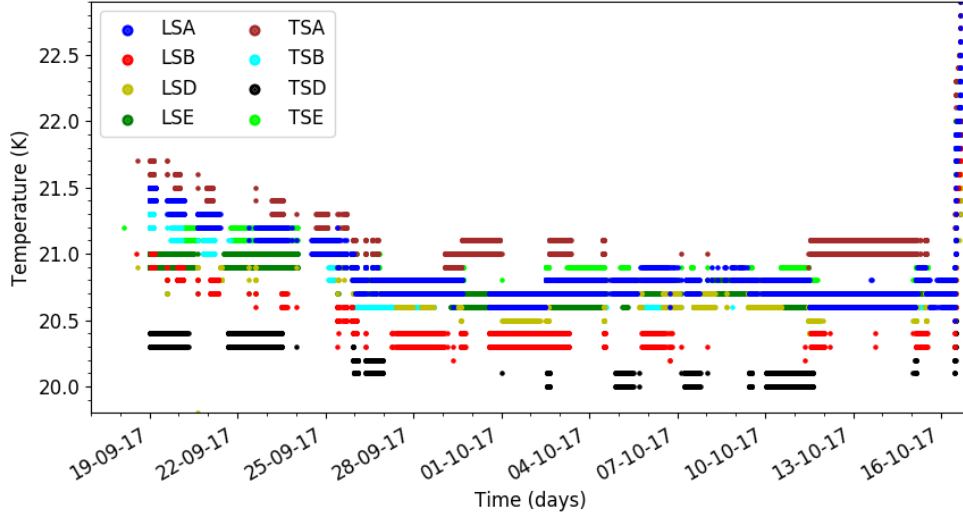


Figure 3.2: Raw readings in the 19.8 to 22.9 K range. In the steady state period from the 26 September 2017 to the 16 October 2017 the temperature readings from the sensors agree to within 1 K.

for four hours. The sensors at the top did show some step increases after about two hours, as the liquid hydrogen in the top of the vessel begins to boil off first, until after four hours the liquid at the bottom of the vessel also boils off and there was a rapid increase in temperature. This is consistent with the time required for a 50 W heater to boil 22 litres of liquid hydrogen. The boiling temperature at the pressure the liquid hydrogen boiled off is 21.692 K. The Cernox 1050 SD temperature sensors are hermetically sealed resistive sensors housed in helium gas. Even though the sensors do not follow a standard response curve, their dimensionless sensitivity $(T/R) \times (dR/dT)$ should remain near constant in our region of interest (19 K to 22 K)¹⁴⁴. However, when Ishimoto et al.¹⁴⁹ performed tests at KEK using the MICE absorber vessel and Cernox temperature sensors at the vapour pressure of liquid hydrogen, they found a small non-linear behaviour in the sensors, both in the gaseous and in the liquid phase. To compare the temperatures across this non-linear region, a correction of 0.06 ± 0.03 K must be made. This is the average temperature the sensors tested by Ishimoto differed by when comparing the vapour temperature for that pressure during boil-off and when the vessel was kept in steady-state (Fig. 3.3).

For data storage considerations, the temperature readings were truncated at a granularity of 0.1 K, with a temperature reading recorded if there was a change in this truncated value. Approximately 100 points were removed from this data during the analysis. There were several occasions when the temperature reading dropped to 4 K from approximately 20 K before returning to the original value, which is not consistent with the physical properties of the absorber.

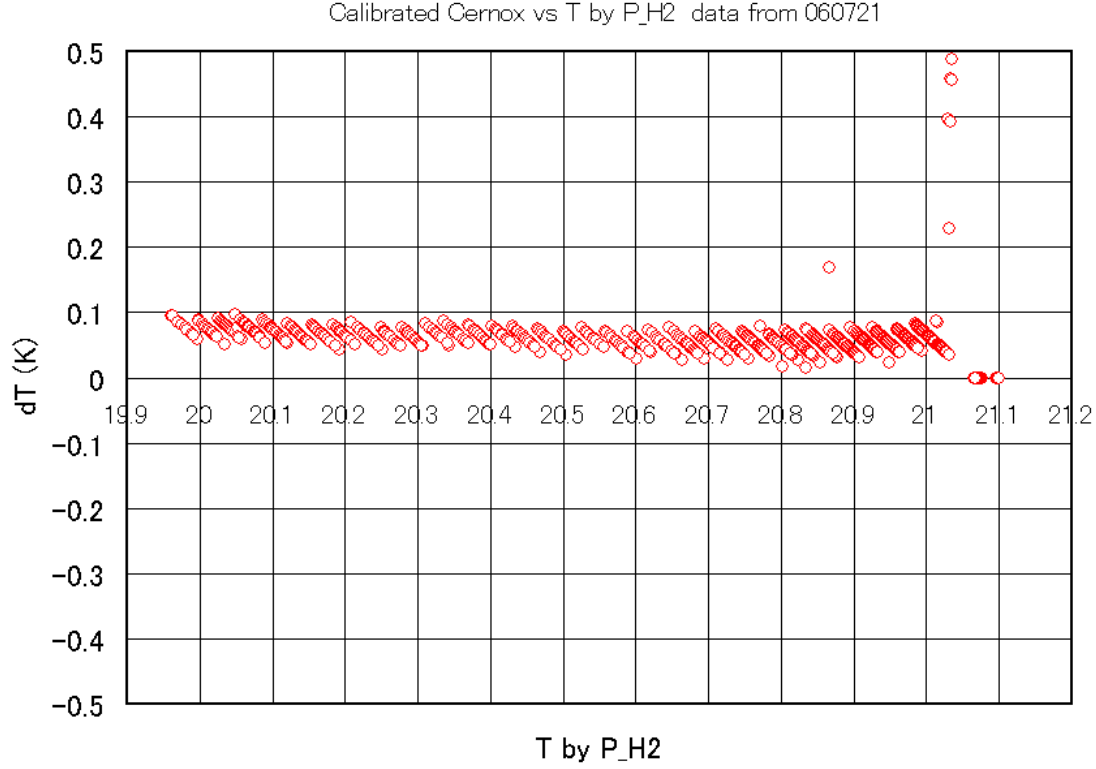


Figure 3.3: Deviation from the expected Vaporisation temperature for the Cernox temperature sensors as the Vaporisation temperature is lowered.

To compare the temperature and pressure readings, a weighted mean was used. In Eqs. 3.4 and 3.5, $T_{average}$ is the average temperature for that time-interval, $T_{previous}$ is the last temperature reading from the previous time-interval and T_{last} is the last temperature reading from the current time-interval. T_i is every temperature reading in that interval bar T_{last} , as it has no corresponding Δt_i . Δt_i is the time between the current temperature reading and the following temperature reading within the time-interval. Δt_{first} is the time between the start of the time-interval and the first temperature reading, while Δt_{last} is the time between the last temperature reading and the end of the time-interval. The time-interval chosen will then also be a sum of all those Δt . Applying the weighted average results in a smoothing of the data (Figures 3.4 and 3.5).

$$T_{average} = \frac{T_{previous}\Delta t_{first} + \sum_i T_i\Delta t_i + T_{last}\Delta t_{last}}{t_{interval}} \quad (3.4)$$

$$t_{interval} = \Delta t_{first} + \sum_i \Delta t_i + \Delta t_{last} \quad (3.5)$$

Eq. 3.6 details the calibration procedure used to calibrate the temperature sensors. The sensors are based on a calibration curve and thus require three points




Figure 3.4: Temperature sensors during boil-off period on the 16 October 2017. The pressure has been scaled by a factor of 60 to allow comparison between the temperature and pressure sensors at the same moments in time.




Figure 3.5: The temperature readings of the eight sensors are shown with the polarity of the focus coil. Grey areas indicate when the focus coil was operated in solenoid mode, while red areas indicate flip mode. Yellow areas indicate when all the magnets are turned off. The white gaps indicate times when the magnets were ramped up or down and times during which no run data was taken.

to perform the calibration. However, the temperature sensors only recorded one calibration point, the boiling point of liquid hydrogen. The calibration procedure is then only valid for temperatures close to that calibration point, and will begin to deviate further away from that calibration point. The calibration procedure is sufficient as during boil-off and steady-state the liquid hydrogen temperatures are at or very close to this calibration point.

The corrected temperature reading is found by calculating and then applying the cut-off correction (as the data is truncated), sensor non-linearity correction, magnetic field correction and then the boiling point scaling factor. The magnet and temperature constants will be calculated. $c_{cut-off}$ is 0.05, $c_{non-linearity}$ is 0.06, I is the focus coil current, while $T_{reading}$ is the temperature value recorded by the sensor.

$$T_{corrected} = \frac{T_{reading} + c_{cut-off} - c_{non-linearity} - c_{magnet}I}{c_{Temperature}} \quad (3.6)$$

The Cernox temperature sensors showed a magnetic field dependence (Fig. 3.5). Two magnet coefficients for Eq. 3.6 were calculated based on the focus coil current, one for each mode the coil was operated in, solenoid and flip, with the correction factors in Table 3.2 corresponding to a linear best fit of the current as a function of temperature during the steady-state period.

Table 3.2: The focus coil current correction coefficients were calculated by plotting the temperature against current as the magnets were ramped up and down for each mode, solenoid and flip. The accuracy of the coefficients is limited by the 0.1 K resolution.

Mode	LSA	LSB	LSD	LSE
Solenoid	3.9424E-4	4.6810E-4	1.2207E-3	5.7725E-5
Flip	5.7000E-4	-6.8095E-4	9.2727E-4	2.0350E-4
Mode	TSA	TSB	TSD	TSE
Solenoid	7.1284E-5	2.8417E-4	4.2315E-4	3.7478E-4
Flip	-3.9965E-4	-6.7594E-4	-1.8550E-4	6.4080E-4

Hydrogen exists as a mixture of parahydrogen and orthohydrogen^{150,151} whose concentration of each varies depending on temperature and time to reach an equilibrium state. At atmospheric pressure, the vaporisation temperature for parahydrogen is 20.271 K, while for orthohydrogen it is 20.380 K. When the hydrogen mixture is cooled, the hydrogen mixture begins to occupy the lowest energy states. This results in a conversion of orthohydrogen to parahydrogen¹⁵². At 20 K nearly all orthohydrogen has been converted to parahydrogen, resulting in a parahydrogen concentration of over 99%. This means the vaporisation temperature of liquid hydrogen is nearly identical to that of parahydrogen.

Table 3.3: The temperature coefficient scaling factor for each sensor, calculated by dividing the temperature reading (adjusted for cut-off coefficient and magnetic field) by the vaporisation temperature at that pressure.

Mode	LSA	LSB	LSD	LSE
$T/T_{Boiling}$	1.010697372	0.989359131	1.003485436	1.00854641
Mode	TSA	TSB	TSD	TSE
$T/T_{Boiling}$	1.027887842	1.003816982	0.978539237	1.015640458




Figure 3.6: The temperature readings after the calibration procedure has been applied. The procedure is limited by the 0.1 K granularity of the truncated temperature readings.

The temperature coefficient in Eq. 3.6 is found by dividing the boiling temperature reading for that sensor by the boiling temperature of parahydrogen at that pressure. The boiling temperature reading was determined as the average temperature from the moment in time the sensors stopped increasing in temperature after the heaters were turned on, until the sensors at the top of the vessel began to increase in temperature again as the liquid hydrogen boiled off. The ratio gives the scaling factor all temperature readings for that sensor should be adjusted by. The pressure sensors also have a ± 5 mbar uncertainty. This results in a 0.014 K uncertainty in the boiling temperature of liquid parahydrogen and subsequently a 0.0006 uncertainty in the scaling factors.

The scaling factor is only appropriate for the liquid phase, as in the gaseous phase the absorber vessel and gas are not in steady state any more. The gas and vessel temperatures rise, but not at the same rate, resulting in localised temperature gradients and thus differing temperature readings.




Figure 3.7: The purple line depicts the boiling temperature at the current pressure in the absorber vessel. The temperature readings have been scaled by the boiling temperature at 1.505 Bar, but do not include the non-linear sensor step correction across the non-linear range, as these are the temperatures and pressures going through the non-linear range. The temperature rises at close to the boiling temperature within the 0.1 K resolution, showing the steady state liquid hydrogen was close to the boiling point at all times.

The boiling temperature at 1.085 Bar is 20.504 K, with the corrected sensor readings slightly higher. There are however a number of uncertainties. The readings are recorded with a granularity of 0.1 K. The sensors add another 17 mK uncertainty (9 mK accuracy + 12 mK stability + 8 mK magnetic field), although the magnetic field error is likely greater. The sensors have a non-linear step between the steady-state and boil-off pressures, adding a 0.03 K uncertainty. The temperature scaling and magnet current correction factors have an associated error, as they are based on the 0.1 K resolution. For example, a calibrated sensor at the boiling temperature and at a pressure of 1.505 Bar should read 21.692 K but can only read 21.65 K (21.6 K cut-off plus 0.05 cut-off correction) i.e. it is off by 0.042 K. The ± 5 mbar uncertainty adds another uncertainty to the temperature calibration constants of ± 0.014 K. Collectively, all these uncertainties add up to 0.177 K ≈ 0.2 K for each sensor.

Fig. 3.7 shows the boiling temperature for a given pressure when the heaters were turned on from the steady state region to the relief valve pressure. The corrected temperature readings follow this line closely, showing that before the heaters were turned on, the liquid hydrogen temperature was already close to the boiling point.

Knowing that in our steady state condition the liquid hydrogen was close to the boiling temperature of liquid parahydrogen at $20.5 \text{ K} \pm 0.2 \text{ K}$ and 1.085 Bar allows us to determine the uncertainty in the liquid hydrogen density as $70.57 \text{ kg/m}^3 \pm 0.24 \text{ kg/m}^3$.

3.4 Total systematic uncertainty on energy loss

In total, there are three main contributions to the systematic uncertainty of the liquid hydrogen absorber on energy loss. The contraction of the absorber and deflection of the absorber window due to internal pressure reduces the central warm bore length by $0.4 \pm 0.2 \text{ mm}$. The combined absorber window thickness variation at the centre of the absorber is 13 microns. The temperature during the steady state period of the experiment when the pressure remained constant at $1085 \pm 5 \text{ mbar}$ is $20.5 \pm 0.2 \text{ K}$ for each sensor. Combining the eight sensors gives an average temperature of $20.51 \pm 0.06 \text{ K}$ (Fig. 3.8) during the steady-state period, corresponding to a liquid hydrogen density of $70.55 \pm 0.07 \text{ kg/m}^3$.




Figure 3.8: Average Temperature of the eight Cernox sensors after the calibration procedure during the steady-state period when the pressure and temperature were being kept constant.

The energy loss is momentum dependent, as each particle will lose a different amount of energy passing through the absorber. Tables 3.4 and 3.5 show the energy loss at various momenta and densities of aluminium and liquid hydrogen^{153–156}. 277 MeV and 344 MeV are the minimum ionization momenta of aluminium and liquid hydrogen, respectively.

Table 3.4: Energy loss for aluminium (Al-6061) at various momenta with a density of 2.699 g/cm³.

Momentum (MeV)	100	140	200	277
Mass Stopping Power (MeV g ⁻¹ cm ²)	1.798	1.688	1.630	1.615
Stopping Power (MeV cm ⁻¹)	4.8528	4.556	4.3994	4.3589

Table 3.5: Energy loss for liquid hydrogen at various densities (0.07048 to 0.0708 g/cm³) and various momenta of muons.

	Density	100	140	200	344
Momentum					
Mass Stopping Power		4.568	4.267	4.104	4.034
Stopping Power	0.07048	0.3220	0.3007	0.2892	0.2843
Stopping Power	0.07055	0.3223	0.3010	0.2895	0.2846
Stopping Power	0.07062	0.3226	0.3013	0.2898	0.2849
Stopping Power	0.0708	0.3234	0.3021	0.29056	0.2856

During the MICE experiment 140, 170, 200 and 240 MeV momenta muon beams were used. The energy loss and its uncertainty were then calculated. The calculation used a central bore length of 349.6 ± 0.2 mm, a total window thickness of 0.785 ± 0.013 mm and a liquid hydrogen density of 70.55 ± 0.07 kg/m³ for a particle travelling straight through the centre of the absorber.

For a 140 MeV muon particle this corresponds to an energy loss of 10.86 ± 0.02 ($\pm 0.2\%$) MeV, while for a 200 MeV muon particle this corresponds to an energy loss of 10.45 ± 0.02 ($\pm 0.2\%$) MeV. In terms of energy loss, the systematic error is 0.2%. This is for a particle travelling along the central axis of the absorber. An actual muon travelling through the absorber with a magnetic field will take a different path and thus have a different path length of aluminium and liquid hydrogen traversed.

Uncertainty of the MICE magnetic field

Both the MICE reconstruction and the demonstration of ionization cooling assume a uniform solenoid field within the tracker regions of the MICE experiment. However, the magnetic fields will never be truly uniform and will therefore carry an associated uncertainty. The magnetic fields of the MICE spectrometer solenoids were extensively measured and modelled by Langlands⁷⁶. This was done using a mapping machine which travelled along the central axis of the spectrometer solenoid and measured the magnetic field using seven Hall probes. The machine travelled in increments of between 2.5 and 5 centimetres, up and down the central axis of each solenoid, making a 20-degree rotation at each end, until a full 360-degree rotation had been completed. This allowed for an extensive mapping of the transverse and longitudinal components of the magnetic fields in the spectrometer solenoids for various magnetic field configurations.

4.1 Dimensions of the spectrometer solenoid coils

Each spectrometer solenoid (Fig. 4.1) consisted of five superconducting coils that were wound using copper matrix niobium-titanium wires (Cu:NbTi) on a 2544 mm long and 711 mm diameter aluminium mandrel housed in a helium bath^{76,157} (This also forms the 'cold mass'). The five coils are labelled as centre coil (CC), end coil (E1 and E2) or match coil (M1 and M2) (Fig. 4.1). The centre and end coils ensured that within each spectrometer, a uniform solenoid field could be created, while the match coils helped to transport the particles from the upstream solenoid to the downstream solenoid.

Table 4.1 gives the warm bore dimensions of the upstream spectrometer solenoid and were the measurements used in MAUS (the outer radius and depth dimensions




Figure 4.1: *Left:* The spectrometer solenoid consists of five NbTi coils wound on to an aluminium mandrel. *Right:* The coils and mandrel sit within the cryostat, which is cooled to 4 Kelvin using liquid Helium⁷⁶.

of the downstream spectrometer solenoid differed by no more than ± 0.7 mm). Using liquid helium, the coils were cooled to approximately 4 Kelvin during the experiment and therefore thermally contracted. The expected dimensions of the cooled upstream spectrometer solenoid⁷⁶ are also given in Table 4.1.

Table 4.1: The warm and cold bore dimensions (in mm) of the upstream spectrometer solenoid coils⁷⁶. The centre of each coil is in reference to the longitudinal position along the mandrel. Each coil is made of a number of layers of copper matrix niobium-titanium wires, with each layer containing a specific number of turns of wires.

Coil	Inner Radius	Outer Radius	Depth	Length	Centre	N_{layers}	N_{turns}
M1 (warm)	258.0	304.1	46.1	201.2	124.0	42	115
M1 (cold)	257.04	301.5	44.5	200.6	123.666		
M2 (warm)	258.0	288.9	30.9	199.4	564.0	28	114
M2 (cold)	256.98	286.6	29.6	198.9	562.116		
E1 (warm)	258.0	318.9	60.9	110.6	964.0	56	64
E1 (cold)	257.09	316.5	59.4	110.2	960.566		
CC (warm)	258.0	280.1	22.1	1314.3	1714.0	20	768
CC (cold)	257.12	278.3	21.2	1310.1	1708.216		
E2 (warm)	258.0	325.7	67.7	110.6	2464.0	62	64
E2 (cold)	257.12	322.9	65.7	110.2	2455.766		

4.2 Centre position of the match coils

The E1, E2 and CC coils shared a power supply. However, to generate a uniform solenoid field across the whole tracker region, the end coils were supplied with additional trim power supplies that could be used to adjust the magnetic fields generated by those coils. When the spectrometer solenoids were being trained in September 2015, they underwent a number of quenches, where the magnets lost their superconductivity and became resistive. This is typically due to a sudden temperature increase. For safety considerations, the trim power supplies weren't used any more (the effect on the uniformity of the magnetic field within the tracker regions will be explored in section 4.5). Additionally, the leads supplying current to the M1 coil of the downstream spectrometer solenoid failed. This meant the M1 coil of the downstream solenoid couldn't be used any more.

Langlands⁷⁶ field mapping study was also used to find the centre of the remaining match coils by powering those coils individually and finding the location where the axial field component was largest. For the upstream solenoid, the centre positions for the M1 and M2 coils are given in Table 4.2. The distance between the centres of the match coils was only 435.83 mm, and is smaller than the warm (440.00 mm) and cold (438.45 mm) bore distances seen in Table 4.1. Additionally, the measured positions of the centre of the M1 and M2 coils from the field mapping study can be compared to the assumed positions from the positional surveys of the spectrometer solenoids that were recorded in the parent geometry files (PGF). Table 4.2 shows that both M2 coils were at least 14 mm closer to the absorber in the magnetic field mapping study than the values assumed in the parent geometry file for the last run taken before the field mapping study (Run 10603).

Table 4.2: The blue shaded cells show the centre of the M1 and M2 coils in the global MICE reference frame (z) from Langlands⁷⁶ field mapping study. The centres for the E1, E2 and CC coils were found relative to the M2 coil using the values from Table 4.1. Comparison is made to the assumed positions in the parent geometry file (PGF) of the last data taking run taken before the field mapping study (Run 10603).

		M1 (mm)	M2 (mm)	E1 (mm)	CC (mm)	E2 (mm)
SSU	z	16106.55	15670.49	15270.29	14519.90	13769.50
SSU	PGF	16095.369	15655.724	15255.724	14505.739	13755.735
SSD	z	17797.05	18237.34	18637.59	18987.81	20138.54
SSD	PGF	17813.615	18253.2598	18653.2597	19403.2444	20153.249

The discrepancy between the field mapping study and the positional survey could be explained by several factors. For example, the forces generated by the magnetic fields could have caused the spectrometers to move. Tarrant¹⁵⁸ however noted that

most of the forces generated by the magnetic fields should have cancelled each other out, while the ground restraints on the spectrometer solenoids should have prevented any large movements. This was seen in a positional survey on the 25 August 2016 where the solenoids moved by only 1-1.5 millimetres, when the spectrometers were operating with a 3 Tesla magnetic field.

To mitigate the solenoids' stray magnetic field from tampering with any of the electrical equipment, a low carbon ($< 0.01\%$) steel Partial Return Yoke (PRY) was installed around the solenoids^{159,160} (left of Fig. 4.2). Outside of the PRY, draw wires were installed (right of Fig. 4.2) that connected to the top of the solenoids via a pulley system. Fig. 4.3 shows the movement of the upstream tracker between the 19 September 2017 and 27 October 2017, when MICE ran under various magnetic field and absorber configurations.




Figure 4.2: *Left:* A low Carbon ($< 0.01\%$) steel Partial Return Yoke (PRY) surrounded the solenoids in the MICE cooling channel¹⁶¹. The magnetic field produced by the solenoids was affected by both the steel and the non-circular shape of the PRY⁷⁶. *Right:* The movement of the solenoids and the PRY was monitored using several draw wires.

The spectrometer solenoids generally moved by no more than 0.5 mm. The small movements of the spectrometers appeared to coincide with changes to the magnetic field currents. It is also apparent that the spectrometer solenoids moved apart from each other or towards each other in some other instances. These instances typically coincided with the strength of the magnetic field changing or when the cooling channel switched from 'solenoid' mode to 'flip' mode and vice versa. In 'solenoid' mode, the upstream and downstream spectrometer solenoids operated with the same polarity magnetic fields, whereas in 'flip' mode they operated with the opposite polarities. Unfortunately, when the magnetic field mapping study was conducted, no draw wire data was collected. To allow the mapping machine to travel along the central axis of the spectrometer solenoids required the focus coil to be removed. This meant the spectrometer solenoids were physically less constrained. Additionally, each

spectrometer, and sometimes each coil, was powered individually. With no counter forces in place, this may lead to some movement of the spectrometers.

However, the movement of the spectrometers is likely to be far too small to explain why the M2 coils were closer to the absorber in the field mapping study than in the positioning survey. A more likely scenario could be the movement of the 'cold mass' in the spectrometer solenoid itself (i.e. the movement of the coils and mandrel in the liquid helium bath). In each spectrometer solenoid, the cold mass was only held suspended in place by eight fibre/epoxy bands with aluminium links⁷⁶. The cold mass in each spectrometer solenoid could therefore sway.




Figure 4.3: *Top:* The movement of the draw wires attached to the upstream (SSU) and downstream (SSD) spectrometer solenoids from their mean position during ISIS cycle 2017/2, covering the 19 September 2017 to 27 October 2017. *Bottom:* The corresponding current supplied to each coil. Note: For the first three weeks of the cycle, the currents for the downstream coils failed to be recorded.

4.3 Centre of the magnetic axis

Similarly, as before when Langlands found the centre of the M2 coils along the longitudinal axis, Langlands⁷⁶ was able to find the centre of the magnetic axis in each spectrometer using the azimuthal magnetic field component (Fig. 4.4). Table 4.3 shows the difference in the magnetic axis between the one obtained from the magnetic field mapping study and the one obtained from the positioning survey.

The distance of the offset of the magnetic axis is similar to the distance of the offset of the centres of the M2 coils between the magnetic field mapping study and the positioning survey. This would indicate that the cold mass has moved slightly and perhaps even swayed as a result of both the thermal contraction of the mandrel and the coils, as well as due to the electromagnetic forces generated by the coils.




Figure 4.4: Uncorrected (left) and corrected (right) vector plots of the transverse field component B_ϕ measured by the Hall probe at a radial position of 0 mm during a whole rotation of the mapper disk⁷⁶. The centre of the corrected plot points to the centre of the magnetic axis.

Table 4.3: After aligning the mapper with MICE global reference frame along the longitudinal axis, the misalignment of the centre of the upstream (SSU) and downstream (SSD) spectrometer solenoids were determined by Langlands⁷⁶ as in Fig. 4.4. The centre of the magnetic axis (x_c, y_c, z_c) is related to the centre of the mapper axis (x_m, y_m, z_m) via $x_c = x_m + \theta_y z_m + p_x$ and $y_c = y_m - \theta_x z_m + p_y$, where $p_{x,y}$ and $\theta_{x,y}$ are the x and y offsets and rotations of the magnetic axis relative to the mapper axis i.e. the MICE global reference frame.

Spectrometer	X Offset (mm)	Y Offset (mm)	X Rotation (mrad)	Y Rotation (mrad)
SSU	-13.05 ± 0.44	-7.64 ± 0.44	-0.61 ± 0.03	0.90 ± 0.03
SSD	-17.40 ± 0.43	-8.43 ± 0.43	0.26 ± 0.02	1.12 ± 0.02

4.4 Deformation of the solenoid coils

In a COMSOL finite element analysis study, Witte et al.^{162,163} found that the contraction of the spectrometer solenoid posed another problem. The thermal expansion coefficient of aluminium is twice as large as that of the niobium-titanium. The aluminium mandrel would therefore contract faster than the five NbTi coils. The study indicated that the aluminium mandrel would deform, while some of the coils attached to the mandrel could become loose and potentially even detach (Fig. 4.5). Due to the electromagnetic forces generated by the coils, a coil could slip in its pocket. The heat dissipated could cause a quench of the solenoid magnets. This may explain why the spectrometer solenoids underwent such a large number of quenches when being trained¹⁶².

To prevent any further quenches, the trim power supplies powering the end coils of each solenoid weren't used any more. This resulted in new magnetic field maps being designed. The end coils could no longer shape the magnetic field produced by the centre coil, which resulted in the magnetic field becoming non-uniform within each tracker region. The new magnetic field maps also lowered the magnetic field strengths from 4 Tesla to approximately 2 or 3 Tesla within the tracker regions.




Figure 4.5: COMSOL finite element analysis study by Witte et al.^{162,163} showing the deformation of the E2 coil and aluminium mandrel. Note: the deformation has been amplified by a factor of ten.

4.5 Tuning the magnetic field

During the MICE experiment, the magnetic field in the MICE cooling channel was measured using six Hall Probes (Table 4.4). Two of these Hall Probes became detached¹⁶⁴, however, leaving only three Hall probes in the upstream spectrometer and one Hall probe in the downstream spectrometer that could accurately measure the magnetic field in the tracker regions¹⁶⁵.

Table 4.4: The longitudinal position and rotation around the longitudinal axis of the Hall Probes in the MICE cooling channel. The Hall probes are located at a radial position of approximately 160 mm. The longitudinal positions and rotations of Hall probes 66 and 67 are uncertain, as at some unknown time they became detached from the tracker.

Hall Probe Number	77	79	65	72	67	66
Position (mm)	14104	14429	14429	19482	15286	18625
Rotation (degrees)	30	270	30	330	150	210

The magnetic field measured by the three Hall probes in the upstream spectrometer during ISIS cycle 2017/3 is shown in Fig. 4.6. The magnetic field strength changes as a function of time due to changes in the magnetic field configuration. However, for a given magnetic field configuration, the measured magnetic field remains constant.

The difference in the measured magnetic field within the upstream tracker region by Hall Probe 77 compared to Hall probes 79 and 65 (i.e. at different longitudinal positions) shows that the magnetic field was non-uniform. There was also a small difference in the measured magnetic field between Hall Probes 79 and 65 which shared the same longitudinal position but were separated by a 120-degree rotation around the longitudinal axis. The difference in the magnetic field strength may have been due to a misalignment of the magnetic axis or due to a positional misalignment of the Hall Probes themselves.

The measured magnetic field however differed from the simulated MICE magnetic field based on the warm bore dimensions of the solenoids. To overcome this problem in MICE, the magnetic fields were tuned in MAUS to match the measured magnetic field¹⁶⁵. This resulted in the E1, CC and E2 coil currents being increased by 2% for the upstream spectrometer and by 1.8% for the downstream spectrometer^{164,166}. However, when the tuned and untuned magnetic fields, along with the Hall probe measurements are compared in Fig 4.7, a discrepancy can still be seen between the tuned magnetic fields and the Hall probe measurements.

The MICE solution of tuning the magnetic fields of the upstream and downstream spectrometer solenoids by varying amounts to match the measured magnetic fields is unsatisfactory. This is due to the reconstructed transverse and longitudinal momentum depending directly on the magnetic field (Eqs. 1.33 and 1.34). Tuning




Figure 4.6: The total magnetic field measured by the three Hall probes (65, 77 and 79) in the upstream spectrometer solenoid at a radial position of 160 mm during ISIS cycle 3 (14 November 2017 to 20 December 2017).

the upstream spectrometer by a larger percentage than the downstream tracker without an adequate physical explanation introduces a bias that could enhance the ionization cooling result.

A possible explanation for the lower simulated magnetic field strength can be found when one considers how the magnetic field is generated in MAUS. The magnetic field for a coil is generated through the *ScaleFactor* term in the parent geometry file. For the M2 coil in the upstream spectrometer, this is given by

$$ScaleFactor = P_{SSUM2} \times C_{SSUM2} \times S_{SSUM2} \times 0.5176 \quad (4.1)$$

where P_{SSUM2} determines the polarity of the coil, C_{SSUM2} is the current delivered to the coil, while S_{SSUM2} determines the percentage the coil has been scaled or tuned by. The magnetic field generated will also depend on the geometry of the coil. For the upstream spectrometer M2 coil, this appears to have been hard coded by the value 0.5176 in MAUS. The parent geometry file does not specify how this value of 0.5176 has been calculated (it will be referred to as N_{geo}), though it appears to have been calculated as follows:

$$N_{geo} = \frac{N_l \times N_t}{l \times d} \quad (4.2)$$

where N_l and N_t are the number of layers and turns of wire in the coil, while l and d is the length and depth of the coil. Tables 4.5 and 4.6 show the N_{geo} values that were used in MAUS for each coil, and the values calculated for each coil based on the warm bore dimensions.




Figure 4.7: The MICE magnetic field within the upstream (top) and downstream (bottom) tracker regions, based on the warm bore dimensions of the solenoids (black line) and the tuned MICE magnetic field (green line) as a function of longitudinal position, for MICE run 10314. Also shown are the Hall Probe measurements (red crosses).

If this description is correct, then the warm bore dimension and MAUS N_{geo} values are very similar, except that the N_{geo} values for the E1 and E2 coils of the upstream spectrometer appeared to have been swapped. This similarly applies to the M1 and M2 coils of the downstream spectrometer. For the downstream spectrometer the N_{geo} values for the E1, E2 and CC coils appear a little high. This may explain why the downstream spectrometer had to be tuned by only 1.8% to match the measured magnetic field, whereas the upstream spectrometer had to be tuned by 2%. In fact, when the cold bore dimensions for the upstream spectrometer solenoid are considered in Table 4.5, then the N_{geo} values for the cold bore dimensions are 1.5 to 2.5 percent larger than the N_{geo} values for the warm bore dimensions. This could explain the difference between the measured magnetic field values from the Hall Probes and those used in MAUS which are based on the warm bore dimensions.

The discrepancies seen between the measured magnetic fields and those expected

Table 4.5: The warm and cold bore dimensions of the upstream spectrometer solenoid coils, along with their calculated N_{geo} values. The N_{geo} value for MAUS has been taken from the parent geometry file. N_l is the number of layers of the copper matrix niobium-titanium wires, while N_t is the number of turns.

Coil	Warm Depth (mm)	Warm Length (mm)	Cold Depth (mm)	Cold Length (mm)	N_l	N_t	MAUS N_{geo}	Warm N_{geo}	Cold N_{geo}
M1	46.1	201.2	44.5	200.6	42	115	0.52	0.5207	0.5411
M2	30.9	199.4	29.6	198.9	28	114	0.5176	0.5181	0.5422
E1	60.9	110.6	59.4	110.2	56	64	0.529	0.5321	0.5475
CC	22.1	1314.3	21.2	1310.1	20	768	0.528	0.5288	0.5530
E2	67.7	110.6	65.7	110.2	62	64	0.532	0.5299	0.5481

Table 4.6: The corresponding values of Table 4.5 for the downstream spectrometer solenoid.

Coil	Warm Depth (mm)	Warm Length (mm)	N_l	N_t	MAUS N_{geo}	Warm N_{geo}
M1	46.4	201.2	42	115	0.52	0.5174
M2	30.6	199.4	28	114	0.5174	0.5231
E1	61.6	110.6	56	64	0.5316	0.5261
CC	22.4	1314.3	20	768	0.52817	0.5217
E2	68.2	110.6	62	64	0.5291	0.5261

by MICE based on the warm bore dimensions in the upstream and downstream spectrometer solenoids were overcome by simply tuning the E1, E2 and CC coils by 2% and 1.8% respectively. Such a tuning ignores the physical reason for such discrepancies and could bias the reconstructed momentum of the particles. Additionally, MICE did not scale the M1 and M2 coil currents. For MICE simulations, this will likely lead to a small error in the transport of particles between the upstream and downstream spectrometers.

A similar scaling was performed for the coil currents of the warm D1 and D2 dipole magnets. The dipole magnets were used to select muons with the appropriate momenta for the MICE cooling channel. For the Monte Carlo simulations, the coil current of the D2 magnet was increased by 4%, so that the momenta of the muons in the Monte Carlo simulations matched the momenta of the muons in the data¹⁶⁷. To check whether such a scaling was appropriate, Franchini^{168,169} measured the magnetic fields produced by the dipole magnets as a function of dipole current. Franchini noted that the difference between the measured magnetic field and the expected magnetic field from the coil currents was never larger than 2%. Simply scaling the D2 coil current by 4% is therefore inappropriate.

4.6 Misalignments

Misalignments can cause a number of issues. If the trackers in MICE are misaligned relative to the magnetic field, then the particles will be reconstructed with the wrong momentum. This is due to the MICE reconstruction assuming a helical particle trajectory which traces out a circle in the transverse position plane. If there is a misalignment of the tracker axis relative to the magnetic field, however, then the particle will trace out an ellipse in the transverse position plane of the tracker reference frame. Additionally, the transverse components of the beam are no longer fully separable from the longitudinal components of the beam in the tracker reference frame. This could bias the ionisation cooling result. The following subsections will consider the misalignment of the solenoid, tracker and magnetic axes.

Misalignment of the solenoid axis

The alignment of the spectrometer solenoids were measured using laser telemetry positional surveys^{74,170}. A laser theodolite was used to measure 16 specific points on the outside of each solenoid. Additionally, the laser theodolite was also used to measure the centre of the flanges at the upstream and downstream ends of each spectrometer solenoid. The flange centres were used to form the solenoid axis of each spectrometer. However, as the centre of the flanges cover an area of free space (i.e. the solenoid bore), it is less clear how these measurements were performed. A three-dimensional grid was formed by the 16 points on the outside of each solenoid, along with the flange centres.

The positional surveys of the spectrometer solenoids were periodically repeated when changes in the experiment occurred, such as when the absorber material was changed^{171–178} (though no positional survey was taken after the installation of the polyethylene wedge). However, due to the presence of the PRY, only four survey points could be accessed on the outside of each solenoid. The movement of these four survey points relative to their positions in the previous survey were used to infer the new positions of the remaining 12 survey points, as well as the flange centres and therefore the solenoid axis in each spectrometer.

Misalignment of the tracker axis

Once the trackers were installed, the position of the trackers were inferred with respect to the end plates of each spectrometer solenoid. To verify the positions of the trackers, Drielsma^{74,170} developed a beam-based detector alignment procedure to determine the central axes of the trackers. The procedure determines the central axes of the trackers with respect to a line joining the centres of the TOF1 and TOF2

stations, the two time-of-flight detectors either side of the spectrometer solenoids. The position of the TOFs were accurately determined from positional surveys^{179,180}. However, the results from the surveys weren't always implemented into MAUS (Fig. 4.8).




Figure 4.8: The transverse offsets of TOF1 and TOF2 in MAUS, taken from positional surveys^{179,180}. Not all results have been implemented, however. For example, a positional survey of TOF1 was completed on the 25 May 2017 (near run 9300), but was never implemented into MAUS.

The tracker axes were determined using straight line tracks from 200 MeV muons, and 300 and 400 MeV pions passing through the TOF detectors. As the particles travel in straight lines from TOF1 to TOF2 (barring for any scattering), the positional hits in each tracker relative to the positional hits in each TOF can be used to determine the offset and rotation of each tracker axes relative to the centre line joining TOF1 to TOF2. The beam based detector alignment procedure was repeated for each ISIS cycle to account for any changes in the MICE experiment. Tables 4.7, 4.8 and 4.9 summarize the measured misalignments of the tracker axes relative to the centre line joining TOF1 to TOF2 as a result of the beam-based alignment procedure.

Table 4.7: The beam based alignment procedure by Drielsma^{74,170} was completed for each ISIS cycle. To prevent an asymmetric sampling bias, not all particle tracks triggering TOF1 were used. Only those particles tracks below a certain gradient and radius in the upstream tracker were selected.

ISIS cycle	Dates	Run Range	Triggers	Selected	% Selected
2015/4	23-26/2/2016	7624 – 7639	1.19×10^6	22146	1.86
2016/2	9-10/7/2016	8001 – 8014	1.55×10^6	47569	3.07
2016/2	16/7/2016	8030 – 8041	0.72×10^6	15779	2.19
2016/3	7-9/10/2016	8381 – 8406	0.59×10^6	15790	2.68
2016/3	19-20/10/2016	8431 – 8434	0.52×10^6	2771	0.53
2016/4	25-27/11/2016	8616 – 8632	1.52×10^6	31104	2.05
2017/1	29-30/5/2017	9353 – 9372	2.11×10^6	79316	3.76
2017/2	19/9/2017	9619 – 9620	0.93×10^6	40872	4.39
2017/3	29-30/11/2017	10303 – 10310	2.08×10^6	93713	4.51

Table 4.8: The x and y offsets and the α and β rotations (rotations around the x and y axes) for the upstream tracker calculated using the beam-based alignment procedure by Drielsma^{74,170} for the straight line track runs defined in Table 4.7^{181–187}.

Dates	x (mm)	y (mm)	α (mrad)	β (mrad)
23-26/2/2016	-0.451 ± 0.149	-0.611 ± 0.146	3.153 ± 0.048	0.855 ± 0.045
9-10/7/2016	2.197 ± 0.088	-0.498 ± 0.104	3.461 ± 0.031	-0.212 ± 0.031
16/7/2016	2.008 ± 0.183	-0.247 ± 0.170	3.545 ± 0.053	-0.270 ± 0.046
7-9/10/2016	3.406 ± 0.182	1.473 ± 0.169	3.865 ± 0.050	-0.604 ± 0.044
19-20/10/2016	2.146 ± 0.530	0.185 ± 0.517	3.696 ± 0.245	-0.492 ± 0.236
25-27/11/2016	1.361 ± 0.133	-1.367 ± 0.136	3.488 ± 0.050	-0.268 ± 0.047
29-30/5/2017	0.147 ± 0.079	-1.616 ± 0.079	3.316 ± 0.025	0.272 ± 0.023
19/9/2017	0.446 ± 0.102	-1.585 ± 0.105	3.006 ± 0.031	0.139 ± 0.028
29-30/11/2017	-0.688 ± 0.071	-1.155 ± 0.072	3.576 ± 0.019	0.666 ± 0.019

Table 4.9: The x and y offsets and the α and β rotations (rotations around the x and y axes) for the downstream tracker calculated using the beam-based alignment procedure by Drielsma^{74,170} for the straight line track runs defined in Table 4.7^{181–187}.

Dates	x (mm)	y (mm)	α (mrad)	β (mrad)
23-26/2/2016	-2.631 ± 0.137	2.638 ± 0.135	-0.614 ± 0.043	1.352 ± 0.043
9-10/7/2016	-2.705 ± 0.113	3.150 ± 0.110	-1.223 ± 0.032	0.903 ± 0.033
16/7/2016	-3.015 ± 0.157	3.009 ± 0.155	-1.113 ± 0.045	1.075 ± 0.045
7-9/10/2016	-2.497 ± 0.149	3.141 ± 0.150	-0.889 ± 0.042	0.897 ± 0.042
19-20/10/2016	-3.138 ± 0.518	2.134 ± 0.518	-0.695 ± 0.225	0.980 ± 0.213
25-27/11/2016	-2.929 ± 0.120	2.162 ± 0.120	-0.957 ± 0.045	1.040 ± 0.045
29-30/5/2017	-2.772 ± 0.069	2.894 ± 0.070	-0.043 ± 0.023	1.264 ± 0.022
19/9/2017	-2.179 ± 0.093	2.672 ± 0.094	0.157 ± 0.029	1.727 ± 0.028
29-30/11/2017	-3.283 ± 0.072	3.329 ± 0.072	-0.660 ± 0.019	0.950 ± 0.019

Misalignment of the magnetic axis

A similar magnetic field mapping study to that of Langlands was performed by Blackmore and Cobb^{188–191} before the installation of the PRY on the 27 March 2015. The offsets of the magnetic axes relative to the solenoid axes are summarised in Table 4.10, though it is unclear how these initial results were implemented in MAUS. From these initial results, the magnetic axes were usually (but not always) rotated by the same values the solenoid axes were rotated by after each positional survey. This can be seen by the yellow line in Fig. 4.9. Fig. 4.9 summarises the implemented rotational misalignments of the tracker, solenoid and magnetic axes in MAUS, as a function of MICE run number. The tracker and magnetic axes misalignments included the misalignments of the solenoid axes, and were removed for the pink and yellow lines, respectively.

Table 4.10: The X and Y offsets show the misalignment of the magnetic axes relative to the solenoid axes, where the upstream and downstream ends of each spectrometer correspond to the longitudinal position of the flange centres used to align the solenoid axes¹⁸⁹. Also shown are the rotation of the magnetic axes around the solenoid axes.

Spectrometer	SSU	SSD
Upstream end X offset (mm)	0.12	0.78
Downstream end X offset (mm)	-0.40	0.55
Upstream end Y offset (mm)	-0.34	-3.90
Downstream end Y offset (mm)	-1.60	-10.95
Rotation around X axis (mrad)	0.180	0.080
Rotation around Y axis (mrad)	0.434	2.434
Upstream end Z position (mm)	13566.62	17432.77
Downstream end Z position (mm)	16482.39	20347.86

4.7 Discussion

How the misalignments of the magnetic axes from the initial field mapping study were implemented into MAUS is not clear. Additionally, the presence of the PRY, and the forces generated by the simultaneous powering of the focus coil and the spectrometer solenoids may have caused the magnetic axes to move¹⁹¹. A tentative comparison can be made with the values from the field mapping study by Langlands⁷⁶ in Table 4.3 which suggest that the cold masses in the spectrometers may have swung slightly towards each other due to the presence of the PRY. For both field mapping studies, the focus coil was not present, however. Therefore, when all the solenoid coils were powered simultaneously, the changes in the alignment of the magnetic axes (due to the forces generated by the coils) were not measured. This also meant that such misalignments were not considered or implemented in MAUS.

As the tracker is housed within the spectrometer solenoid, the tracker would not be expected to move relative to the spectrometer solenoid. However, this is not the case (pink lines in Fig. 4.9). This is due to the tracker axes depending on the beam-based alignment procedure by Drielsma, though it is not clear how these results were implemented into MAUS. In fact, not all results of the alignment procedure were implemented, e.g. ISIS cycles 2016/4 and 2017/2. More concerning, however, the variation of the tracker axes relative to the solenoid axes as a function of MICE run number can be greater than 0.4 degrees (or 7 mrad). This implies that the tracker has moved by 7 mrad relative to the spectrometer solenoid while being housed within the spectrometer solenoid, which does not seem plausible. The misalignments implemented into MAUS (and used for the MICE reconstruction) are therefore not completely trustworthy. The following chapter will consider the effect of misalignments and a non-uniform magnetic field on the MICE reconstruction.

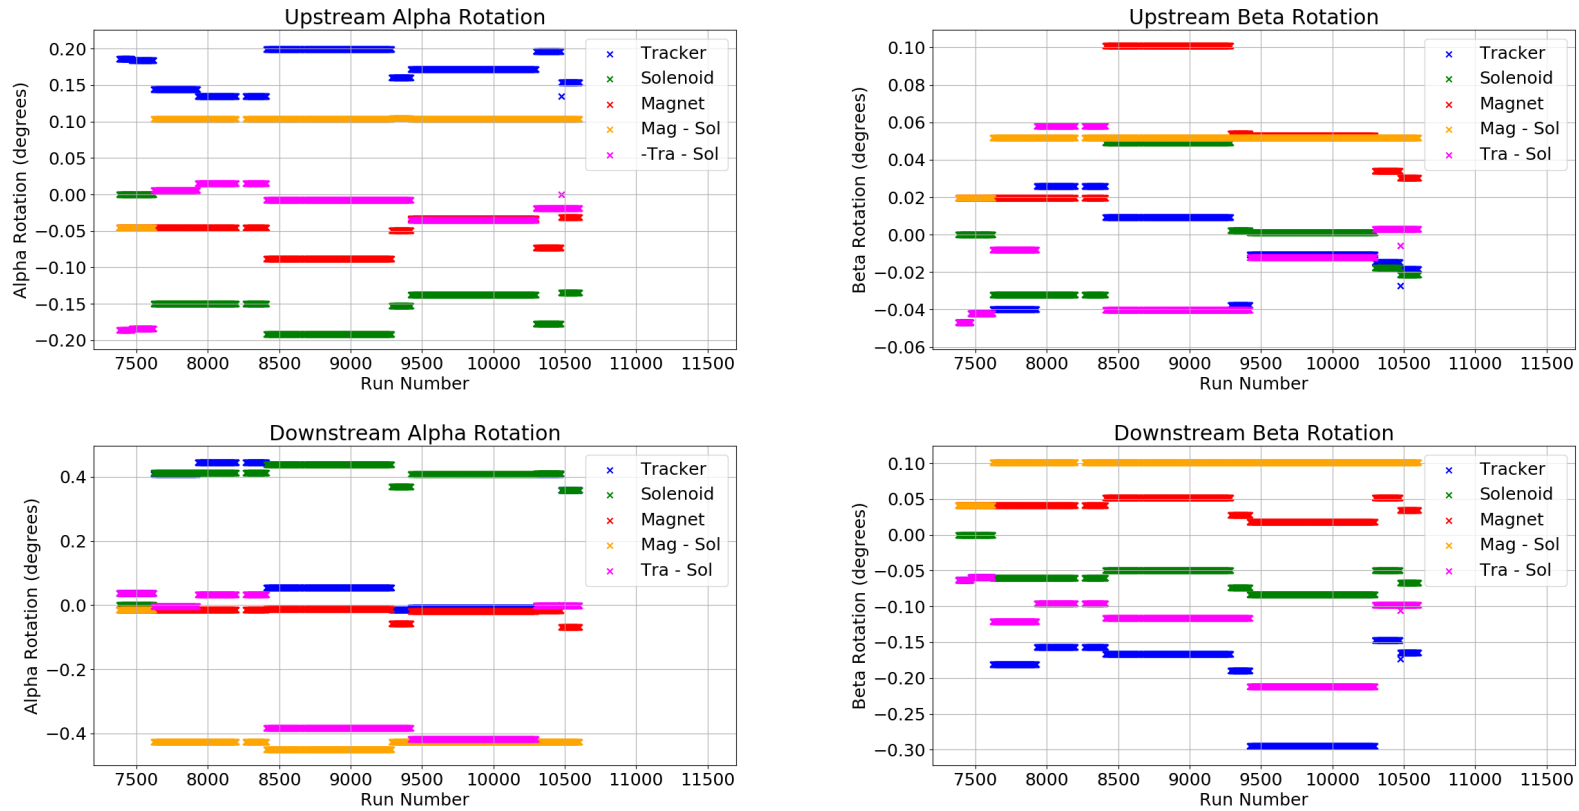


Figure 4.9: The alpha (left) and beta (right) rotations of the upstream (top) and downstream (down) tracker, solenoid and magnetic axes for each MICE run. The rotation of the tracker and magnetic axes include the rotation of the solenoid axes, though this has been removed for the pink and yellow markers, respectively. Note, for the pink markers, a factor of -1 has been included for the upstream alpha rotation so that those values match the values from the parent geometry files.

Momentum reconstruction at MICE

The spectrometer solenoids of the MICE experiment were designed to operate with a uniform four Tesla solenoid field across the tracker regions. The uniform field meant the transverse and longitudinal components of the assembled muon beam could be treated separately⁷⁸, a necessity for the demonstration of ionisation cooling. The strength of the field ensured the expected errors on the reconstructed momentum remained small (Table 5.1), though they were still statistically significant.

Table 5.1: The reconstructed transverse (p_t) and longitudinal (p_z) momentum residuals for the upstream and downstream trackers, before the failure of one of the coils in the downstream solenoid⁸⁹.

	p_t mean	p_t RMS	p_z mean	p_z RMS
TKU	-0.095 ± 0.009	0.910 ± 0.015	-0.102 ± 0.013	1.683 ± 0.018
TKD	0.067 ± 0.009	0.919 ± 0.015	0.468 ± 0.013	1.660 ± 0.018

The failure of one of the coils in the downstream solenoid meant the magnetic field lattices had to be reconfigured. To prevent risk of further damage, the spectrometers were operated with a maximum magnetic field of 3 Tesla. This was achieved by altering the coil currents supplied to each coil.

To prevent any further quenches, the end coils and centre coils in each spectrometer were powered using the same power supplies and therefore had the same coil currents. This resulted in the magnetic fields no longer being uniform within the tracker regions (despite being claimed to be and presented as uniform⁸). This can be seen in Fig. 5.1 where at the upstream tracker reference plane, the strength of the magnetic field deviated by 2% from the assumed uniform magnetic field used in the MICE reconstruction. At the downstream tracker reference plane, which typically operated using a 2 Tesla magnetic field, this deviation increased to 3%.




Figure 5.1: The magnetic field in the tracker regions of the MICE experiment is claimed to be and presented as uniform⁸ (*top*). However, a closer inspection¹⁶⁵ reveals that the magnetic field strength at the upstream tracker reference plane differs by 2% from the assumed uniform magnetic field strength used in the MICE reconstruction (*bottom left*) and by 3% at the downstream tracker reference plane (*bottom right*). The blue lines in each plot refer to the longitudinal position of the stations in each tracker, while the dark green and light green lines show the magnetic field strength along the central beam axis and 160 mm away from the central beam axis.

The effect the new magnetic field lattices had on the reconstructed momentum can be seen in Fig. 5.2. The figure shows that the reconstruction depends on the momentum of the particles. Additionally, the resolution of the reconstructed total momentum of the particle is worse for the downstream tracker than for the upstream tracker. This is due to the lower magnetic field strength used for the downstream spectrometer.

The major problem of the MICE reconstruction however is that the upstream tracker underestimates the total momentum of the particles while the downstream tracker overestimates the total momentum of the particles. The reconstruction effectively adds energy to the particles between the trackers. This similarly occurs for the reconstructed transverse momentum, though to a lesser extent. The biases of the momentum reconstruction can also bias energy loss measurements as well as the ionisation cooling result.

How the MICE reconstruction is affected by issues such as the non-uniformity of the MICE magnetic field, misalignments, energy loss and scattering will be investigated over the following sections.

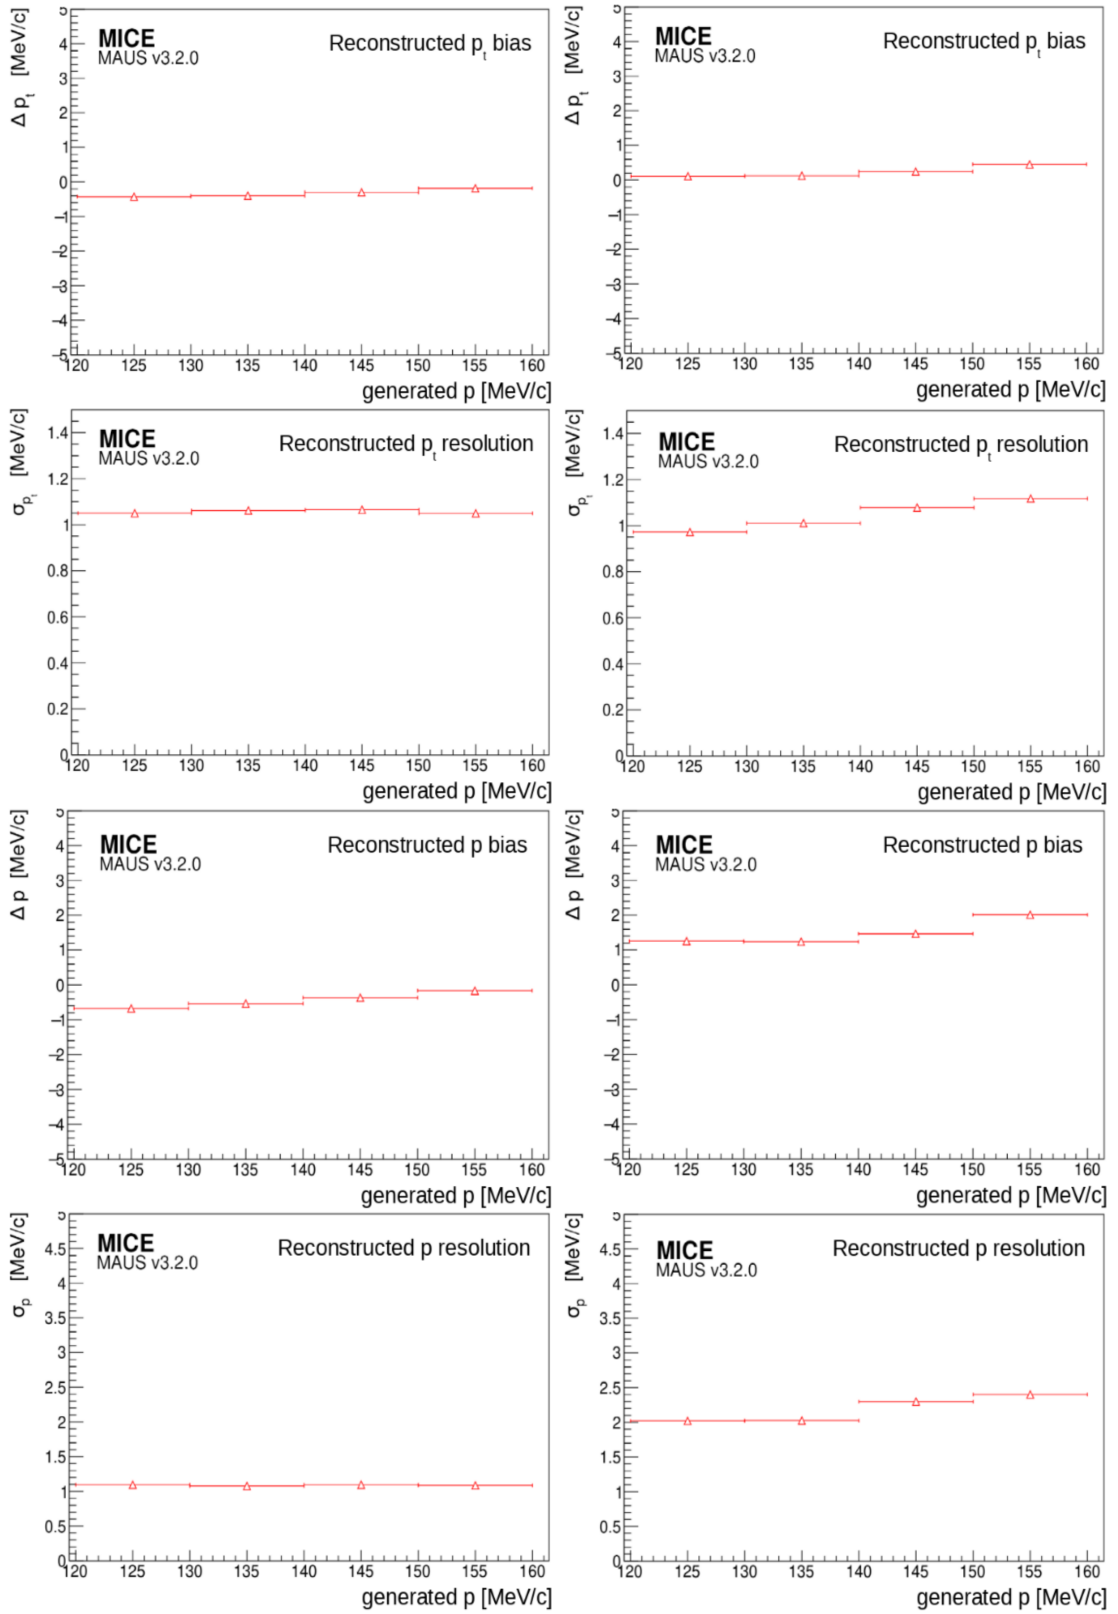


Figure 5.2: The transverse momentum bias (*top row*) and resolution (*second row*), as well as the total momentum bias (*third row*) and resolution (*bottom row*) for the upstream (*left*) and downstream (*right*) tracker³.

5.1 Linear solenoid transfer matrix assumptions

The ionization cooling result assumes a uniform magnetic field. This allows the transverse components of the beam to be treated separately from the longitudinal components of the beam⁷⁸. This is based upon applying the linear approximation to Eq. 1.19. However, Fig. 5.2 already showed that the magnetic field within the tracker regions is not uniform. This means a transfer between the longitudinal and transverse components of the beam (and vice versa) can take place that is not due to ionisation cooling, but rather due to the non-uniformity of the solenoid field.

The linear approximation of the solenoid transfer matrix also assumed that the longitudinal momentum was far greater than the transverse momentum. This means particles that have a large transverse momentum may be poorly reconstructed or not reconstructed at all. This can be seen in Fig. 5.3 showing the percentage of reconstructed particles decreasing as the transverse momentum increased. A drop-off in efficiency is also seen for very low transverse momenta particles, but this is mainly a resolution effect as at lower transverse momenta, particles trace out helices with smaller radii whose momenta are far more difficult to accurately reconstruct.

The transverse momentum dependence of the track finding efficiency poses two problems. If high transverse momenta particles are not reconstructed in the downstream tracker, then the ionisation cooling effect may be overestimated. However, if low transverse momenta particles are not reconstructed, then the cooling effect may be underestimated. Similar issues apply to the upstream tracker. The upstream tracker can create a selection bias if low transverse momenta particles that may have been heated or high transverse momenta particles that may have been cooled, have been excluded.




Figure 5.3: The track finding efficiency¹⁶⁵ in the downstream tracker as a function of the transverse momentum coordinates p_x and p_y for input beams with nominal momenta of 140 MeV and input emittances of 6 mm (left) and 10 mm (right).

5.2 Biases of the MICE reconstruction

The MICE reconstruction assumes a uniform solenoid field, with the particles travelling along a helical trajectory through the spectrometer solenoids. The particles may deviate from this helical trajectory for several reasons. Scattering can send the particles onto a lower or higher radius helical path. Energy loss, which slows particles down, will send the particles onto lower radius helical paths.

Similarly, if the magnetic field is changing within the tracker regions, then the particles will no longer follow a helical trajectory, but will instead follow a spiral trajectory as the radius of the helical path traced out in the transverse position space has a dependence on the strength of the magnetic field (Eq. 1.33).

If there are misalignments of the tracker, solenoid or magnetic axes, then the helical trajectory path will no longer trace out a circle in transverse position space, but rather an ellipse. The effect each scenario can have on the MICE reconstruction will be considered in turn.

5.2.1 Scattering and energy loss

Scattering will result in the radius of the particle's helical trajectory changing. This also results in a transfer between the transverse and longitudinal momenta, or vice versa. However, as the MICE reconstruction performs a χ^2 cut on the circle fit of the helical trajectory path traced out in transverse position space, only those particles that pass the χ^2 fit are reconstructed, i.e. particles that encounter little to no scattering in the trackers. The cut is to ensure that particles that have significantly scattered in the tracker stations are removed from the analysis, as the error on their reconstructed momentum will be large.

The particles will however undergo energy loss at the tracker stations, resulting in a periodic loss of transverse momentum. Selection biases can therefore occur. For a muon with a momentum of 140 MeV, the particle will lose approximately 0.6 MeV per station or approximately 3 MeV per tracker. This corresponds to a 2% reduction in the transverse momentum and therefore a 2% reduction in the radius of the circle traced out by the particle in transverse position space as it passes through the tracker. Due to the finite resolution of the scintillating fibres in the tracker, the effect on low transverse momenta particles will be small. For large transverse momenta particles, the effect can be significant. If a particle does undergo energy loss, then the particle will hit a station a few channels away from where it would have hit the station if it didn't undergo energy loss.

If the particle is on a high radius helical trajectory path, then the energy loss will cause a significant distortion to the circle traced out in transverse position

space, resulting in some particles no longer passing the χ^2 cut and therefore not being reconstructed any more. However, if some of the particles now encounter scattering, then the change in their helical path and the subsequent transfer between their longitudinal and transverse momenta may result in some particles being reconstructed that previously wouldn't have been. Similarly, some particles that were significantly scattered to a larger radius helical path that previously wouldn't have been reconstructed will now be reconstructed as their helical trajectory path will be altered due to the energy loss experienced at the tracker stations. Similar scenarios exist with particles no longer being reconstructed that previously would have been reconstructed due to encountering energy loss and scattering.

If the particle is on a helical trajectory whose path takes it beyond the aperture of the tracker, similar issues on the types of particles that are or aren't reconstructed apply in terms of scattering and energy loss. If the magnetic field is non-uniform, then these effects may be exacerbated, as the shape of the circle traced out in transverse position space is further distorted. A further selection bias is introduced if the non-uniform magnetic field strength differs between the two trackers, as the radius of the helical trajectory path depends on the magnetic field strength.

Whenever the particle encounters a tracker station, the particle will end up on a new helical trajectory path whose radius and circle centre will change. The pattern recognition step of the MICE reconstruction doesn't consider this however and instead performs a circle fit of the spacepoints in each tracker, finding the circle centre and the radius of a circle for those spacepoints. However, this circle is distorted whenever the particle experiences energy loss, with the circle centre shifting to a new position each time the particle experiences energy loss at a station.

5.2.2 Including the MICE magnetic field

As the MICE magnetic field is not uniform, the particles will deviate even further from their helical trajectory paths. This can be seen in Figures 5.5 and 5.6 where the blue bars consider the percentage of particles whose transverse position laid within the circle of the circle fit (Fig. 5.5) or outside of the circle of the circle fit (Fig. 5.6) at every station of each tracker and for four different absorber material scenarios. The additional coloured bars (green, yellow, red) show if the particles were more than 1, 2 or 3 millimetres inside or outside of the circle fit respectively.

Some trends can be immediately seen. In each tracker, the particle tends to start outside of the circle fit before ending up inside the circle fit, i.e. the particle spirals inwards. The first and last stations of each tracker show a significant number of particles that were more than 1, 2 and 3 mm inside or outside of the circle fit. This can be explained by the magnetic field deviating significantly from a uniform

magnetic field at those stations (Fig. 5.1), and therefore deviating the trajectory of the particle from a circular path in transverse position space.

A momentum dependence for the pattern recognition stage can be seen when absorbers are present in the MICE cooling channel. For all four scenarios (no absorber, wedge, lithium hydride and liquid hydrogen), the percentage of particles inside or outside the circle fit at a station is similar for the upstream tracker, however, differences can be seen for the downstream tracker. When a 140 MeV muon passes through either the lithium hydride or liquid hydrogen absorbers, it loses approximately 7-8% of its momentum. The lower transverse momentum results in the particle travelling along a lower radius orbital path in transverse position space. However, as the particle also has a lower longitudinal momentum, the particle will trace out a larger percentage of the circular path in the transverse position plane, i.e. the phase advance between stations increases. Additionally, as the magnetic field is also non-uniform, the deviation from the circular path for the lower longitudinal momentum particle will also increase. This can be seen when the last station of the downstream tracker is considered in Figures 5.5 and 5.6. The percentage of particles that are more than three millimetres inside or outside of the circle fit for the lithium hydride and liquid hydrogen absorber scenarios is twice as large as for the no absorber scenario.

If the Kalman filter had appropriately corrected the particle hits at every station for a helical trajectory, then there would be a fifty-fifty chance of each particle laying inside or outside of the circle fit at every station. The Kalman filter is however not performing as expected. This can be seen in Table 5.2. The table shows the mean and RMS energy loss experienced by each particle after passing through each station of the upstream and downstream trackers. The energy loss experienced by a particle is a parameter that is free to change in the Kalman filter, though should remain similar to the value calculated from the mean energy loss of the Bethe-Bloch curve. This results in the mean energy loss being similar at every station. However, the RMS of the energy loss is significantly smaller for the stations closer to the absorber material. This shows the Kalman filter has preferentially pulled the particles in some stations differently than in other stations to achieve a particular energy loss.

The issues with the Kalman filter can cause significant problems. As the particles follow a spiral trajectory instead of a helical trajectory, the circle fit creates a momentum bias that is not adequately corrected by the Kalman filter for a non-uniform solenoid field and explains the asymmetry in the momentum bias between the upstream and downstream trackers of Fig. 5.2. This is due to the reference planes being the last station of the upstream tracker and the first station of the downstream tracker. As the particle follows a spiral trajectory instead of a helical trajectory, particles near the start of the spiral have their momentum overestimated,

Table 5.2: The mean energy loss experienced by a particle passing through each station of the upstream and downstream trackers. It considers the difference in energy of a particle at that station compared to the previous station.

	S5 \leftrightarrow S4	S4 \leftrightarrow S3	S3 \leftrightarrow S2	S2 \leftrightarrow S1
TKU mean (MeV)	0.6042	0.6009	0.6138	0.6105
TKU RMS (MeV)	0.2462	0.235	0.1978	0.03991
TKD mean (MeV)	0.5631	0.6087	0.6188	0.6134
TKD RMS (MeV)	0.2414	0.2235	0.2117	0.05951

while particles near the end of the spiral have their momentum underestimated. This is reflected in the percentage of particles laying inside or outside of the circle fit at those stations. It is also responsible for the apparent increase in momentum between the reference planes.

In fact, it can be seen that the momentum bias has a dependence on the actual momentum of the particle itself (Fig. 5.4), making a correction for the momentum bias even more difficult.




Figure 5.4: The momentum residuals at downstream tracker reference plane as a function of the reconstructed momentum. The residuals have been calculated for each particle between the true and reconstructed momenta, using the first 100 files of the MICE Monte Carlo simulation¹⁹² 000247.

Figure 5.5: The blue bars show the percentage of particles whose transverse position laid within the circle of the circle fit at every station of the upstream and downstream trackers for the four different absorber scenarios: no absorber (top left), polyethylene wedge (top right), lithium hydride (bottom left) and liquid hydrogen (bottom right). The additional coloured bars (green, yellow, red) show if the particles were more than 1, 2 or 3 millimetres inside the circle fit respectively.

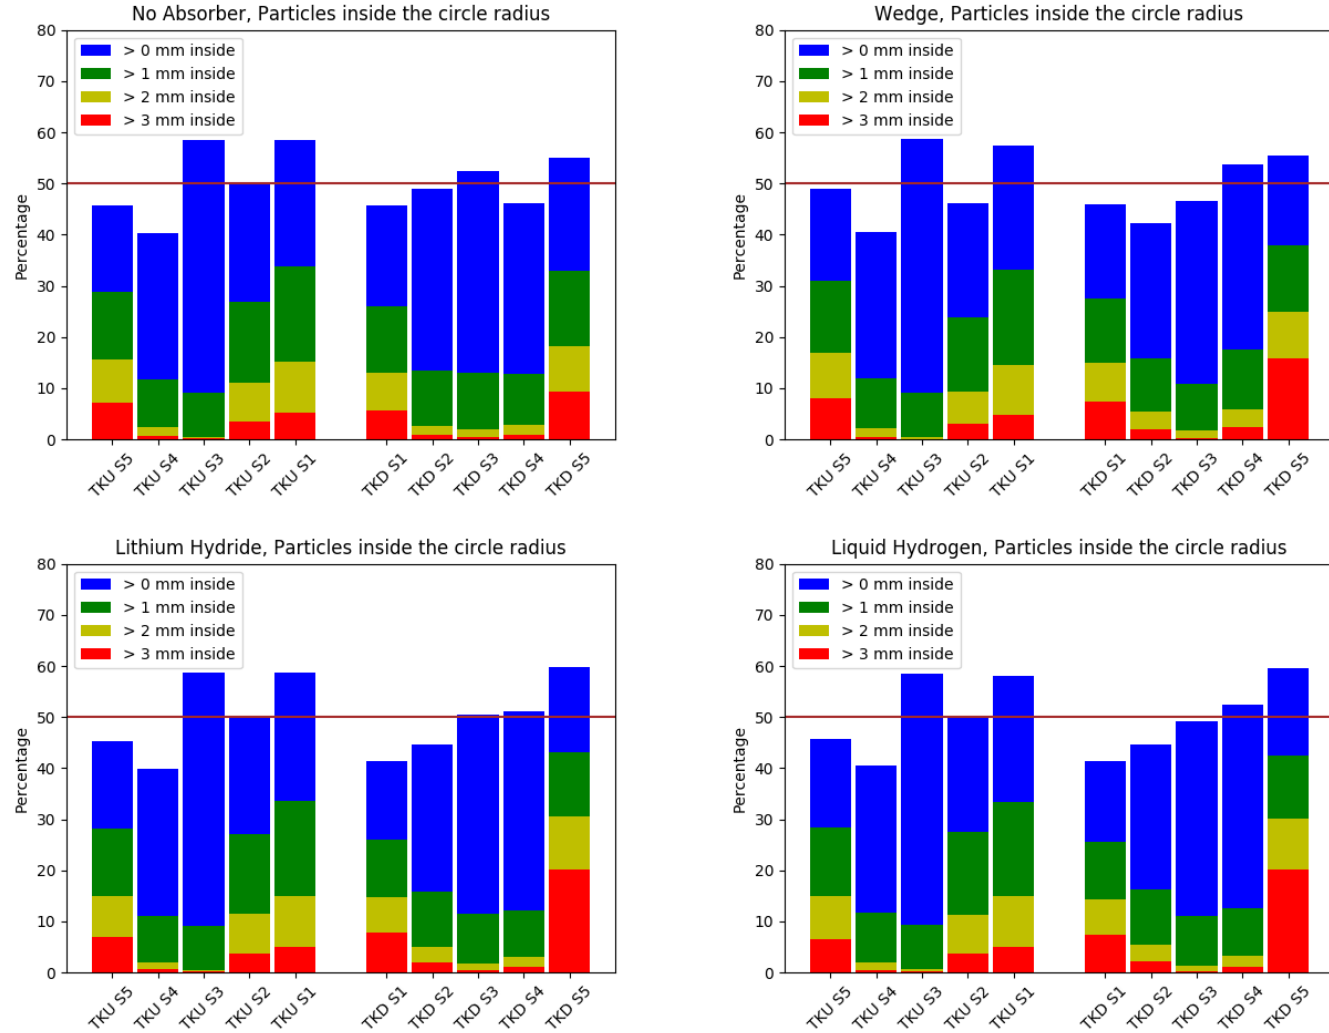
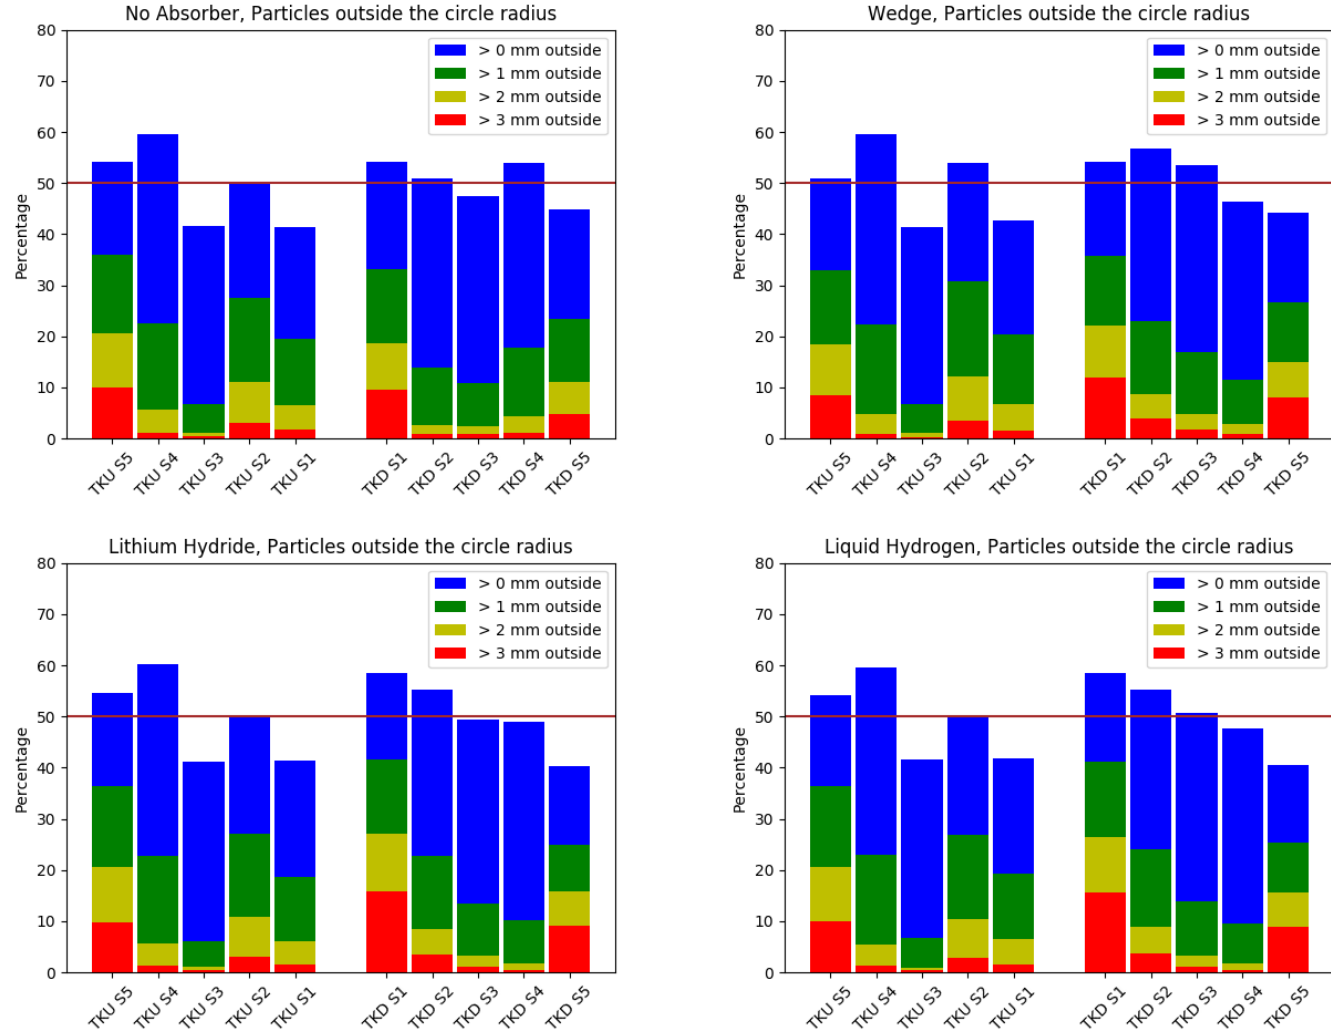


Figure 5.6: The blue bars show the percentage of particles whose transverse position laid outside the circle of the circle fit at every station of the upstream and downstream trackers for the four different absorber scenarios: no absorber (top left), polyethylene wedge (top right), lithium hydride (bottom left) and liquid hydrogen (bottom right). The additional coloured bars (green, yellow, red) show if the particles were more than 1, 2 or 3 millimetres outside of the circle fit respectively.



5.2.3 Including misalignments

For a uniform solenoid field, if the tracker axes were misaligned from the magnetic field axes, then the trajectory traced out by the particle in transverse position space will no longer be a circle, but rather an ellipse. The tilt also means that the transverse and longitudinal momentum components of the particle were no longer fully separated.

The misalignment between the tracker and magnetic axes also changed throughout the experiment as different absorber materials were placed into the beamline and as different magnetic field configurations were used. Additionally, the MICE magnetic field was non-uniform. This means a transfer between a particle's transverse and longitudinal momentum components took place as the particle traversed the non-uniform field of the tracker region.

5.3 Quantifying the momentum bias

The momentum bias due to the MICE reconstruction will be quantified for a number of different scenarios in this section. This was done by changing the MICE magnetic field or by introducing different misalignments of the solenoid, tracker and magnetic axes. Each scenario used the same input particles taken from the first 100 files of the MICE Monte Carlo¹⁹² simulation 247. The simulation corresponded to a 'solenoid' mode run (i.e. no field flip), where no absorber material was present in the MICE cooling channel and the input beam had a nominal emittance and momentum of 6 mm and 140 MeV respectively. From this Monte Carlo simulation, the muons at a virtual plane just before the most upstream plane of the upstream tracker were extracted. These extracted muons acted as the starting position for each scenario considered in this section. These muons were then propagated through the MICE cooling channel for each scenario.

The first scenario considered when there were no changes made to the implemented misalignments and the implemented magnetic field used in MAUS i.e. a repeat of MICE Monte Carlo simulation 247. Figures 5.7 and 5.8 show the mean (top) and RMS (bottom) transverse and longitudinal momenta residuals, at all 30 planes of the upstream and downstream trackers. The transverse momentum residuals show distinctive U-shapes, similar to the magnetic field U-shapes within the trackers (Fig. 5.1). The longitudinal momentum mean residuals become progressively larger as a function of longitudinal position, i.e. momentum is continuously added to the muon as a function of longitudinal position. Interestingly, though, the RMS stayed relatively constant within each tracker. For this scenario, a similar amount of momentum was added to each muon at the tracker reference planes.

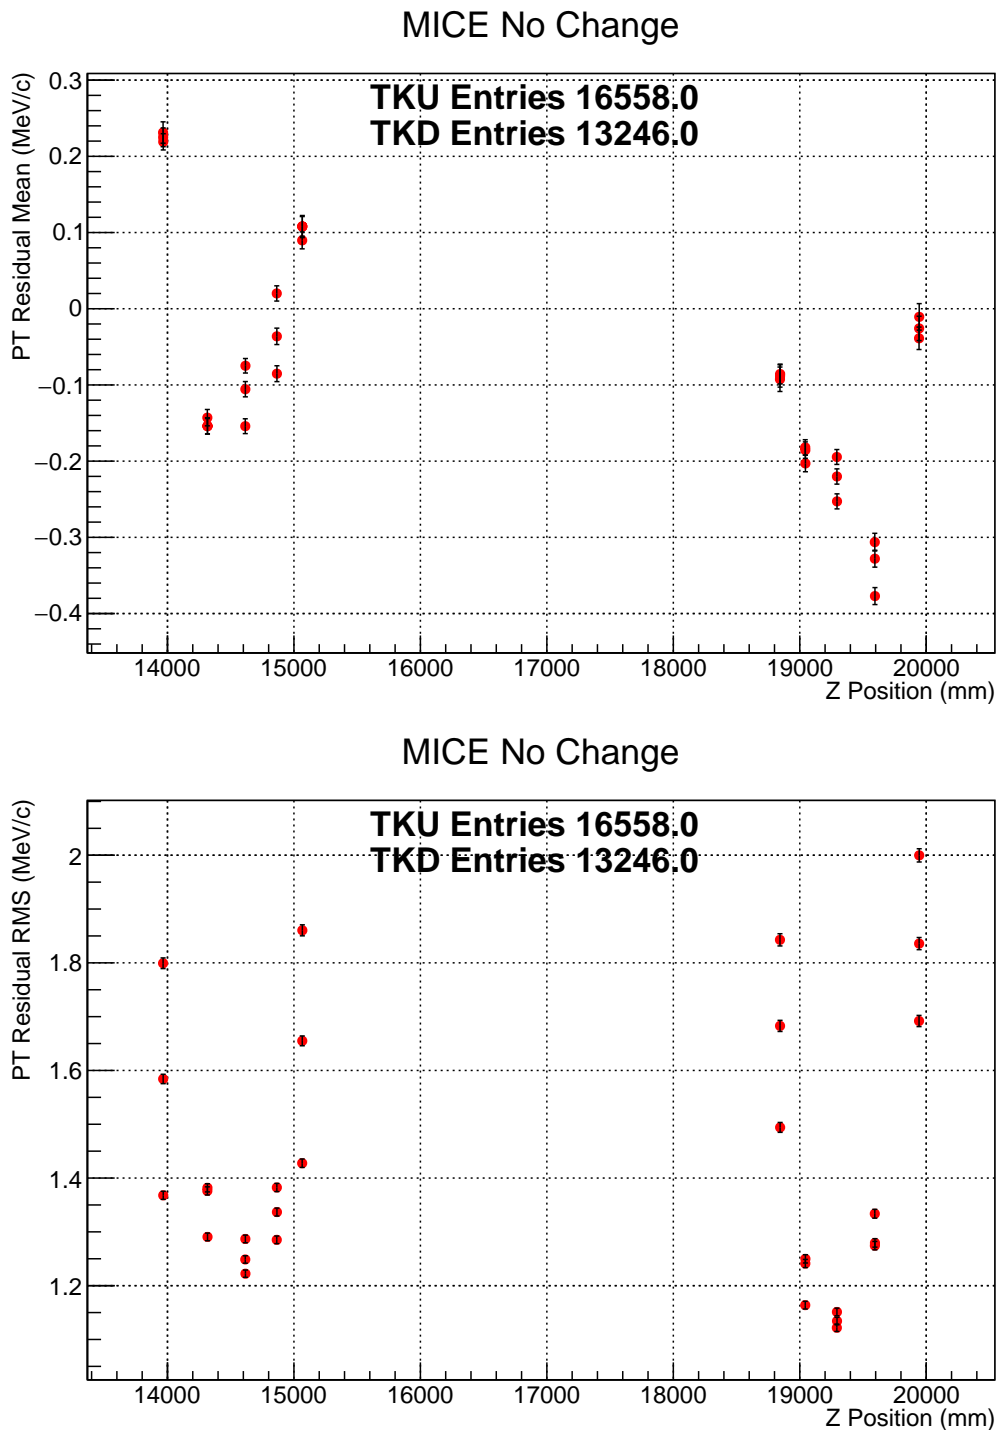


Figure 5.7: The mean (top) and RMS (bottom) residual transverse momentum at every plane of every station in the upstream and downstream trackers as a function of longitudinal position in the MICE cooling channel.




Figure 5.8: The mean (top) and RMS (bottom) residual longitudinal momentum at every plane of every station in the upstream and downstream trackers as a function of longitudinal position in the MICE cooling channel.

5.3.1 Misalignment of the MICE magnetic field

The transverse and longitudinal momentum residuals will be considered for a total of 22 different scenarios at each plane of the upstream and downstream trackers. The 22 scenarios consider two different magnetic field configurations of various field strengths, namely the MICE magnetic field and a constant solenoid field across both tracker regions. The 22 scenarios also consider various misalignments of the solenoid, tracker and magnetic axes. A summary of the 22 scenarios is given in Table 5.3. The table additionally shows the percentage of transmitted and reconstructed particles. A summary of the transverse and longitudinal momentum residuals at the tracker reference planes are given in Tables 5.4 and 5.5.

Figures 5.9 and 5.10 show the mean (top) and RMS (bottom) transverse and longitudinal momentum residuals for the MICE “No Change” scenario in comparison to the scenario where the coils of the upstream and downstream spectrometers as well as the solenoid and tracker axes have been aligned to the MICE global reference axis i.e. no misalignments (MICE 0° Sol, 0° Tr, 0° MF).

The transverse momentum residuals show distinctive U-shapes, while the longitudinal momentum residuals show the muons are effectively gaining momentum as a function of longitudinal position as a result of the MICE reconstruction. The larger misalignments of the downstream tracker are reflected by the noticeable difference in the residual momentum bias between the “No Change” and “No Misalignment” scenarios, especially for the residual longitudinal momentum RMS. The change in the mean longitudinal momentum residual now results in momentum being lost between the two reference planes as a result of the reconstruction.

The MICE (1.02 TKU 1.018 TKD), MICE (0.66 TKU 0.66 TKD) and MICE (1.5 TKU 1.5 TKD) scenarios consider the same misalignments as for the “No Change” scenario. It however scales the coil currents to different values. The MICE (1.02 TKU 1.018 TKD) scenario considers the scaling of the coil currents used for the MICE Nature paper^{8,165}, i.e. an increase in the coil currents of 2% for the upstream coils and an 1.8% increase for the downstream coils. The 66% scaling scenario results in 2 Tesla and 1.5 Tesla fields in the upstream and downstream trackers, while the 150% scenario results in 4.5 Tesla and 3 Tesla fields in the upstream and downstream trackers.

An appreciable difference is noted for the mean and RMS transverse momentum residuals for the 150% scenario. Unsurprisingly, the scaling used for the MICE Nature paper only produces a small change from the “No Change” scenario. The effect on the longitudinal momentum residual is more noticeable, however. For the “No Change” scenario, the mean residuals between the reference planes show no net change (TKU: -0.768 ± 0.028 MeV, TKD: -0.775 ± 0.042 MeV), while for the MICE

Nature paper scenario, a net momentum increase is experienced (TKU: -0.715 ± 0.027 MeV, TKD: -0.839 ± 0.042 MeV). Interestingly, when similar magnetic field strengths are considered for the upstream and downstream trackers, then the RMS longitudinal momentum residual remained similar, i.e. the 2 Tesla case of the 66% scenario and the 3 Tesla case of the 150% scenario.

It should be noted that the residuals are based on the surviving particles. The optics of the cooling channel will deviate from an optimal particle transport scenario due to the changes in the magnetic field strength or the magnet axis alignment as a result of the specific scenario considered. This can result in many particles being lost. It can also result in some particles simply not being reconstructed any more. For example, the 150% scenario had a similar number of virtual particles passing through the upstream tracker as the “No Change” scenario (17409 vs 17633), however approximately 1500 fewer particles were reconstructed (14917 vs 16558). While the residuals can show the bias of the momentum reconstruction, it can only do so for the distribution considered. Therefore, if the momentum reconstruction has a bias dependent on the distribution considered, then an acceptance bias is introduced if not all of the particles are reconstructed.

The MICE (0° Sol, 0° Tr, 0° MF) scenario considers the MICE magnetic field when there were no misalignments of any of the tracker, solenoid or magnetic axes. Various scenarios were then considered where each of these three axes were rotated by up to 2 degrees (approximately 35 mrad). The previous chapter showed that some of the MICE misalignments were greater than 0.4 degrees, though it wasn’t always clear if all of the misalignments were always implemented. It is unlikely however that the sum of all of the misalignments would be greater than 1 degree. The scenario that considers a 2-degree rotation of the magnetic axes considers what would happen in an extreme scenario.

For the scenarios that only consider the rotation of the magnetic axes, the number of particles travelling from TKU to TKD begins to decrease as the rotation of the magnetic axes increases (Table 5.3). The transmission losses could place an effective upper bound on the maximal possible misalignment of the magnetic axes when the real data is considered. Additionally, for the misaligned magnetic field scenarios, the muons gain less momentum as a result of the reconstruction and in fact lose a significant amount of momentum for the 2 degree scenario.

Similar scenarios were considered for the rotation of the solenoid and tracker axes, though it should be noted that the tracker axes are defined relative to the solenoid axes i.e. if the solenoid axes are misaligned by one degree, then the tracker axes need to be misaligned by minus one degree to keep the tracker axes aligned with the MICE global reference axis.

Further scenarios were also considered for the rotations of the tracker axes or for

the rotation of several axes. For the MICE magnetic field, the mean residuals do not change by large amounts unless the misalignments either become too large or too numerous (the percentage of transmitted and reconstructed particles should be kept in mind, however). The primary effect can be seen on the RMS residuals, which change significantly as a result of the misalignments.

To show that the bias of MICE reconstruction is not solely due to the MICE magnetic field, Monte Carlo simulations were run where the MICE magnetic field was replaced by a 9905 mm long constant solenoid field with a field radius of 516 mm. The centre of the constant solenoid field was placed at the centre of the absorber material location. The length of the solenoid field ensured that the trackers were housed within a constant solenoid field and that the muons were transported between the trackers using a constant solenoid field. Again, various magnetic field strengths were considered (2T, 3T and 4T) along with various rotations of the solenoid, tracker and magnetic axes. A TKU or TKD rotation of the magnetic axis corresponds to the same rotations the coils were assumed to be rotated by in the MICE experiment.

Similarly, as for the MICE magnetic field scenarios, the mean transverse momentum residuals showed distinctive non-zero biases, while the RMS residuals showed distinctive U-shapes (Figures 5.11 and 5.12). This shows that the reconstruction finds it difficult to obtain the correct transverse momentum of a particle at the first and last stations of each tracker. This is due to the MICE reconstruction's extended (linear) Kalman Filter not accounting for the changing trajectory of a muon from its perfect helical trajectory in a uniform solenoid field when it encounters energy loss, i.e. when the particles pass through the stations of the tracker. For the same reason, the incorrect phase advance is calculated for the muon, resulting in the incorrect longitudinal momentum being calculated. This results in the reconstruction effectively adding momentum to the particle as a function of longitudinal position within each tracker.

When considering like for like scenarios between the MICE magnetic field cases and the uniform solenoid field cases in Table 5.5, the magnitude of the bias is typically smaller for the uniform solenoid field scenarios than for the MICE magnetic field scenarios, showing that the non-uniformity of the MICE magnetic field does have an effect on the bias of the momentum reconstruction.

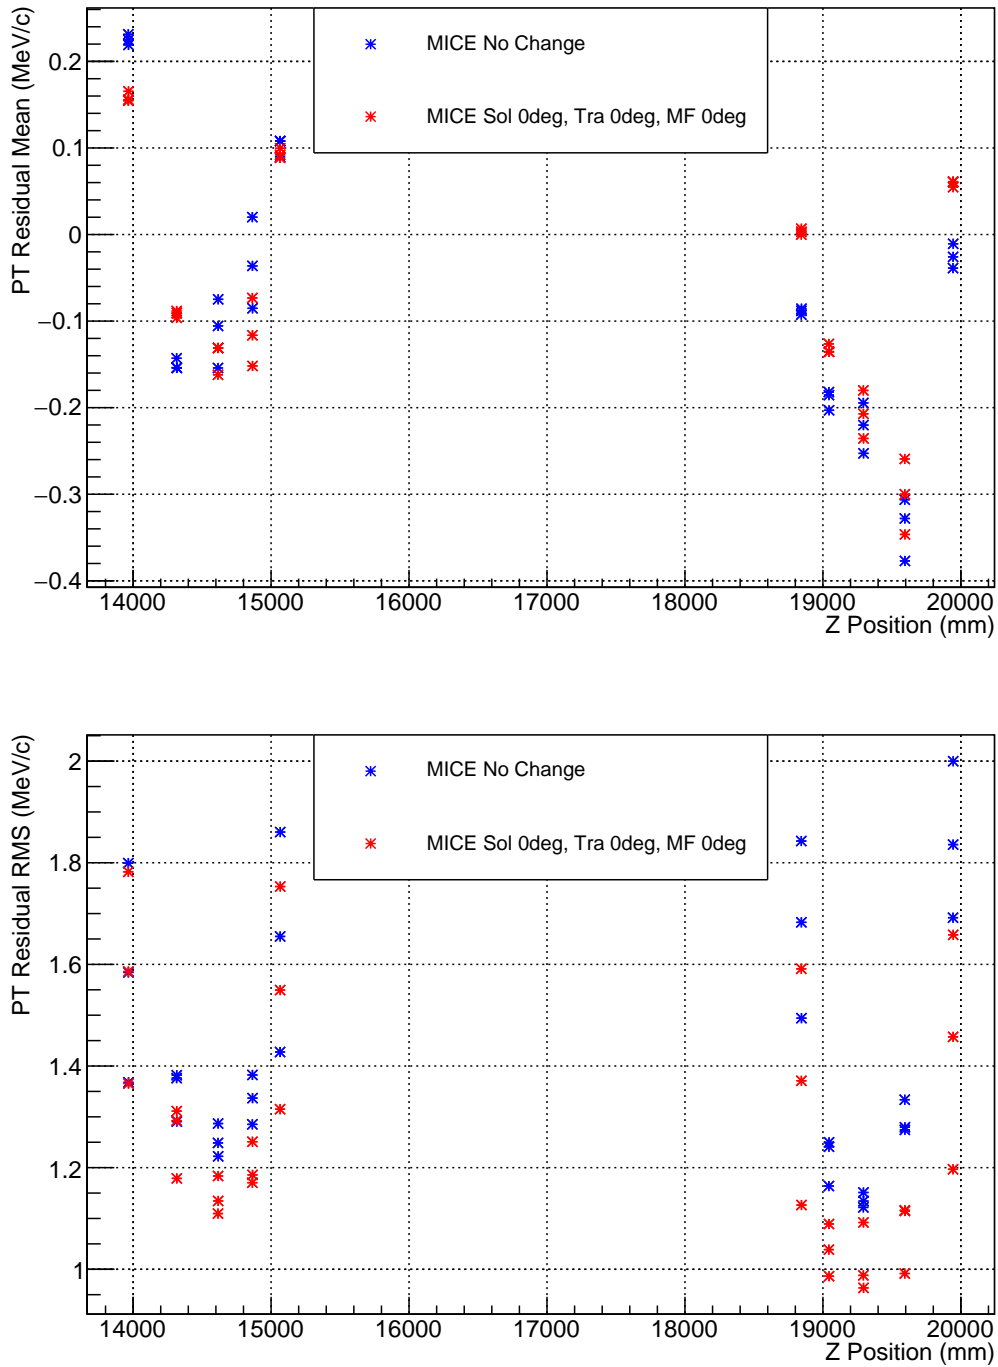


Figure 5.9: The mean (top) and RMS (bottom) transverse momentum residual at every plane of every station in the upstream and downstream trackers as a function of longitudinal position in the MICE cooling channel using the MICE alignment (blue) and when there were no misalignments (red).

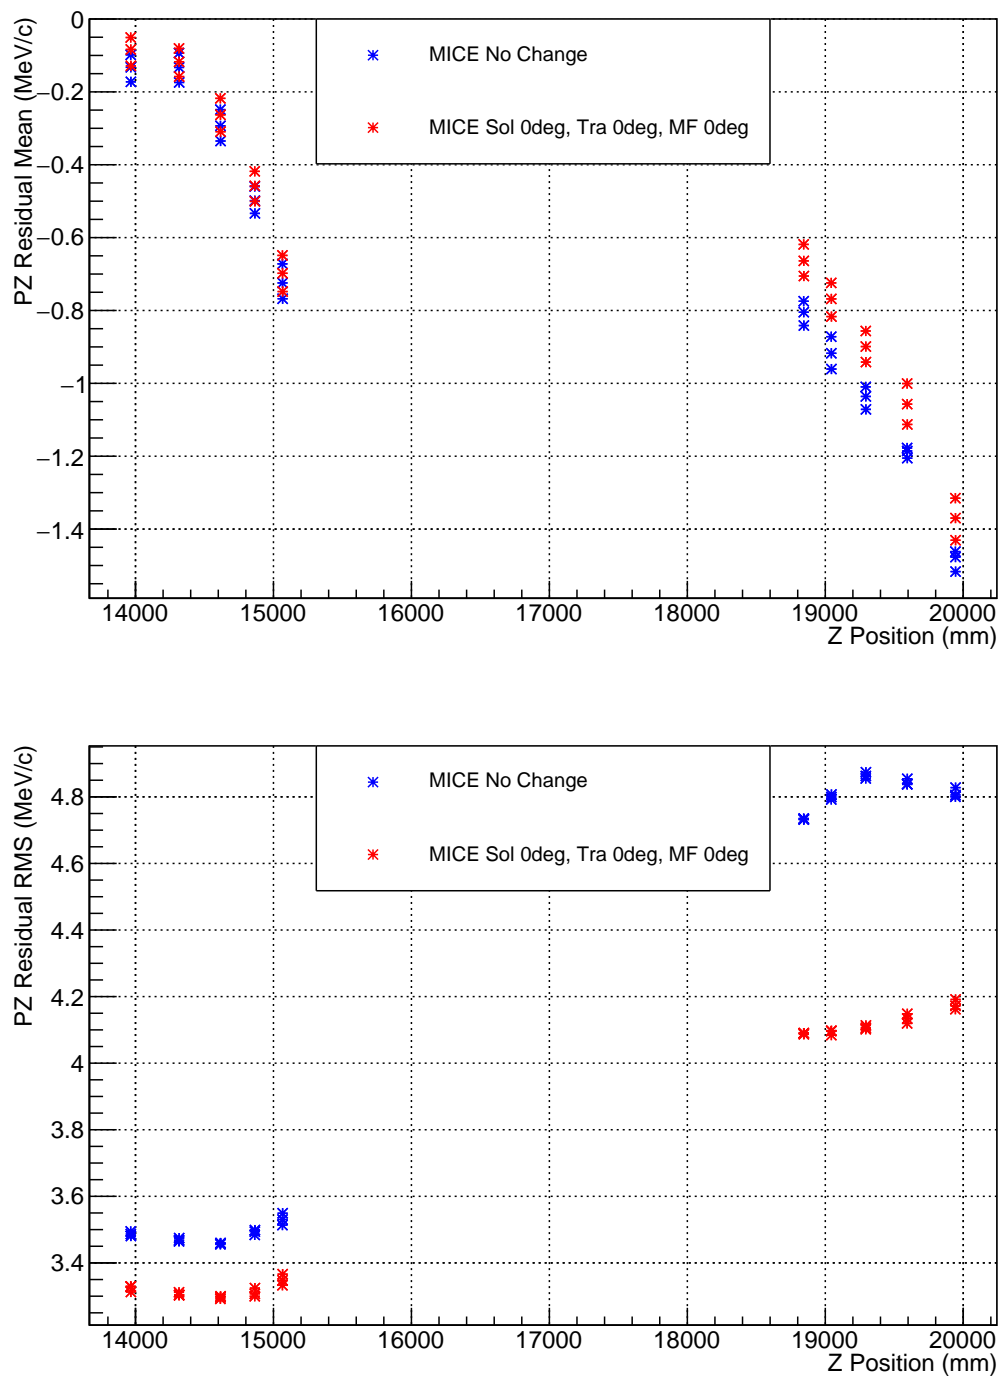


Figure 5.10: The mean (top) and RMS (bottom) longitudinal momentum residual at every plane of every station in the upstream and downstream trackers as a function of longitudinal position in the MICE cooling channel using the MICE alignment (blue) and when there were no misalignments (red).




Figure 5.11: The mean (top) and RMS (bottom) transverse momentum residual at every plane of every station in the upstream and downstream trackers as a function of longitudinal position in the MICE cooling channel. The MICE magnetic field has been replaced by a 2 Tesla (yellow), 3 Tesla (blue) or 4 Tesla (magenta) constant solenoid field, with the constant fields rotated by the same rotation as the magnetic axis misalignment of the upstream tracker.

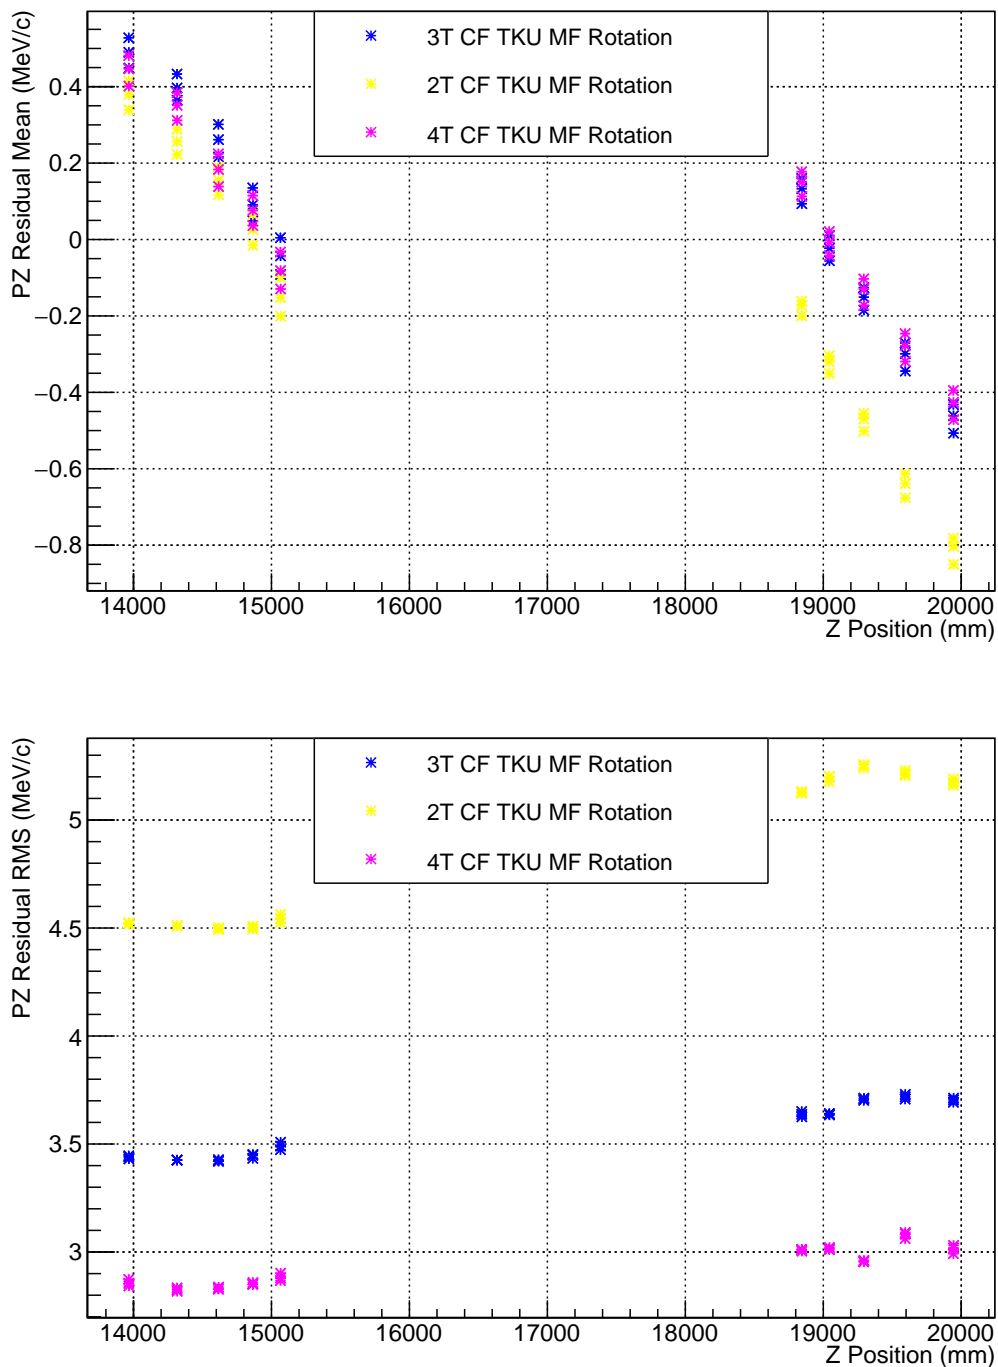


Figure 5.12: The mean (top) and RMS (bottom) longitudinal momentum residual at every plane of every station in the upstream and downstream trackers as a function of longitudinal position in the MICE cooling channel. The MICE magnetic field has been replaced by a 2 Tesla (yellow), 3 Tesla (blue) or 4 Tesla (magenta) constant solenoid field, with the constant fields rotated by the same rotation as the magnetic axis misalignment of the upstream tracker.

Table 5.3: For each scenario considered, the table shows the number and percentage of virtual particles that were reconstructed in each tracker, the percentage of particles that were transmitted between the trackers, and the number of particles that were reconstructed in the downstream tracker compared to the upstream tracker. MICE = MICE magnetic field, "x.xx TKU/TKD" = coil current scaling, CF = Constant Field, Sol = Solenoid Axis, Tr = Tracker Axis, MF = Magnetic Field Axis. Angles refer to rotations of those axes.

Scenario	TKU	TKD	TKU	TKD	TKU	TKD	Virtual	Recon
	Virtual	Virtual	Recon	Recon	Recon/	Recon/	Transmission	Transmission
	Entries	Entries	Entries	Entries	Virtual	Virtual		
MICE No Change	17633	13613	16558	13246	93.90 %	97.30 %	77.20 %	80.00 %
MICE 1.02 TKU 1.018 TKD	17641	13695	16487	13330	93.46 %	97.33 %	77.63 %	80.85 %
MICE 0.66 TKU 0.66 TKD	16919	7537	15878	6617	93.85 %	87.79 %	44.55 %	41.67 %
MICE 1.5 TKU 1.5 TKD	17409	9927	14917	9604	85.69 %	96.75 %	57.02 %	64.38 %
MICE 0° Sol, 0° Tr, 0° MF	17786	13732	16664	13438	93.69 %	97.86 %	77.21 %	80.64 %
MICE 0° Sol, 0° Tr, 1° MF	17734	12580	16372	12066	92.32 %	95.91 %	70.94 %	73.70 %
MICE 0° Sol, 0° Tr, 2° MF	17535	9693	13893	8664	79.23 %	89.38 %	55.28 %	62.36 %
MICE 1° Sol, -1° Tr, 0° MF	17562	13602	16466	13261	93.76 %	97.49 %	77.45 %	80.54 %
MICE 1° Sol, -1° Tr, 1° MF	17500	12000	16151	11574	92.29 %	96.45 %	68.57 %	71.66 %
MICE 1° Sol, 1° Tr, 1° MF	17530	12074	16131	11537	92.02 %	95.55 %	68.88 %	71.52 %
MICE 1° Sol, 1° Tr, 2° MF	17207	8933	16082	8445	93.46 %	94.54 %	51.91 %	52.51 %
3T CF TKU MF Rotation	17805	17107	16715	16438	93.88 %	96.09 %	96.08 %	98.34 %
3T CF No MF Rotation	17986	17287	16892	16544	93.92 %	95.70 %	96.11 %	97.94 %
3T CF TKD MF Rotation	17820	17146	16732	16496	93.89 %	96.21 %	96.22 %	98.59 %
3T CF 1° MF Rotation	17877	16228	16398	13518	91.73 %	83.30 %	90.78 %	82.44 %
3T CF 2° MF Rotation	16772	2545	11098	419	66.17 %	16.46 %	15.17 %	3.78 %
2T CF TKU MF Rotation	17432	16352	16444	15781	94.33 %	96.51 %	93.80 %	95.97 %
4T CF TKU MF Rotation	17889	17241	16259	16038	90.89 %	93.02 %	96.38 %	98.64 %
3T CF 1° Sol, -1° Tra, 0° MF	17799	17067	16742	16398	94.06 %	96.08 %	95.89 %	97.95 %
3T CF 1° Sol, -1° Tra, 1° MF	17615	15516	16115	13720	91.48 %	88.42 %	88.08 %	85.14 %
3T CF 1° Sol, 1° Tra, 1° MF	17588	15460	16331	12930	92.85 %	83.64 %	87.90 %	79.17 %
3T CF 1° Sol, 1° Tra, 2° MF	16696	3748	15484	916	92.74 %	24.22 %	22.45 %	5.92 %

Table 5.4: For each scenario considered, the table shows the mean and RMS transverse momentum residuals at the reference planes of the upstream and downstream trackers. MICE = MICE magnetic field, "x.xx TKU/TKD" = coil current scaling, CF = Constant Field, Sol = Solenoid Axis, Tr = Tracker Axis, MF = Magnetic Field Axis. Angles refer to rotations of those axes.

Scenario	TKU Entries	TKD Entries	TKU Mean	TKD Mean	TKU RMS	TKD RMS
MICE No Change	16558	13246	0.108 ± 0.014	-0.093 ± 0.016	1.860 ± 0.010	1.843 ± 0.011
MICE 1.02 TKU 1.018 TKD	16487	13330	0.109 ± 0.014	-0.102 ± 0.016	1.777 ± 0.010	1.857 ± 0.011
MICE 0.66 TKU 0.66 TKD	15878	6617	0.061 ± 0.014	0.073 ± 0.025	1.805 ± 0.010	2.004 ± 0.017
MICE 1.5 TKU 1.5 TKD	14917	9604	0.297 ± 0.019	-0.011 ± 0.019	2.342 ± 0.014	1.874 ± 0.014
MICE 0° Sol, 0° Tr, 0° MF	16664	13438	0.100 ± 0.014	-0.000 ± 0.014	1.753 ± 0.010	1.591 ± 0.010
MICE 0° Sol, 0° Tr, 1° MF	16372	12066	-0.011 ± 0.019	0.175 ± 0.021	2.404 ± 0.013	2.306 ± 0.015
MICE 0° Sol, 0° Tr, 2° MF	13893	8664	-0.631 ± 0.039	1.190 ± 0.043	4.510 ± 0.028	3.987 ± 0.031
MICE 1° Sol, -1° Tr, 0° MF	16466	13261	0.116 ± 0.013	-0.024 ± 0.014	1.615 ± 0.009	1.560 ± 0.010
MICE 1° Sol, -1° Tr, 1° MF	16151	11574	0.036 ± 0.019	-0.269 ± 0.022	2.377 ± 0.013	2.358 ± 0.016
MICE 1° Sol, 1° Tr, 1° MF	16131	11537	-0.892 ± 0.030	-0.595 ± 0.028	3.703 ± 0.021	2.970 ± 0.020
MICE 1° Sol, 1° Tr, 2° MF	16082	8445	-0.168 ± 0.018	-0.210 ± 0.015	2.216 ± 0.012	1.405 ± 0.011
3T CF TKU MF Rotation	16715	16438	-0.002 ± 0.014	-0.161 ± 0.016	1.821 ± 0.010	2.066 ± 0.011
3T CF No MF Rotation	16892	16544	-0.008 ± 0.013	-0.043 ± 0.014	1.730 ± 0.009	1.736 ± 0.010
3T CF TKD MF Rotation	16732	16496	-0.028 ± 0.014	-0.107 ± 0.015	1.822 ± 0.001	1.889 ± 0.010
3T CF 1° MF Rotation	16398	13518	-0.088 ± 0.023	-0.406 ± 0.026	2.951 ± 0.016	3.005 ± 0.019
3T CF 2° MF Rotation	11098	419	-0.858 ± 0.059	-1.773 ± 0.331	5.998 ± 0.042	6.438 ± 0.234
2T CF TKU MF Rotation	16444	15781	-0.003 ± 0.014	-0.135 ± 0.016	1.749 ± 0.010	1.958 ± 0.011
4T CF TKU MF Rotation	16259	16038	-0.113 ± 0.017	-0.221 ± 0.020	2.108 ± 0.012	2.536 ± 0.014
3T CF 1° Sol, -1° Tra, 0° MF	16742	16398	-0.013 ± 0.013	-0.053 ± 0.013	1.700 ± 0.009	1.700 ± 0.009
3T CF 1° Sol, -1° Tra, 1° MF	16115	13720	0.029 ± 0.023	-0.476 ± 0.026	2.903 ± 0.016	3.059 ± 0.019
3T CF 1° Sol, 1° Tra, 1° MF	16331	12930	-0.849 ± 0.024	-0.758 ± 0.027	3.016 ± 0.017	3.002 ± 0.019
3T CF 1° Sol, 1° Tra, 2° MF	15484	916	0.060 ± 0.015	-0.242 ± 0.094	1.841 ± 0.010	2.785 ± 0.067

Table 5.5: For each scenario considered, the table shows the mean and RMS longitudinal momentum residuals at the reference planes of the upstream and downstream trackers. MICE = MICE magnetic field, "x.xx TKU/TKD" = coil current scaling, CF = Constant Field, Sol = Solenoid Axis, Tr = Tracker Axis, MF = Magnetic Field Axis. Angles refer to rotations of those axes.

Scenario	TKU Entries	TKD Entries	TKU Mean	TKD Mean	TKU RMS	TKD RMS
MICE No Change	16558	13246	-0.768 ± 0.028	-0.775 ± 0.042	3.549 ± 0.020	4.735 ± 0.030
MICE 1.02 TKU 1.018 TKD	16487	13330	-0.715 ± 0.027	-0.839 ± 0.042	3.430 ± 0.019	4.706 ± 0.029
MICE 0.66 TKU 0.66 TKD	15878	6617	-0.909 ± 0.037	-0.466 ± 0.083	4.617 ± 0.026	6.514 ± 0.058
MICE 1.5 TKU 1.5 TKD	14917	9604	-0.811 ± 0.023	-0.687 ± 0.036	2.717 ± 0.017	3.482 ± 0.026
MICE 0° Sol, 0° Tr, 0° MF	16664	13438	-0.748 ± 0.026	-0.619 ± 0.036	3.366 ± 0.019	4.090 ± 0.025
MICE 0° Sol, 0° Tr, 1° MF	16372	12066	-0.632 ± 0.043	-0.077 ± 0.060	5.195 ± 0.030	6.298 ± 0.042
MICE 0° Sol, 0° Tr, 2° MF	13893	8664	0.755 ± 0.096	3.888 ± 0.106	9.547 ± 0.068	8.872 ± 0.075
MICE 1° Sol, -1° Tr, 0° MF	16466	13261	-0.737 ± 0.026	-0.667 ± 0.036	3.397 ± 0.019	4.066 ± 0.025
MICE 1° Sol, -1° Tr, 1° MF	16151	11574	-0.967 ± 0.044	-1.407 ± 0.063	5.309 ± 0.031	6.461 ± 0.044
MICE 1° Sol, 1° Tr, 1° MF	16131	11537	-0.005 ± 0.054	-1.034 ± 0.072	6.345 ± 0.038	7.311 ± 0.051
MICE 1° Sol, 1° Tr, 2° MF	16082	8445	-0.538 ± 0.028	-0.894 ± 0.047	3.490 ± 0.020	4.284 ± 0.033
3T CF TKU MF Rotation	16715	16438	-0.092 ± 0.028	0.160 ± 0.029	3.508 ± 0.020	3.651 ± 0.021
3T CF No MF Rotation	16892	16544	0.035 ± 0.027	0.224 ± 0.025	3.485 ± 0.019	3.179 ± 0.018
3T CF TKD MF Rotation	16732	16496	-0.147 ± 0.028	0.181 ± 0.028	3.538 ± 0.020	3.532 ± 0.020
3T CF 1° MF Rotation	16398	13518	0.129 ± 0.048	-0.108 ± 0.053	5.748 ± 0.034	5.753 ± 0.037
3T CF 2° MF Rotation	11098	419	2.592 ± 0.126	3.261 ± 0.890	10.277 ± 0.089	11.360 ± 0.629
2T CF TKU MF Rotation	16444	15781	-0.201 ± 0.036	-0.162 ± 0.042	4.563 ± 0.026	5.123 ± 0.030
4T CF TKU MF Rotation	16259	16038	-0.129 ± 0.023	0.178 ± 0.025	2.902 ± 0.017	3.013 ± 0.018
3T CF 1° Sol, -1° Tra, 0° MF	16742	16398	0.008 ± 0.027	0.211 ± 0.026	3.429 ± 0.019	3.27 ± 0.018
3T CF 1° Sol, -1° Tra, 1° MF	16115	13720	-0.085 ± 0.049	-0.094 ± 0.052	5.805 ± 0.034	5.657 ± 0.037
3T CF 1° Sol, 1° Tra, 1° MF	16331	12930	0.374 ± 0.048	-0.573 ± 0.052	5.782 ± 0.034	5.512 ± 0.037
3T CF 1° Sol, 1° Tra, 2° MF	15484	916	0.248 ± 0.028	-0.610 ± 0.159	3.419 ± 0.020	4.498 ± 0.113

5.3.2 Momentum dependence of the MICE reconstruction

For the MICE experiment, a number of cuts were applied to the data in the MICE analysis. These included applying a momentum cut to only consider particles within a narrow momentum range (135 – 145 MeV). Assuming a narrow momentum range and a uniform magnetic field allowed MICE to assume that the transverse and longitudinal components of the beam could be separated⁷⁸ (i.e. that a linear transfer matrix could be used), and that the time component could be ignored (due to the narrow momentum range). Therefore, for the assumptions MICE used, the demonstration of transverse ionisation cooling at MICE only required the transverse components of the beam (x, y, p_x, p_y) .

Tables 5.6, 5.7 and 5.8 have applied similar cuts to the virtual muons. Only virtual muons that had a momentum of between 130 MeV and 150 MeV at the end of the upstream tracker were considered. They also must have passed through all of the tracker planes within a tracker, at a tracker plane radius of less than 150 mm. This results in the number of transmitted particles drastically reducing. However, of the particles that were transmitted, the percentage of virtual particles that were reconstructed has increased. It should be noted that the radial and momentum cuts were performed on the virtual particles. If a high number of virtual particles were transmitted to the downstream tracker (i.e. low losses), then the number of reconstructed particles in the downstream tracker could be larger than the number of reconstructed particles in the upstream tracker (leading to a Recon Transmission greater than 100% in Table 5.6).

A noticeable effect on the transverse and longitudinal momentum residuals can be seen, though the trends aren't always clear. For the MICE “No Change” scenario and the scaled coil currents scenario considered for the MICE Nature paper (“MICE 1.02 TKU 1.018 TKD”), the magnitude of the residuals have become stronger, while the magnitude of the momentum added to the muons between the trackers has also increased. When no misalignments were considered, however, the differences on the residuals as a result of the radial and momentum cuts appear minor.

The MICE reconstruction therefore has an inherent momentum bias even when there are no misalignments. This bias is then enhanced or changed when misalignments are introduced. This bias can also be dependent on the input distribution. This can be more clearly seen when the momentum distributions are considered.

Fig. 5.13 shows the transverse momentum distributions at the upstream tracker reference plane (after the radial and momentum cuts have been applied) for the MICE magnetic field scenarios when there were no misalignments (top), and when the magnetic fields in the trackers were rotated by 2 degrees (bottom). As the magnetic field misalignment increases, the plots show that the mean of the reconstructed

transverse momentum becomes increasingly larger than the mean of the virtual transverse momentum. This is primarily driven by low transverse momenta muons being reconstructed with a higher transverse momentum. As the virtual planes and tracker planes are aligned, the effect is solely due to the rotation of the magnetic axes. As the magnetic axes are misaligned from the tracker axes, the muon's transverse and longitudinal momentum components are no longer fully separable in the tracker reference frame, and thus create a transverse momentum kick. As only particles in a narrow momentum range were considered (130-150 MeV), the transverse momentum kick is primarily noticeable for low transverse momenta particles, as the transverse momentum kick is a far greater percentage of the particle's initial transverse momentum. If the misalignment is large enough, then a hole is formed at the centre of the reconstructed transverse momentum distribution in (p_x, p_y) space where very few particles are reconstructed, even if the (virtual) distribution contained a significant number of low transverse momenta particles. This can be seen at the bottom of Fig. 5.13. The low transverse momentum hole seen in the MICE data could therefore be explained by a misalignment of the magnetic axes.

Similar plots for the longitudinal momentum distribution are shown in Fig. 5.14. When there are no misalignments, the mean bias on the reconstructed momentum is just over 1 MeV. This seems to be primarily driven by a few particles being reconstructed with a higher momentum pulling the distribution. The overall reconstructed momentum distribution (red line) is similar to the corresponding virtual momentum distribution (blue line), though the distribution has a lower peak and a wider spread, indicated by the larger RMS. For the 2 degree magnetic axis rotation, the RMS drastically increases. Even more particles are reconstructed at a higher longitudinal momentum, which further increases the gap between the mean residual of the virtual and reconstructed distributions.

The main part of the reconstructed distribution no longer coincides with the corresponding virtual distribution. A double peaked distribution has formed, with some particles losing momentum, while some other particles have gained a significant amount of momentum as a result of the reconstruction.

For many of the scenarios considered, a significant number of particles were either not transmitted through the MICE cooling channel or were not reconstructed in the trackers. Such transmission losses can bias any ionisation cooling result and will be studied in chapter 7.

For the demonstration of emittance exchange, a time component is required, however the MICE reconstruction doesn't reconstruct a time component. The following section will outline how the time component of a particle was reconstructed for the analysis in chapter 7.

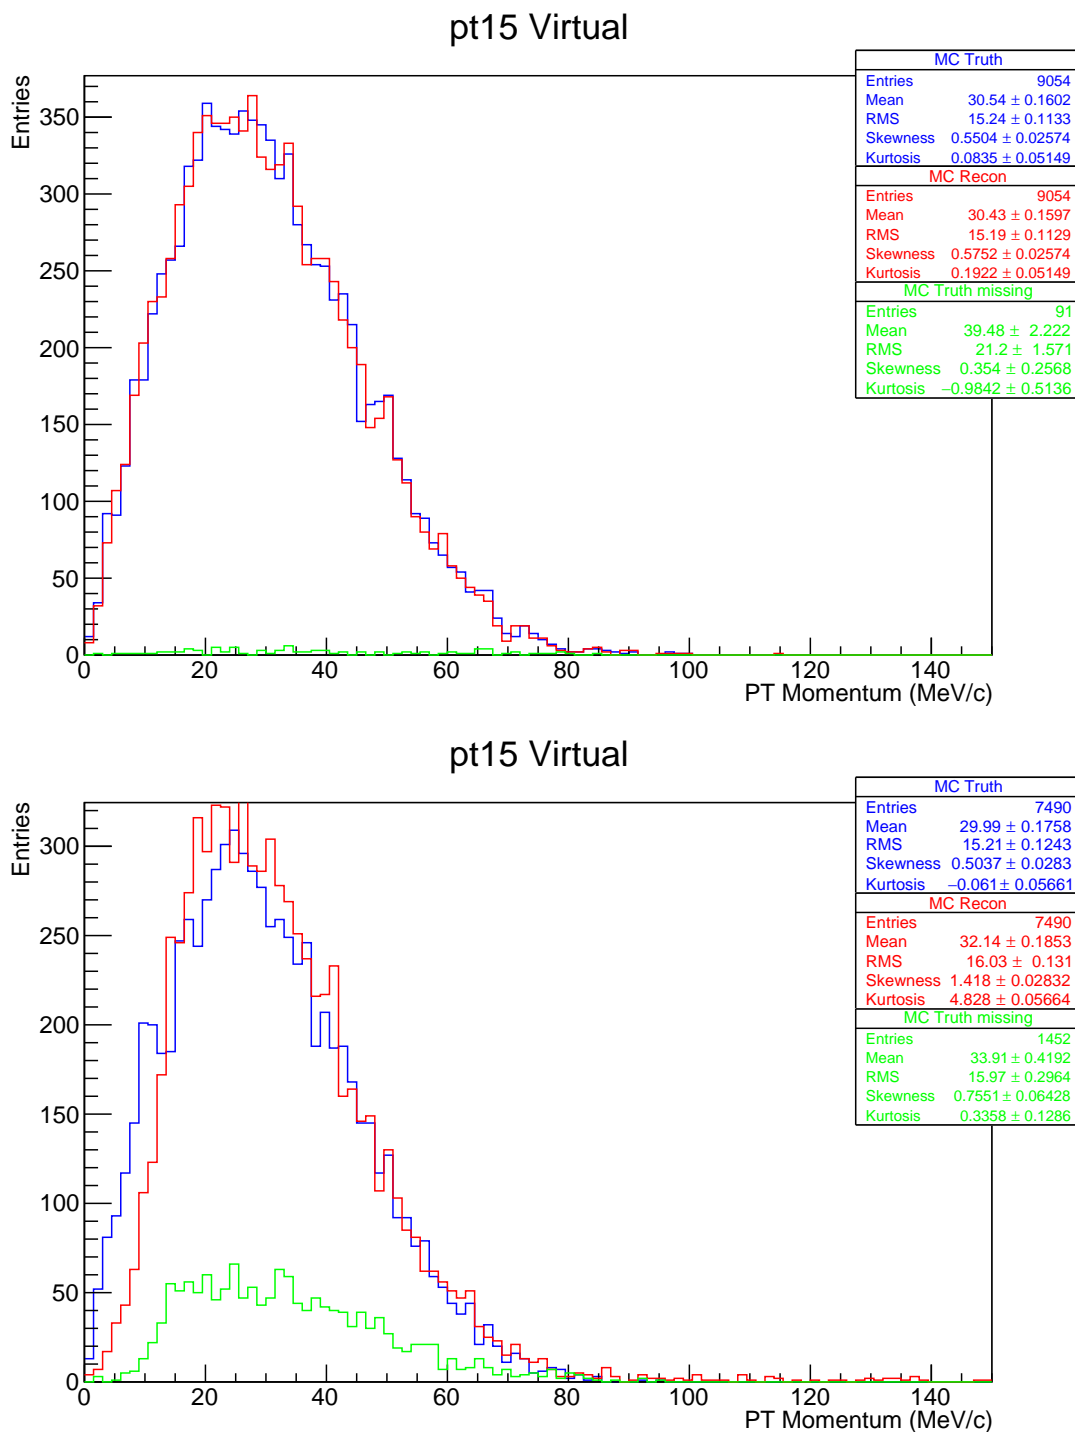


Figure 5.13: The transverse momentum distributions at the upstream tracker reference plane for the no misalignment scenario (MICE 0° Sol, 0° Tr, 0° MF) (top) and the 2 degree magnetic axes rotation scenario (MICE 0° Sol, 0° Tr, 2° MF) (bottom), after the radial (less than 150 mm at a tracker plane, for all tracker planes within a tracker) and momentum (130 – 150 MeV) cuts have been applied.




Figure 5.14: The longitudinal momentum distributions at the upstream tracker reference plane for the no misalignment scenario (MICE 0° Sol, 0° Tr, 0° MF) (top) and the 2 degree magnetic axes rotation scenario (MICE 0° Sol, 0° Tr, 2° MF) (bottom), after the radial (less than 150 mm at a tracker plane, for all tracker planes within a tracker) and momentum (130 – 150 MeV) cuts have been applied.

Table 5.6: The table is the same as Table 5.3, except that it only considers muons that have passed through all of the tracker planes within a tracker, at a tracker plane radius of less than 150 mm, that also had a momentum of between 130 MeV and 150 MeV at the end of the upstream tracker. For cases where the virtual transmission is high, the recon transmission can be greater than 100% if the reconstruction efficiency is larger in the downstream tracker than in the upstream tracker.

Scenario	TKU	TKD	TKU	TKD	TKU	TKD	Virtual Transmission	Recon Transmission
	Virtual Entries	Virtual Entries	Recon Entries	Recon Entries	Recon/ Virtual	Recon/ Virtual		
MICE No Change	9115	7765	9022	7693	98.98 %	99.07 %	85.19 %	85.27 %
MICE 1.02 TKU 1.018 TKD	9186	7790	9078	7692	98.82 %	98.74 %	84.80 %	84.73 %
MICE 0.66 TKU 0.66 TKD	8319	3111	8171	3050	98.22 %	98.04 %	37.40 %	37.33 %
MICE 1.5 TKU 1.5 TKD	9022	5880	8689	5830	96.31 %	99.15 %	65.17 %	67.10 %
MICE 0° Sol, 0° Tr, 0° MF	9145	7776	9054	7713	99.00 %	99.19 %	85.03 %	85.19 %
MICE 0° Sol, 0° Tr, 1° MF	9141	6829	8925	6751	97.64 %	98.86 %	74.71 %	75.64 %
MICE 0° Sol, 0° Tr, 2° MF	8942	4629	7490	4475	83.76 %	96.67 %	51.77 %	59.75 %
MICE 1° Sol, -1° Tr, 0° MF	9096	7760	8989	7683	98.82 %	99.01 %	85.31 %	85.47 %
MICE 1° Sol, -1° Tr, 1° MF	9074	6353	8828	6277	97.29 %	98.80 %	70.01 %	71.10 %
MICE 1° Sol, 1° Tr, 1° MF	9274	6524	9043	6394	97.51 %	98.01 %	70.35 %	70.71 %
MICE 1° Sol, 1° Tr, 2° MF	9068	3401	8916	3358	98.32 %	98.74 %	37.51 %	37.66 %
3T CF TKU MF Rotation	9231	9214	9132	9138	98.93 %	99.18 %	99.82 %	100.07 %
3T CF No MF Rotation	9271	9264	9167	9192	98.88 %	99.22 %	99.92 %	100.27 %
3T CF TKD MF Rotation	9227	9218	9116	9136	98.80 %	99.11 %	99.90 %	100.22 %
3T CF 1° MF Rotation	9194	6009	8909	5901	96.90 %	98.20 %	65.36 %	66.24 %
3T CF 2° MF Rotation	8483	40	5929	32	69.89 %	80.00 %	0.47 %	0.54 %
2T CF TKU MF Rotation	8707	8601	8565	8485	98.37 %	98.65 %	98.78 %	99.07 %
4T CF TKU MF Rotation	9307	9301	9172	9193	98.55 %	98.84 %	99.94 %	100.23 %
3T CF 1° Sol, -1° Tra, 0° MF	9243	9231	9127	9167	98.74 %	99.31 %	99.87 %	100.44 %
3T CF 1° Sol, -1° Tra, 1° MF	9150	5843	8849	5719	96.71 %	97.88 %	63.86 %	64.63 %
3T CF 1° Sol, 1° Tra, 1° MF	9356	6005	9159	5813	97.89 %	96.80 %	64.18 %	63.47 %
3T CF 1° Sol, 1° Tra, 2° MF	8682	46	8492	46	97.81 %	100.00 %	0.53 %	0.54 %

Table 5.7: For each scenario considered, the table shows the mean and RMS transverse momentum residuals at the reference planes of the upstream and downstream trackers for muons that have passed through all of the tracker planes within a tracker, at a tracker plane radius of less than 150 mm, that also had a momentum of between 130 MeV and 150 MeV at the end of the upstream tracker. MICE = MICE magnetic field, "x.xx TKU/TKD" = coil current scaling, CF = Constant Field, Sol = Solenoid Axis, Tr = Tracker Axis, MF = Magnetic Field Axis. Angles refer to rotations of those axes.

Scenario	TKU Entries	TKD Entries	TKU Mean	TKD Mean	TKU RMS	TKD RMS
MICE No Change	9022	7693	0.194 ± 0.017	-0.111 ± 0.021	1.654 ± 0.012	1.836 ± 0.015
MICE 1.02 TKU 1.018 TKD	9078	7692	0.196 ± 0.017	-0.131 ± 0.021	1.644 ± 0.012	1.795 ± 0.015
MICE 0.66 TKU 0.66 TKD	8171	3050	0.017 ± 0.020	0.149 ± 0.037	1.772 ± 0.014	2.057 ± 0.026
MICE 1.5 TKU 1.5 TKD	8689	5830	0.391 ± 0.020	-0.016 ± 0.024	1.839 ± 0.014	1.805 ± 0.017
MICE 0° Sol, 0° Tr, 0° MF	9054	7713	0.136 ± 0.017	-0.006 ± 0.017	1.570 ± 0.012	1.532 ± 0.012
MICE 0° Sol, 0° Tr, 1° MF	8925	6751	0.113 ± 0.024	0.252 ± 0.028	2.288 ± 0.017	2.303 ± 0.020
MICE 0° Sol, 0° Tr, 2° MF	7490	4475	-0.398 ± 0.051	1.467 ± 0.059	4.360 ± 0.036	3.941 ± 0.042
MICE 1° Sol, -1° Tr, 0° MF	8989	7683	0.152 ± 0.015	-0.015 ± 0.018	1.441 ± 0.011	1.562 ± 0.013
MICE 1° Sol, -1° Tr, 1° MF	8828	6277	0.036 ± 0.024	-0.274 ± 0.031	2.231 ± 0.017	2.429 ± 0.022
MICE 1° Sol, 1° Tr, 1° MF	9043	6394	-0.781 ± 0.037	-0.702 ± 0.037	3.520 ± 0.026	2.975 ± 0.026
MICE 1° Sol, 1° Tr, 2° MF	8916	3358	-0.122 ± 0.022	-0.264 ± 0.027	2.030 ± 0.015	1.554 ± 0.019
3T CF TKU MF Rotation	9132	9138	0.099 ± 0.017	-0.114 ± 0.020	1.668 ± 0.012	1.884 ± 0.014
3T CF No MF Rotation	9167	9192	0.053 ± 0.017	-0.028 ± 0.017	1.580 ± 0.012	1.615 ± 0.012
3T CF TKD MF Rotation	9116	9136	0.066 ± 0.018	-0.082 ± 0.018	1.682 ± 0.012	1.756 ± 0.013
3T CF 1° MF Rotation	8909	5901	0.114 ± 0.030	-0.479 ± 0.038	2.797 ± 0.021	2.907 ± 0.027
3T CF 2° MF Rotation	5929	32	-0.497 ± 0.080	-6.075 ± 0.913	5.983 ± 0.057	4.655 ± 0.646
2T CF TKU MF Rotation	8565	8485	-0.049 ± 0.018	-0.172 ± 0.021	1.690 ± 0.013	1.943 ± 0.015
4T CF TKU MF Rotation	9172	9193	-0.072 ± 0.018	-0.194 ± 0.023	1.748 ± 0.013	2.232 ± 0.017
3T CF 1° Sol, -1° Tra, 0° MF	9127	9167	0.048 ± 0.016	-0.018 ± 0.016	1.520 ± 0.011	1.490 ± 0.011
3T CF 1° Sol, -1° Tra, 1° MF	8849	5719	0.065 ± 0.030	-0.598 ± 0.039	2.771 ± 0.021	2.941 ± 0.028
3T CF 1° Sol, 1° Tra, 1° MF	9159	5813	-0.767 ± 0.030	-0.895 ± 0.039	2.847 ± 0.021	2.941 ± 0.028
3T CF 1° Sol, 1° Tra, 2° MF	8492	46	0.102 ± 0.018	-0.042 ± 0.099	1.623 ± 0.012	0.643 ± 0.070

Table 5.8: For each scenario considered, the table shows the mean and RMS longitudinal momentum residuals at the reference planes of the upstream and downstream trackers for muons that have passed through all of the tracker planes within a tracker, at a tracker plane radius of less than 150 mm, that also had a momentum of between 130 MeV and 150 MeV at the end of the upstream tracker. MICE = MICE magnetic field, "x.xx TKU/TKD" = coil current scaling, CF = Constant Field, Sol = Solenoid Axis, Tr = Tracker Axis, MF = Magnetic Field Axis. Angles refer to rotations of those axes.

Scenario	TKU Entries	TKD Entries	TKU Mean	TKD Mean	TKU RMS	TKD RMS
MICE No Change	9022	7693	-0.807 ± 0.039	-0.913 ± 0.058	3.630 ± 0.027	4.935 ± 0.041
MICE 1.02 TKU 1.018 TKD	9078	7692	-0.727 ± 0.037	-0.992 ± 0.057	3.504 ± 0.026	4.924 ± 0.041
MICE 0.66 TKU 0.66 TKD	8171	3050	-0.874 ± 0.054	-0.152 ± 0.126	4.805 ± 0.038	6.730 ± 0.089
MICE 1.5 TKU 1.5 TKD	8689	5830	-0.815 ± 0.028	-0.671 ± 0.043	2.607 ± 0.020	3.249 ± 0.031
MICE 0° Sol, 0° Tr, 0° MF	9054	7713	-0.714 ± 0.036	-0.660 ± 0.050	3.376 ± 0.025	4.343 ± 0.035
MICE 0° Sol, 0° Tr, 1° MF	8925	6751	-0.593 ± 0.058	0.100 ± 0.085	5.224 ± 0.041	6.664 ± 0.060
MICE 0° Sol, 0° Tr, 2° MF	7490	4475	1.093 ± 0.133	5.149 ± 0.147	9.665 ± 0.094	8.702 ± 0.104
MICE 1° Sol, -1° Tr, 0° MF	8989	7683	-0.699 ± 0.036	-0.729 ± 0.049	3.425 ± 0.026	4.262 ± 0.035
MICE 1° Sol, -1° Tr, 1° MF	8828	6277	-1.099 ± 0.060	-1.468 ± 0.093	5.402 ± 0.042	6.965 ± 0.065
MICE 1° Sol, 1° Tr, 1° MF	9043	6394	0.515 ± 0.070	-1.058 ± 0.104	6.215 ± 0.050	7.824 ± 0.074
MICE 1° Sol, 1° Tr, 2° MF	8916	3358	-0.490 ± 0.038	-0.859 ± 0.087	3.510 ± 0.027	4.908 ± 0.061
3T CF TKU MF Rotation	9132	9138	-0.088 ± 0.038	0.181 ± 0.038	3.590 ± 0.027	3.587 ± 0.027
3T CF No MF Rotation	9167	9192	0.060 ± 0.037	0.243 ± 0.033	3.518 ± 0.026	3.158 ± 0.024
3T CF TKD MF Rotation	9116	9136	-0.207 ± 0.039	0.186 ± 0.038	3.676 ± 0.028	3.551 ± 0.027
3T CF 1° MF Rotation	8909	5901	0.189 ± 0.064	-0.238 ± 0.080	5.690 ± 0.045	5.713 ± 0.056
3T CF 2° MF Rotation	5929	32	3.310 ± 0.170	-13.506 ± 1.341	10.128 ± 0.120	3.793 ± 0.948
2T CF TKU MF Rotation	8565	8485	-0.190 ± 0.052	-0.175 ± 0.058	4.719 ± 0.037	5.222 ± 0.041
4T CF TKU MF Rotation	9172	9193	-0.112 ± 0.030	0.237 ± 0.031	2.858 ± 0.021	2.933 ± 0.022
3T CF 1° Sol, -1° Tra, 0° MF	9127	9167	0.031 ± 0.037	0.219 ± 0.035	3.509 ± 0.026	3.330 ± 0.025
3T CF 1° Sol, -1° Tra, 1° MF	8849	5719	-0.323 ± 0.067	-0.223 ± 0.083	5.896 ± 0.047	5.790 ± 0.058
3T CF 1° Sol, 1° Tra, 1° MF	9159	5813	0.754 ± 0.063	-0.754 ± 0.080	5.705 ± 0.045	5.659 ± 0.057
3T CF 1° Sol, 1° Tra, 2° MF	8492	46	0.279 ± 0.038	0.083 ± 0.827	3.453 ± 0.027	5.360 ± 0.585

5.4 Time reconstruction

If two muons in the MICE cooling channel have the same overall momentum, then the muon's time of flight between every station of the upstream and downstream trackers will be approximately the same. However, when the two muons encounter an absorber material, the muons will undergo energy loss and hit the stations of the downstream tracker at different times, i.e. a time dependence. For individual particles, the time dependence doesn't play a significant role. However, for a collection of particles this is not necessarily the case. Consider a collection of particles that hit the upstream tracker reference plane at the same time. After encountering the absorber material, the collection of particles will hit the downstream tracker reference plane at all different times depending on their momentum.

This poses a problem for the emittance measurement in MICE, which ignores the time dependence of the beam and only depends on the transverse components of the beam. Such an approximation can be made when there is no absorber material present in the MICE cooling channel, as the muons will hit both reference planes at approximately the same time. When an absorber material is present, however, there is a time delay between when the fastest and slowest muons hit the downstream tracker reference plane. This means the transverse beam components of the slower muons will continue to change as they continue on their spiral trajectory, until they reach the downstream tracker reference plane. This poses a problem for the emittance measurement which assumes there is only a narrow spread in the momenta of the muon beam, i.e. no time dependence. The issues this poses will be further investigated in chapter 7. For now, though, a reconstruction procedure for the time component will be introduced (MICE does not reconstruct the time component).

Using the equations for the energy and longitudinal momentum of a muon (whose longitudinal momentum is far greater than its transverse momentum),

$$E = \gamma m_0 c^2 \quad (5.1)$$

$$p_z = \gamma m_0 v_z \quad (5.2)$$

the longitudinal velocity of the muon can then be given by

$$\frac{\Delta z}{\Delta t} = v_z = \frac{c p_z}{E} \quad (5.3)$$

where γ is the Lorentz factor, m_0 is the rest mass of the muon, c is the speed of light and Δz is the longitudinal distance traversed by the muon in a time period Δt .

The time period can then be expressed as

$$\Delta t \approx \frac{\Delta z E}{c p_z} \quad (5.4)$$

The time period between two tracker stations is then simply given by the longitudinal distance between the stations, and the momentum and energy of the muon as it leaves the first station. For the reconstruction, the time coordinate will be given in a reference system that is triggered by the TOF1 station preceding the upstream tracker. The time coordinate at the first station of the upstream tracker is then given by the time period between TOF1 and the first station of the upstream tracker, with the time coordinate at each subsequent station taking the time coordinate at the previous station and then adding the time period between those two stations. At the upstream tracker reference plane, the time coordinate is given by

$$t_{S1}^{TKU} = \Delta t_{TOF1 \rightarrow S5} + \Delta t_{S5 \rightarrow S4} + \Delta t_{S4 \rightarrow S3} + \Delta t_{S3 \rightarrow S2} + \Delta t_{S2 \rightarrow S1} \quad (5.5)$$

$$\begin{aligned} t_{S1}^{TKU} = & \frac{\Delta z_{TOF1 \rightarrow S5} (E_{S5} + (E_{S5} - E_{S4}))}{c(p_{z,S5} + (p_{z,S5} - p_{z,S4}))} + \frac{\Delta z_{S5 \rightarrow S4} E_{S5}}{c p_{z,S5}} \\ & + \frac{\Delta z_{S4 \rightarrow S3} E_{S4}}{c p_{z,S4}} + \frac{\Delta z_{S3 \rightarrow S2} E_{S3}}{c p_{z,S3}} + \frac{\Delta z_{S2 \rightarrow S1} E_{S2}}{c p_{z,S2}} \end{aligned} \quad (5.6)$$

Note, as TOF1 is not a tracker station, the energy and momentum at the exit of TOF1 was calculated slightly differently. It takes the energy and momentum of the muon at station 5, and adds the lost energy and momentum between stations 5 and 4 of the upstream tracker to the muon, to assume the energy and momentum of the muon before it has struck station 5. This treatment ignores the energy lost by the muon at the diffuser, though this isn't a major concern at the moment.

The time coordinates for the downstream tracker stations are calculated slightly differently. As the time of flight is measured between TOF1 and TOF2, it uses this time period as the starting point. The time period between TOF2 and each station of the downstream tracker is then subtracted from the initial starting point to arrive at the time coordinate for each station of the downstream tracker. The time coordinate at the downstream tracker reference plane is then given by

$$t_{S1}^{TKD} = \Delta t_{TOF1 \rightarrow TOF2} - \Delta t_{S5 \rightarrow TOF2} - \Delta t_{S5 \rightarrow S4} - \Delta t_{S4 \rightarrow S3} - \Delta t_{S3 \rightarrow S2} - \Delta t_{S2 \rightarrow S1} \quad (5.7)$$

$$\begin{aligned} t_{S1}^{TKD} = & \Delta t_{TOF1 \rightarrow TOF2} - \frac{\Delta z_{S5 \rightarrow TOF2} E_{S5}}{c p_{z,S5}} - \frac{\Delta z_{S4 \rightarrow S5} E_{S4}}{c p_{z,S4}} \\ & - \frac{\Delta z_{S3 \rightarrow S4} E_{S3}}{c p_{z,S3}} - \frac{\Delta z_{S2 \rightarrow S3} E_{S2}}{c p_{z,S2}} - \frac{\Delta z_{S1 \rightarrow S2} E_{S1}}{c p_{z,S1}} \end{aligned} \quad (5.8)$$

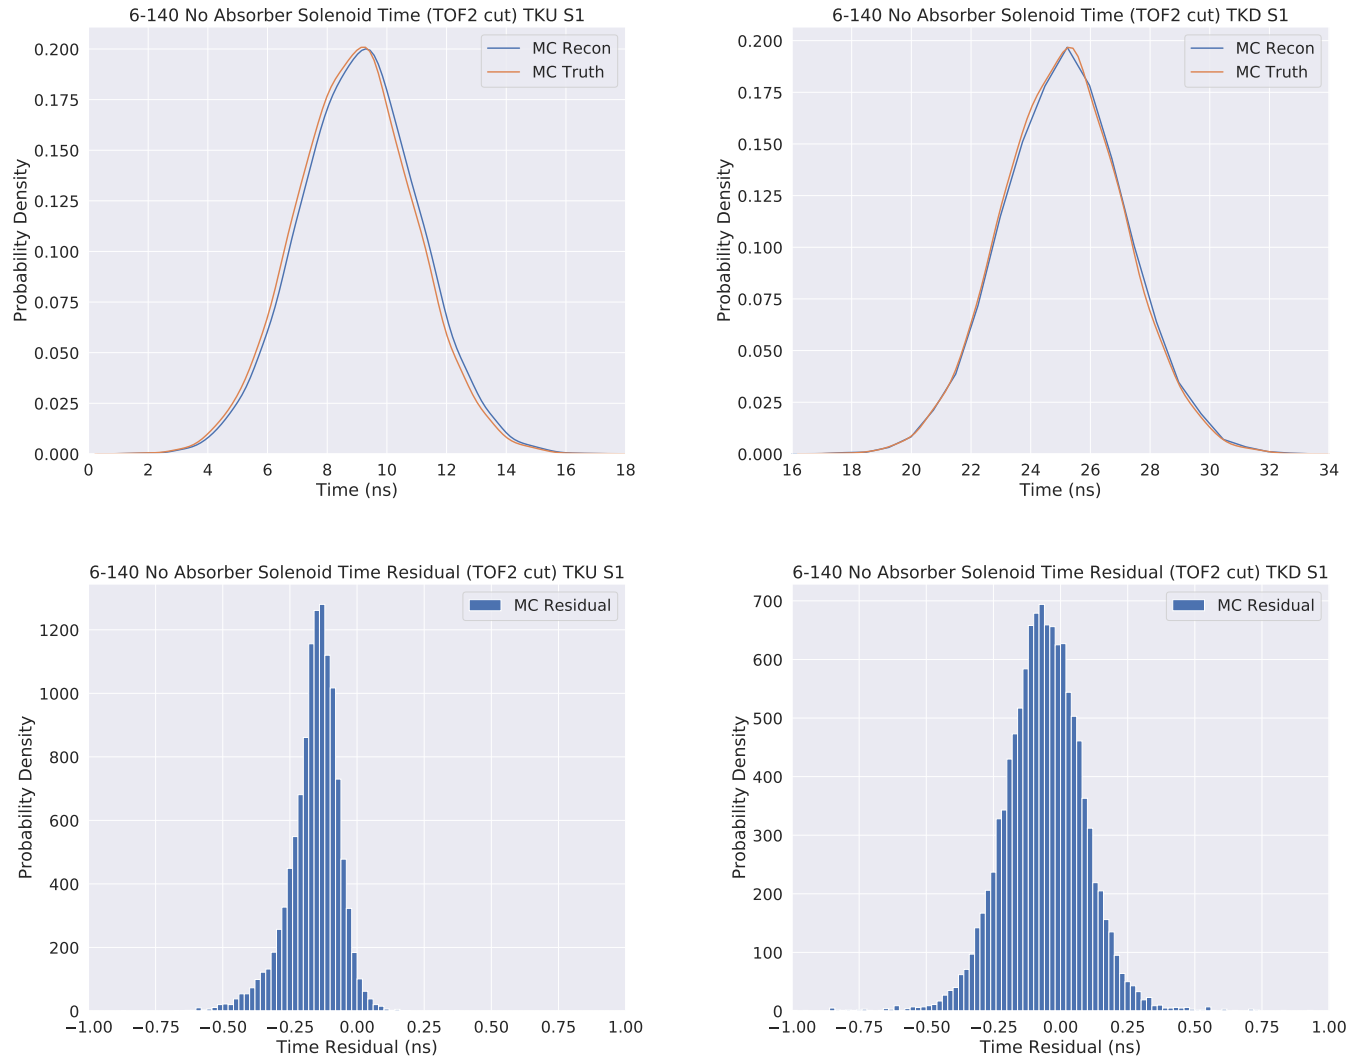
Using the MICE Monte Carlo simulation 247 (the same simulation that was used previously in this chapter), each muon's virtual longitudinal momentum and energy was used to reconstruct the time component of each muon, with Fig. 5.15 showing the distributions of the true and reconstructed time components (top), as well as the residuals (residuals) between the true and reconstructed time component at the upstream (left) and downstream (right) tracker reference planes of all the muons that made it to TOF2. The use of the TOF2 station creates a transmission issue, as not all of the particles at the downstream tracker reference plane will have their time component reconstructed, if those particles were lost between the downstream tracker reference plane and TOF2. This is not an issue for particles at the upstream tracker reference plane, which relies only on TOF1.

Fig. 5.15 shows a small mean offset in the reconstructed time component from the true time component, as the energy lost in the diffuser hasn't been accounted for. Not all muons may pass through the diffuser, however, making such a correction more tricky. A possible correction could involve virtually tracking the muons from the upstream tracker to the diffuser to see if those muons would have passed through the diffuser (this requires an accurate knowledge of the correct magnetic field between the diffuser and the upstream tracker). The error is small in any case in comparison to the error that would be introduced if the reconstructed longitudinal momentum and energy from the MICE reconstruction were used instead of the virtual components.

The spread in the residual time distributions is due to the muon's longitudinal momentum not remaining constant in the MICE cooling channel. The longitudinal and transverse components of the beam transfer between each other as a result of the non-uniformity of the MICE magnetic field, misalignments in the MICE geometry, or when muons enter or exit a solenoid⁸⁷ i.e. the radial magnetic field changes.

How the six dimensional position-momentum phase-space of a particle beam changes as it goes through the MICE cooling channel will be considered in the following chapter in terms of transfer matrices. These transfer matrices will then be used in chapter 7 to fully understand the ionisation cooling result.

Figure 5.15: The true and reconstructed time distributions (top) and residuals (bottom) at the upstream (left) and downstream (right) tracker reference planes for MICE Monte Carlo simulation 247. Only muons that make it to TOF2 are included in the distributions. The distributions show the time it takes the muons to travel from TOF1. The reconstructed time at the upstream tracker reference plane has not accounted for the energy loss the muon beam experiences in the diffuser, resulting in a small offset.



Transfer matrices

For a lens system consisting of a series of magnets (e.g. dipoles, quadrupoles, solenoids), the transfer matrix describing the system can be determined through a combination of the transfer matrices describing each individual magnet. However, aberrations can cause a non-linear transport of the particles through a lens (magnet) system¹⁹³. If the alignment and aberrations of each individual magnet is not known to a great accuracy, then it may be easier to measure the transfer matrix describing the system. For example, if the linear transfer matrix of a solenoid was not known (Eq. 1.32), then it could be determined by knowing that each individual particle is transported via some transport matrix \mathbf{M} , from some initial point $v_u(x, y, p_x, p_y)$ at the upstream end of the solenoid to some final point $v_d(x+dx, y+dy, p_x+dp_x, p_y+dp_y)$ at the downstream end of the solenoid i.e. $v_d = \mathbf{M}v_u$. The components of the transfer matrix \mathbf{M} can then be written in terms of the particle components as⁷⁷

$$\begin{pmatrix} x_d \\ y_d \\ p_{x,d} \\ p_{y,d} \end{pmatrix} = \begin{pmatrix} M_{00} & M_{01} & M_{02} & M_{03} \\ M_{10} & M_{11} & M_{12} & M_{13} \\ M_{20} & M_{21} & M_{22} & M_{23} \\ M_{30} & M_{31} & M_{32} & M_{33} \end{pmatrix} \begin{pmatrix} x_u \\ y_u \\ p_{x,u} \\ p_{y,u} \end{pmatrix} \quad (6.1)$$

Each individual particle component at the end of the solenoid can be written in terms of the initial components at the start of the solenoid. For example, the x component would then be given by

$$x_d = M_{00}x_u + M_{01}y_u + M_{02}p_{x,u} + M_{03}p_{y,u} \quad (6.2)$$

If there are n particles in the distribution, then a set of n linear equations can be found. The optimal values for the M_{00} , M_{01} , M_{02} and M_{03} components are found by

finding the linear least squares solution of the set of n linear equations⁷⁷. This can be repeated for the other beam components of the transport matrix as well.

The same principle can be applied to higher order transfer matrices as well. The approach can also be extended to include the longitudinal components of the beam (z, p_z), or any constant offsets (a_0, a_1, \dots) e.g. if there are any misalignments. The second order transfer matrix between two longitudinal z positions is then given by:

$$\begin{pmatrix} x_d \\ y_d \\ p_{x,d} \\ p_{y,d} \\ p_{z,d} \\ x_d^2 \\ \vdots \\ p_{y,d}p_{z,d} \end{pmatrix} = \begin{pmatrix} M_{00} & M_{01} & M_{02} & M_{03} & M_{04} & M_{05} & \dots & M_{0n} \\ M_{10} & M_{11} & M_{12} & M_{13} & M_{14} & M_{15} & \dots & M_{1n} \\ M_{20} & M_{21} & M_{22} & M_{23} & M_{24} & M_{25} & \dots & M_{2n} \\ M_{30} & M_{31} & M_{32} & M_{33} & M_{34} & M_{35} & \dots & M_{3n} \\ M_{40} & M_{41} & M_{42} & M_{43} & M_{44} & M_{45} & \dots & M_{4n} \\ M_{50} & M_{51} & M_{52} & M_{53} & M_{54} & M_{55} & \dots & M_{5n} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ M_{n0} & M_{n1} & M_{n2} & M_{n3} & M_{n4} & M_{n5} & \dots & M_{nn} \end{pmatrix} \begin{pmatrix} x_u \\ y_u \\ p_{x,u} \\ p_{y,u} \\ p_{z,u} \\ x_u^2 \\ \vdots \\ p_{y,u}p_{z,u} \end{pmatrix} + \begin{pmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ \vdots \\ a_n \end{pmatrix} \quad (6.3)$$

Similar matrices for even higher orders can be derived, though the size of the matrices will quickly grow. Choosing the order of the matrix is not arbitrary, however. If the order of the matrix is too low, then the matrix will under-fit the data. The left plot of Fig. 6.1 shows an example of a data set that has been under-fitted¹⁹⁴. The actual distribution of the data set is given by the blue line, while the circles are a set of measurements of the distribution. Based on the measurements, a simple model is used to describe the system (red line). In this case it is a zeroth-order matrix with the model producing a constant output. If the measurement were to be repeated, the model would still only return the constant value. This means the variance of the model describing the data is zero, however the bias of the model is large¹⁹⁴.

Conversely, if the order of the matrix is too high, then the matrix will over-fit the data. The right plot of Fig. 6.1 shows an example of a data set that has been over-fitted. Again, the actual distribution of the data set is given by the blue line, the circles are a set of measurements of the distribution, while the red line is the output of the model. The model however has too many free parameters as it tries to fit to every single data point, and misses the underlying distribution. The optimal model choice (i.e. which matrix order to use) can be found by finding an optimal balance between the variance and bias of the fitted data sets, e.g. minimising the Mean Squared Error (MSE).

Spurious data points can pull the components of the transfer matrix. For MICE, these would be particles that have been affected by scattering or energy loss which do not describe the transport of those particles through a magnet. As the primary




Figure 6.1: Both plots show the underlying distribution (blue line) and a set of measurements (circles). The left plot has under-fitted the data by using an overly simplistic model to fit to the data. The right plot has over-fitted the data by using an overly complex model to try to fit every data point, and therefore misses the underlying distribution.¹⁹⁴.

concern is to obtain a transfer matrix describing the transport of the particles through a magnet or magnet system, these highly deviating particles can be removed. When a transfer matrix is applied to these particles, they will produce large residuals from the expected distribution. A limit can therefore be placed on the resulting residuals, with the transfer matrix recalculated for the particles that lie within the limit of the residuals. The process can then be repeated until no more particles fall outside the limit of the accepted residuals. The transport matrix improves with every iteration, unless the limit on the residuals is too strict (i.e. too many particles are removed), there are too many free parameters (Over-Fitting), or the number of particles significantly affected by scattering and energy loss is too high e.g. particles passing through several tracker stations or the absorber material.

The effectiveness of the transport matrix can be evaluated by testing the transport matrix on an independent data set and analysing the residuals obtained between the true beam components at the end of the magnet system compared to the beam component values calculated by the transport matrix. The following sections will consider the transport of particles between two stations within a tracker and the transport of particles between the two tracker reference planes.

6.1 Transfer matrix between two stations

First, second, third and fourth order transfer matrices were calculated between station 2 and station 1 of the upstream tracker for the scenario that considered the MICE magnetic field, and when there were no misalignments of the solenoid, tracker or magnetic axes. Muons whose residual beam components exceeded ± 10 mm or

± 10 MeV were continuously excluded from the repeated calculation of the transfer matrix, until no muons remained whose residual beam components exceeded ± 10 mm or ± 10 MeV. This removed approximately 0.1% of the particles. The resulting transfer matrices were then applied to an independent set of data.

The following plots compare the true, reconstructed and predicted (using the transfer matrices) distributions at station 1 of the upstream tracker. Additionally, the residual plots compare the residual between the true and predicted beam components, compared to the residual between the true and reconstructed beam components.

Figures 6.2 and 6.3 show the distributions and the residuals for the x and p_x components. The predicted distributions nearly match the true and reconstructed distributions when a second transfer matrix was used, except that both the reconstructed and predicted p_x components differed slightly from the true distribution at the centre of the distribution. For both the x and p_x components, the residuals showed no noticeable improvements when transfer matrices beyond the second order were used. Additionally, the predicted p_x residuals were slightly smaller than the reconstructed residuals, showing that if the beam components were accurately known at station 2, the transfer matrix could predict the p_x momentum at station 1 slightly better than the MICE reconstruction could.

Fig. 6.4 shows the distributions and the residuals for the p_z component. In this case, the MICE reconstruction performed poorly and did not match the true distribution. The predicted distribution nearly matched the true distribution and only showed a slight deviation from the true distribution near 135 MeV. This was similarly seen in the residual plots, where the residuals did not improve beyond the first order. Again, if the beam components were accurately known at station 2, then the transfer matrix could predict the p_z momentum at station 1 far better than the MICE reconstruction could.

The predicted x and p_x residuals showed a significant improvement when a second order transfer matrix was used. As the distance between the two stations was only 200 mm, it is unlikely for any aberrations to have had a large enough effect to significantly alter the transfer matrix. The improved residuals can however be explained by a non-uniform magnetic field between the two stations of the tracker. Fig. 5.1 showed that the magnetic field had a u-shape in the tracker region. As no improvements in the residuals were obtained by using a transfer matrix beyond the second order, this would suggest that the constant magnetic field term used for the reconstruction of the transverse and longitudinal momenta (Eqs. 1.33 and 1.34) could be replaced by a second order equation describing the U-shape of the magnetic field within the tracker region. This similarly implies that the radial and phase advance constants in those equations should be replaced by second order equations describing the radial and phase advance components.

Figure 6.2: Transfer matrices between the last two stations of the upstream tracker were calculated using a set of muons transported through the MICE cooling channel by the MICE magnetic field. The resulting transfer matrices were then applied to an independent data set, with the plots showing the predicted (Pred), true (True) and MICE reconstructed (Reco) x position distributions (top) and the x position residuals between the true and predicted (Pred) x positions, and between the true and MICE reconstructed (Reco) x positions at the last station of the upstream tracker when first (left) and second (right) order transfer matrices were used.

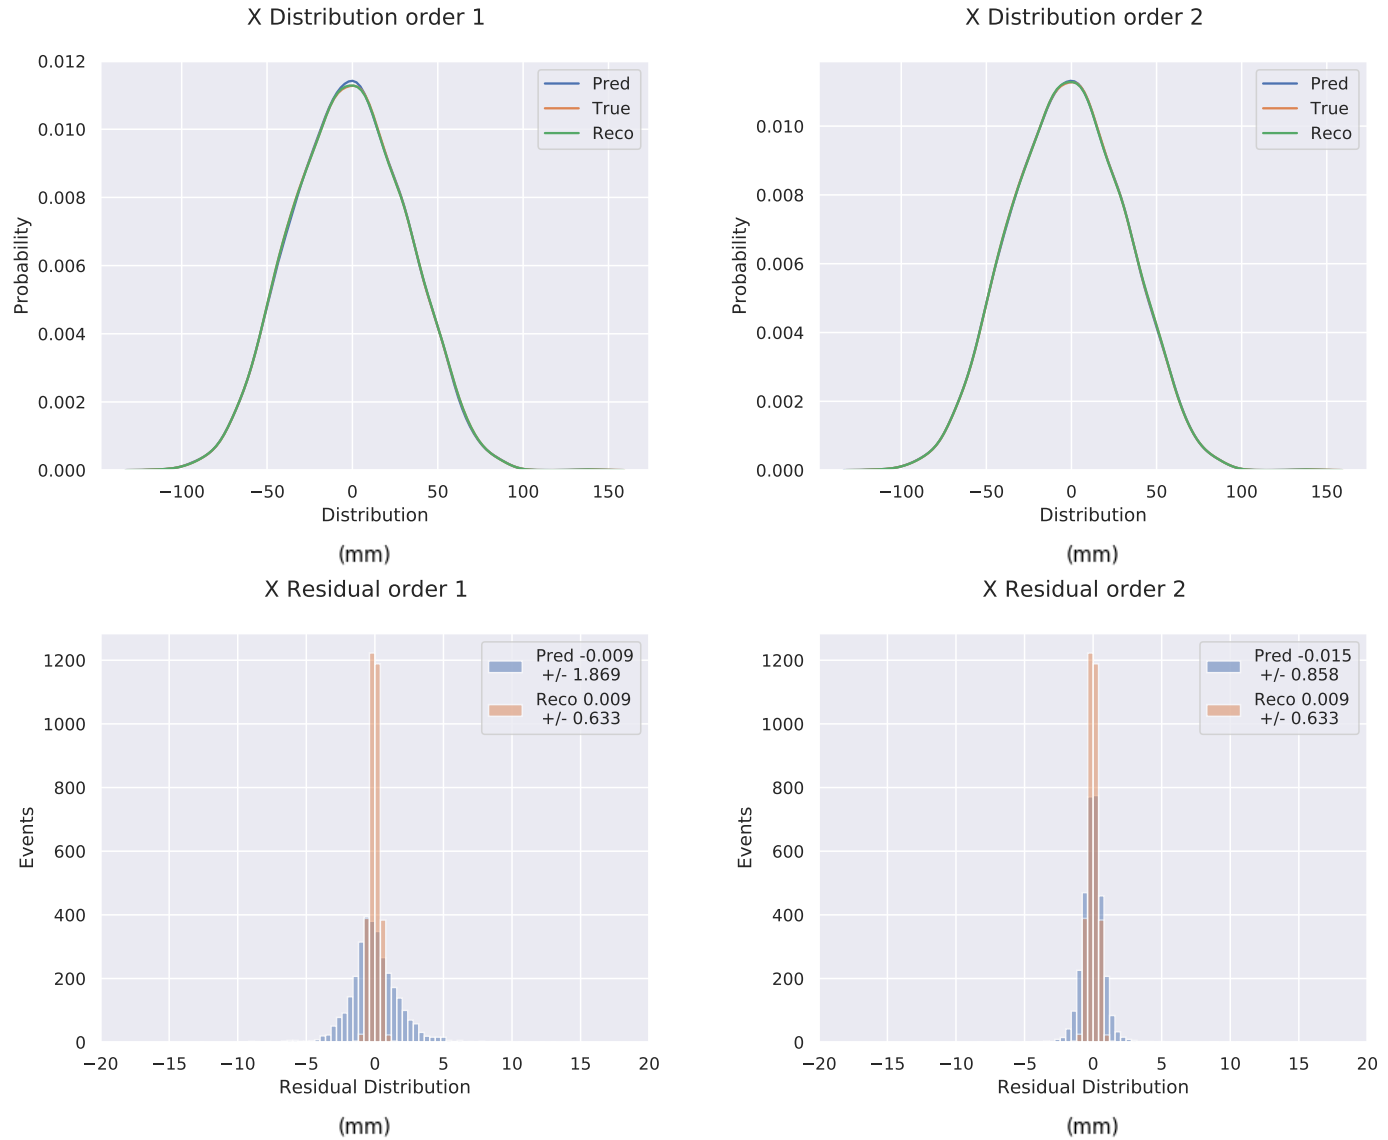


Figure 6.3: Transfer matrices between the last two stations of the upstream tracker were calculated using a set of muons transported through the MICE cooling channel by the MICE magnetic field. The resulting transfer matrices were then applied to an independent data set, with the plots showing the predicted (Pred), true (True) and MICE reconstructed (Reco) p_x momentum distributions (top) and the p_x momentum residuals between the true and predicted (Pred) p_x momenta, and between the true and MICE reconstructed (Reco) p_x momenta at the last station of the upstream tracker when first (left) and second (right) order transfer matrices were used.

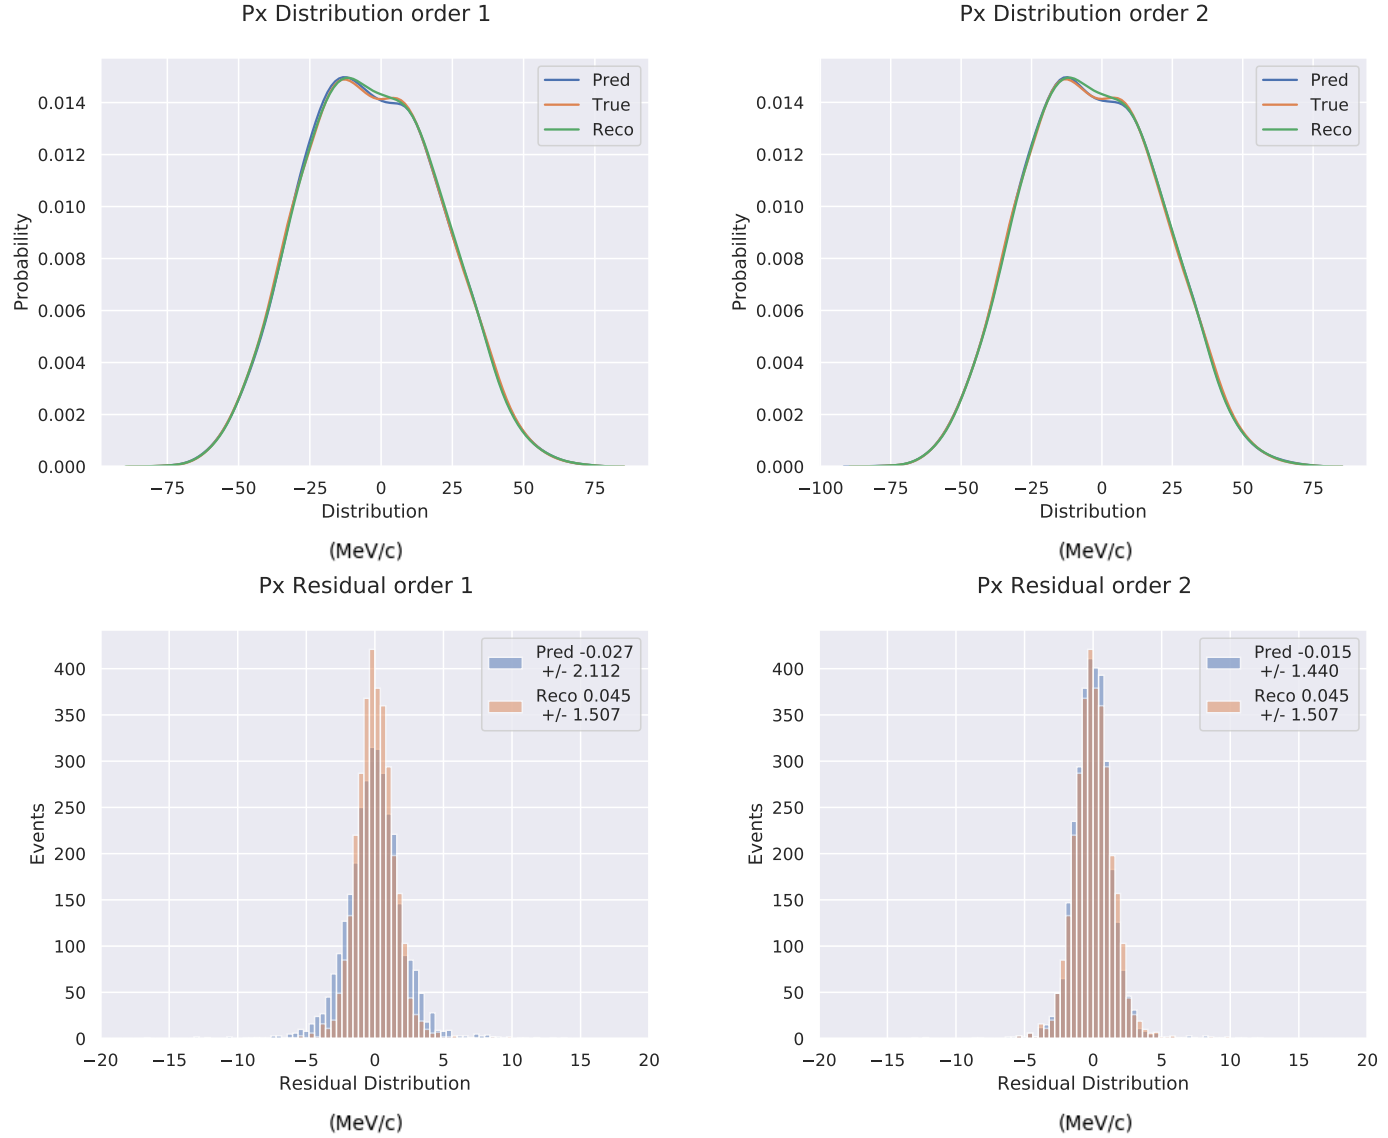
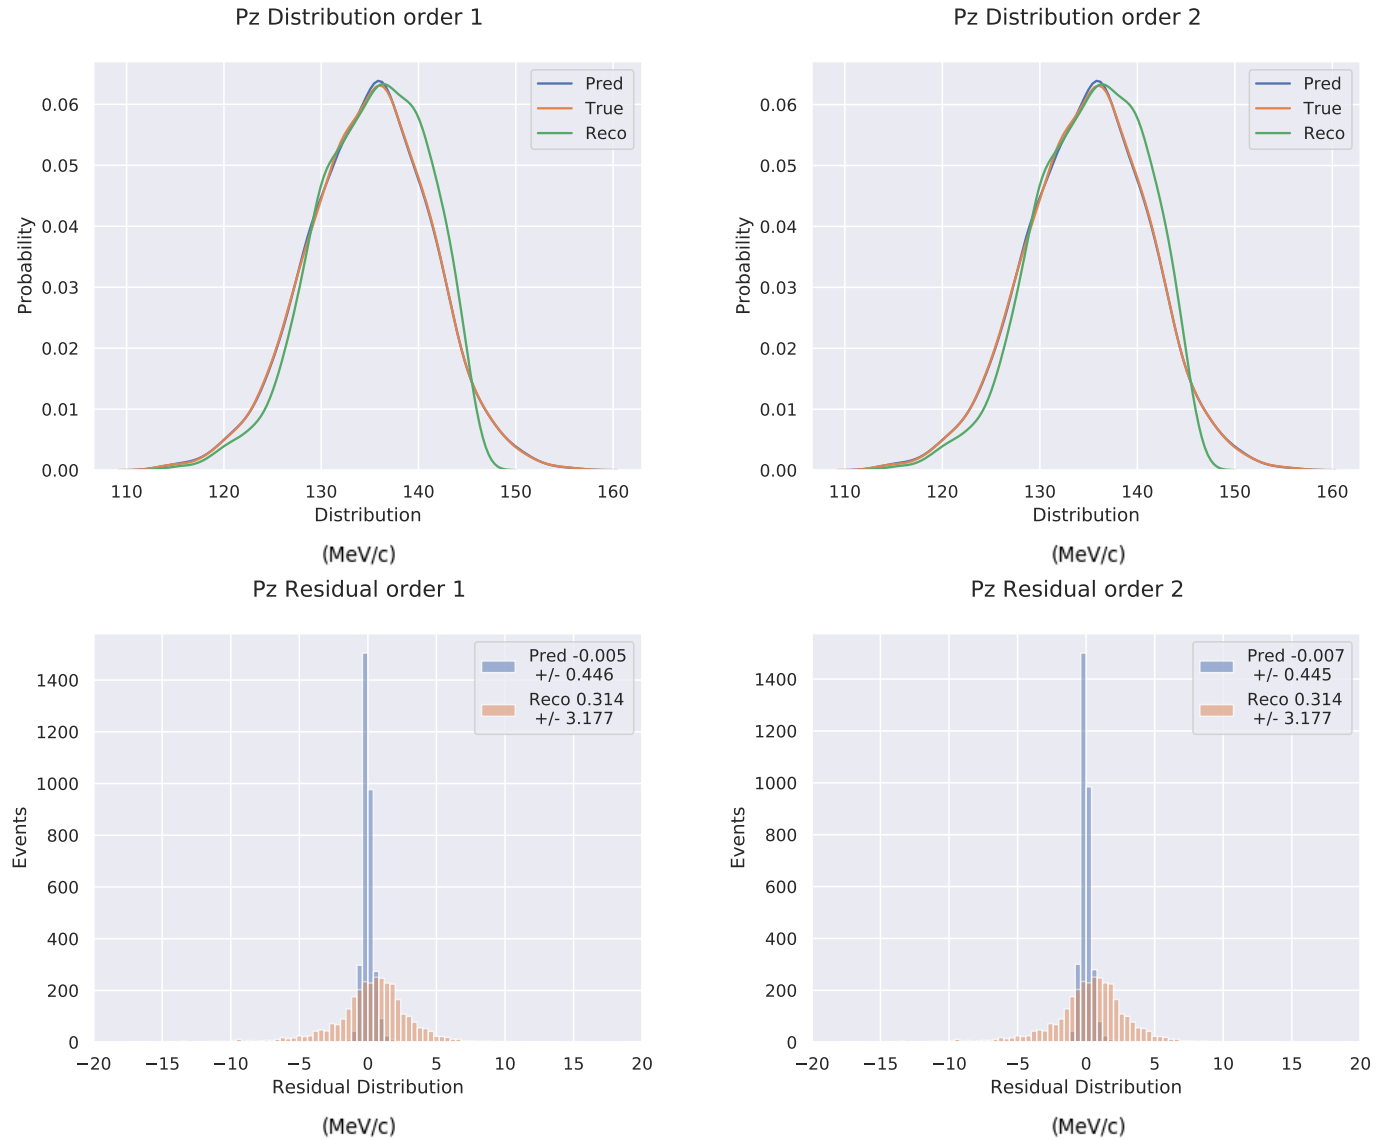


Figure 6.4: Transfer matrices between the last two stations of the upstream tracker were calculated using a set of muons transported through the MICE cooling channel by the MICE magnetic field. The resulting transfer matrices were then applied to an independent data set, with the plots showing the predicted (Pred), true (True) and MICE reconstructed (Reco) p_z momentum distributions (top) and the p_z momentum residuals between the true and predicted (Pred) p_z momenta, and between the true and MICE reconstructed (Reco) p_z momenta at the last station of the upstream tracker when first (left) and second (right) order transfer matrices were used.



6.2 Transfer matrix between two trackers

Transfer matrices up to the fourth order were calculated between the reference planes of the upstream and downstream trackers for the scenario that considered the MICE magnetic field and when there were no misalignments of the solenoid, tracker or magnetic axes. Figures 6.5, 6.6, 6.7 and 6.8 show the distributions and the residuals for the x and p_x components. The predicted distributions show some slight deviations from the true distribution, but the predicted distribution continuously improves as the matrix order is increased. This is similarly seen for the residual plots.

For the x component, the transfer matrix approach performs significantly worse than the MICE reconstruction. However, for this scenario, the transfer matrix was applied over a longitudinal distance of approximately 3780 mm. The transfer matrix also traverses the non-uniform magnetic fields generated by the match coils of the spectrometer solenoids and the focus coils. When a fourth order transfer matrix was used for the p_x component, the predicted residuals were only slightly worse than the reconstructed residuals. The RMS is however significantly larger for the predicted residuals as for a few muons, the predicted momentum differs significantly from the true momentum. This could be due to scattering or a significant amount of energy loss. An energy loss cut could be applied to obtain more representative means and RMSs. It should be noted though that not all virtual muons were reconstructed, i.e. muons that would alter the reconstructed mean or RMS.

By design, the transfer matrix approach tries to obtain residuals whose mean is equal to zero. The benefit of such an approach becomes apparent for the p_z residuals (Fig. 6.10), where the bias of the MICE reconstruction can be significant. The p_z distribution plots (Fig. 6.9) show that the reconstructed momentum differs significantly from the true momentum, while the predicted distribution matches the true distribution more closely as the matrix order increases. This is not too surprising. The predicted p_z primarily depends on the initial p_z at the upstream tracker reference plane, while the reconstructed p_z depends on the non-helical trajectory of the muon through the downstream tracker.

The transfer matrix approach has its advantages and disadvantages. The linear least squares approach means it is inherently unbiased compared to the MICE reconstruction, though it performs poorly for particles that undergo large energy losses or scattering. The transverse positions can be poorly predicted by the transfer matrix. However, as the transverse positions are measured directly, these can be used instead. The transverse momenta are predicted slightly worse than the reconstructed momenta, though the reconstruction fails to reconstruct all particles. The longitudinal momentum is predicted to a much higher accuracy than the reconstructed momentum, though it does require knowledge of the initial longitudinal momentum.

Figure 6.5: Transfer matrices between the tracker reference planes were calculated using a set of muons transported through the MICE cooling channel by the MICE magnetic field. The solenoid, tracker and magnetic axes were not misaligned in the Monte Carlo simulation. The resulting transfer matrices were then applied to an independent data set, with the plots showing the predicted (Pred), true (True) and MICE reconstructed (Reco) x position distributions at the reference plane of the downstream tracker in increasing order of the applied transfer matrix.

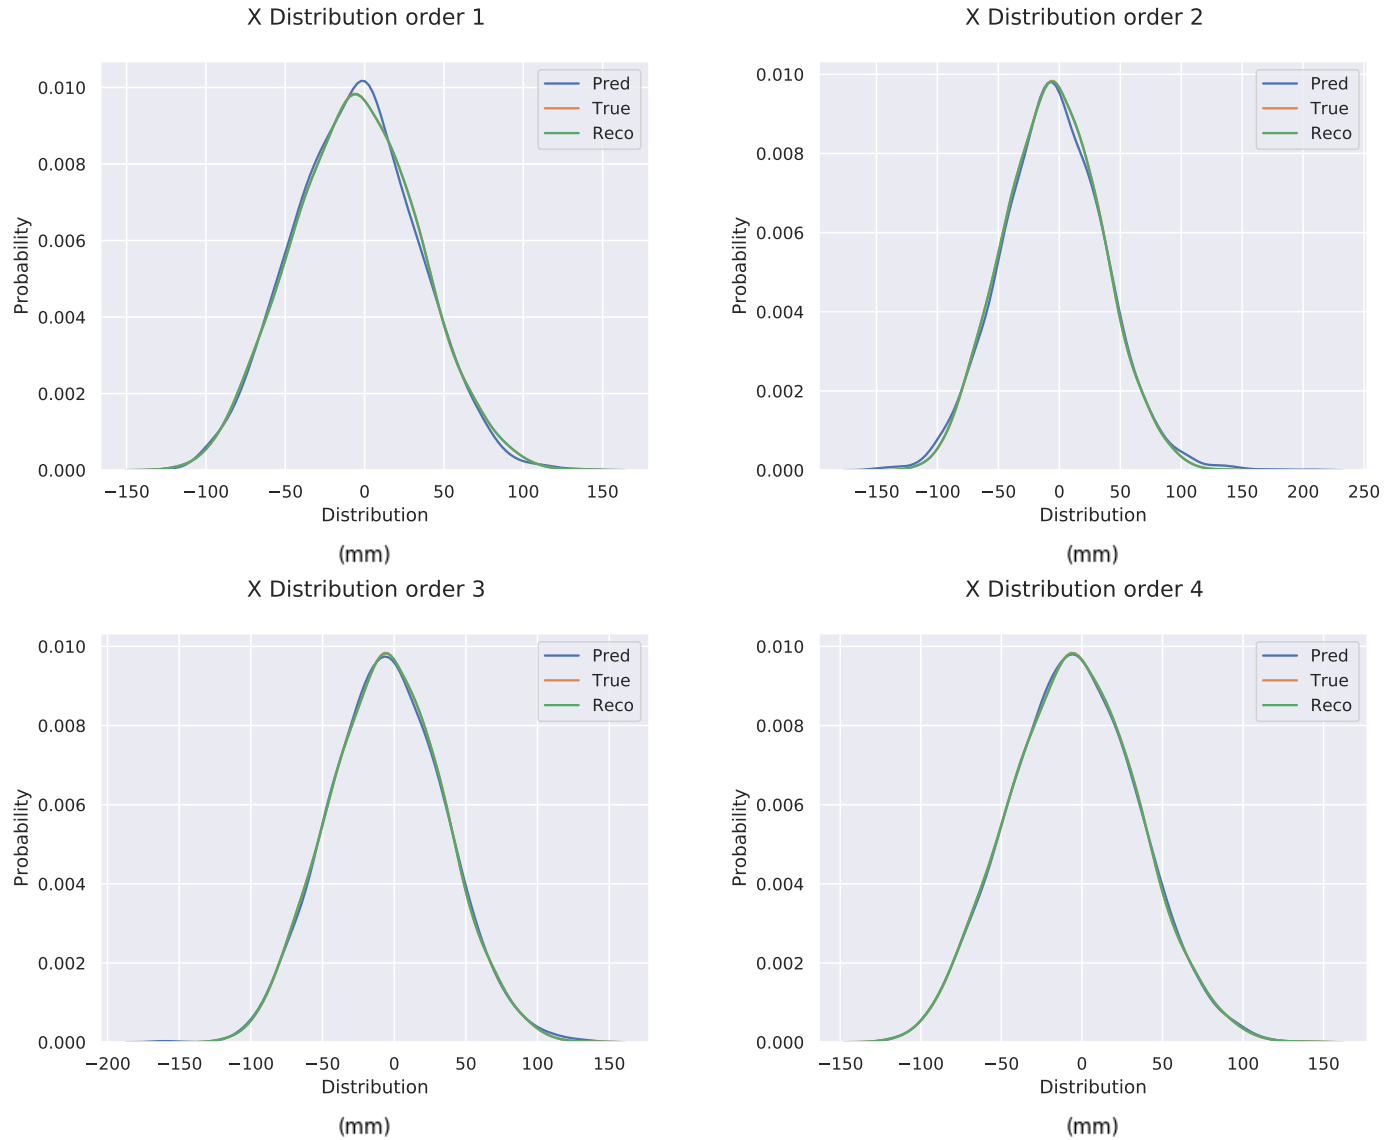


Figure 6.6: Using the same transfer matrices applied to the same independent data set as for Fig. 6.5, the plots show the x position residuals between the true and predicted (Pred) x positions, and between the true and MICE reconstructed (Reco) x positions at the reference plane of the downstream tracker in increasing order of the applied transfer matrix. The legend shows the corresponding mean \pm the RMS for each residual distribution.

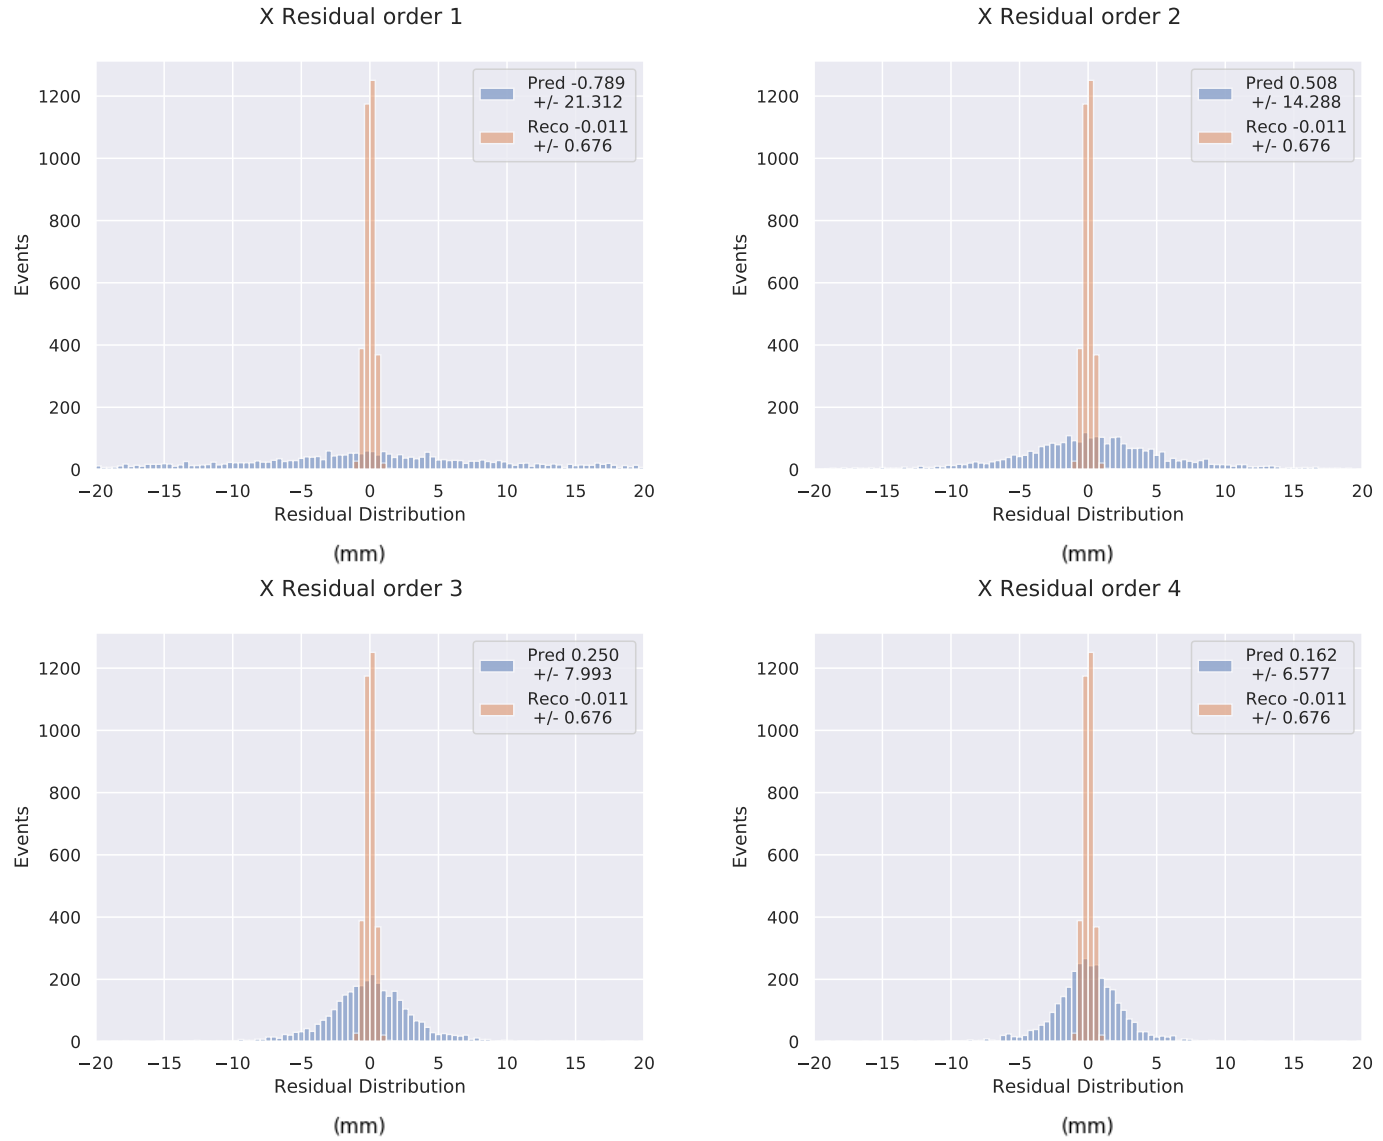


Figure 6.7: Transfer matrices between the tracker reference planes were calculated using a set of muons transported through the MICE cooling channel by the MICE magnetic field. The solenoid, tracker and magnetic axes were not misaligned in the Monte Carlo simulation. The resulting transfer matrices were then applied to an independent data set, with the plots showing the predicted (Pred), true (True) and MICE reconstructed (Reco) p_x momentum distributions at the reference plane of the downstream tracker in increasing order of the applied transfer matrix.

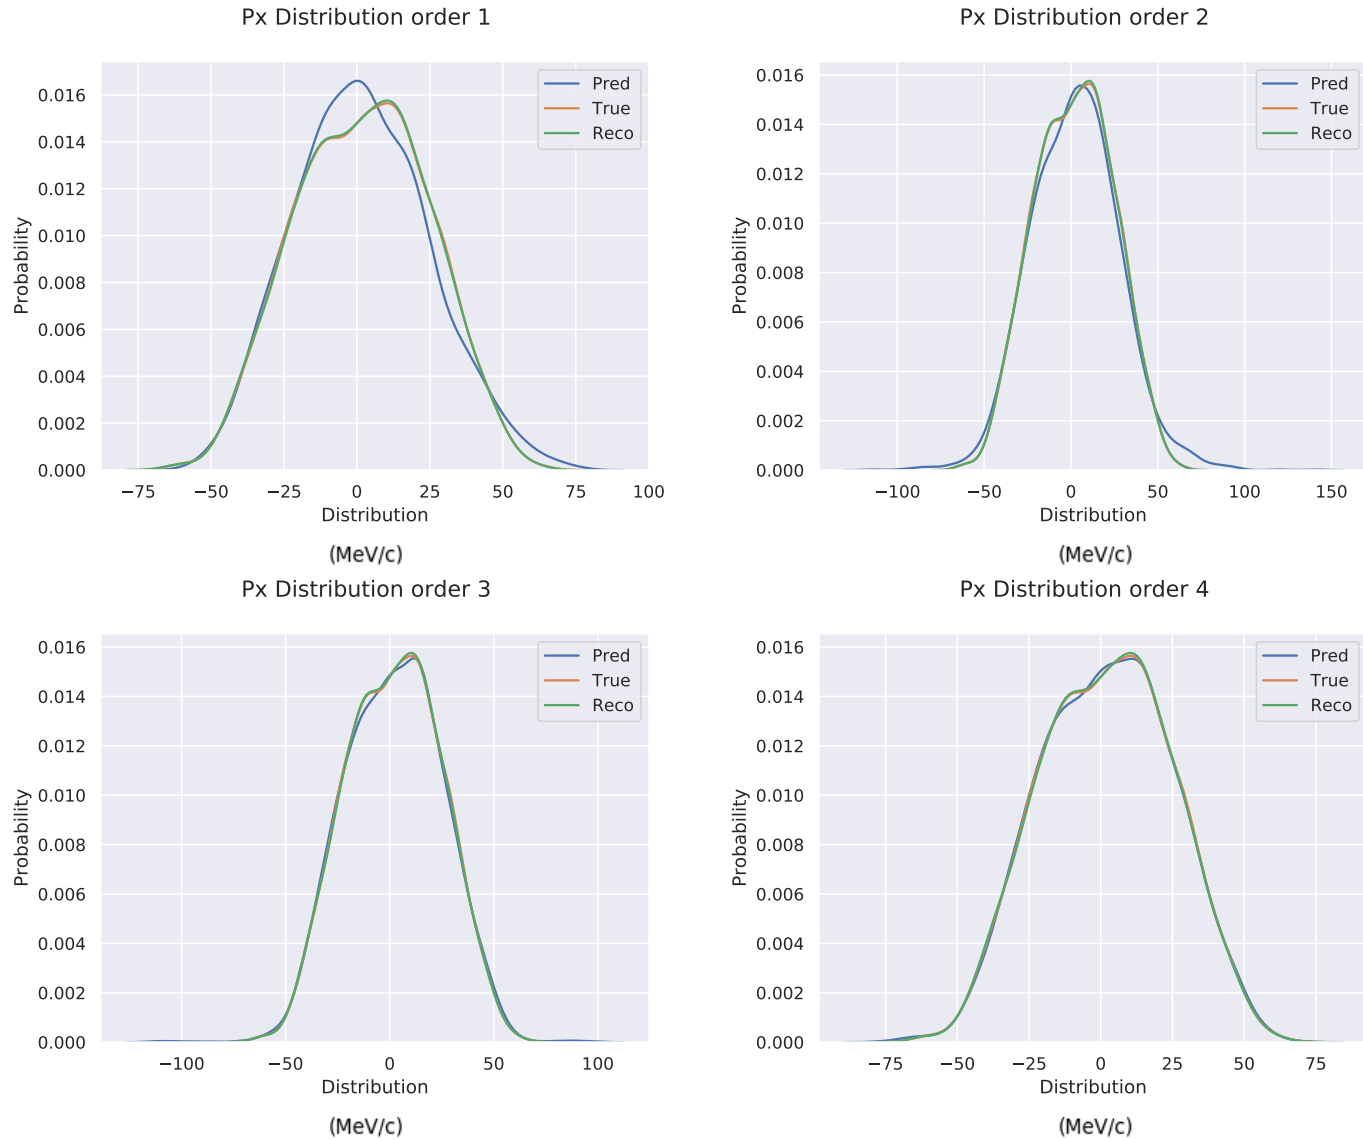


Figure 6.8: Using the same transfer matrices applied to the same independent data set as for Fig. 6.7, the plots show the p_x momentum residuals between the true and predicted (Pred) p_x momentum, and between the true and MICE reconstructed (Reco) p_x momentum at the reference plane of the downstream tracker in increasing order of the applied transfer matrix. The legend shows the corresponding mean \pm the RMS for each residual distribution.

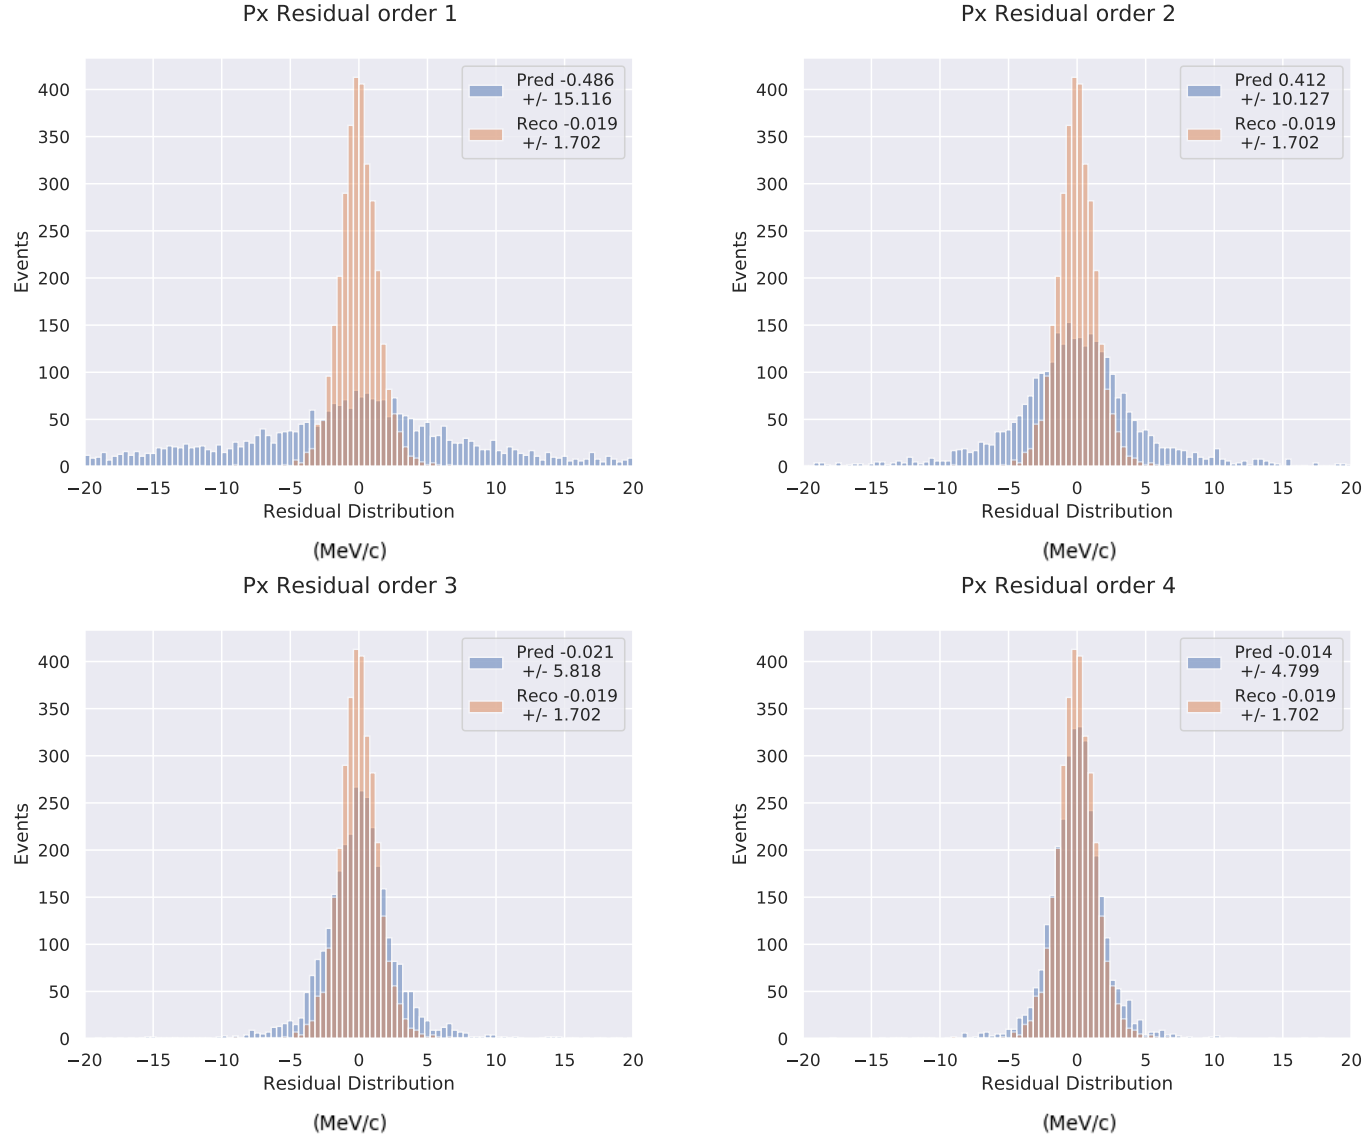


Figure 6.9: Transfer matrices between the tracker reference planes were calculated using a set of muons transported through the MICE cooling channel by the MICE magnetic field. The solenoid, tracker and magnetic axes were not misaligned in the Monte Carlo simulation. The resulting transfer matrices were then applied to an independent data set, with the plots showing the predicted (Pred), true (True) and MICE reconstructed (Reco) p_z momentum distributions at the reference plane of the downstream tracker in increasing order of the applied transfer matrix.

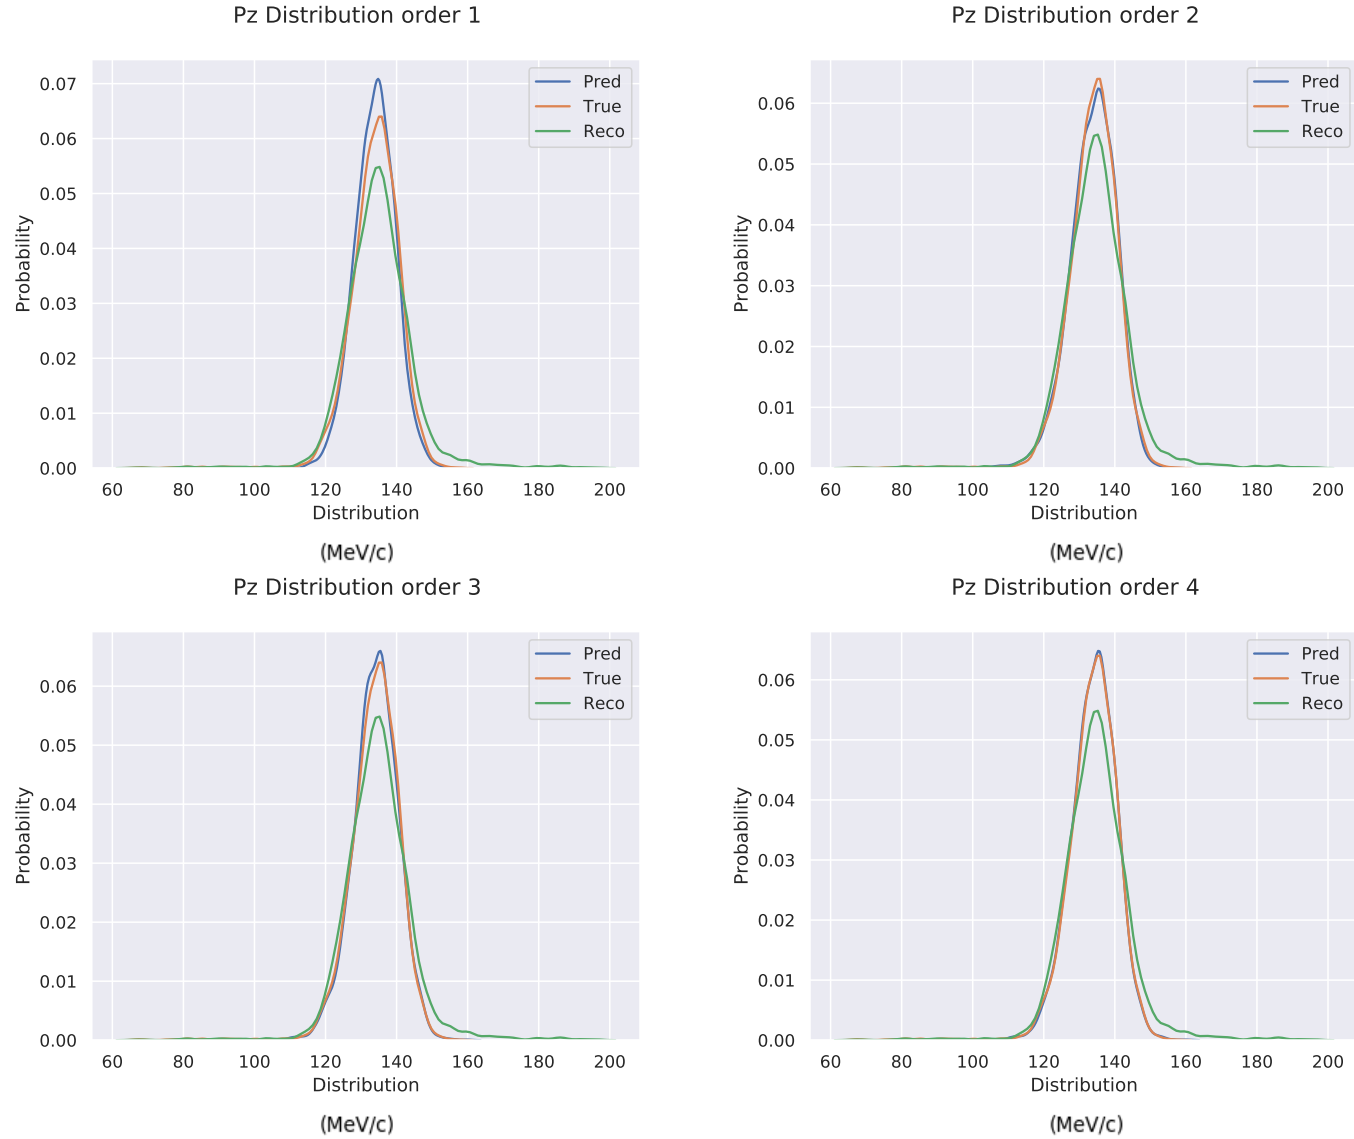
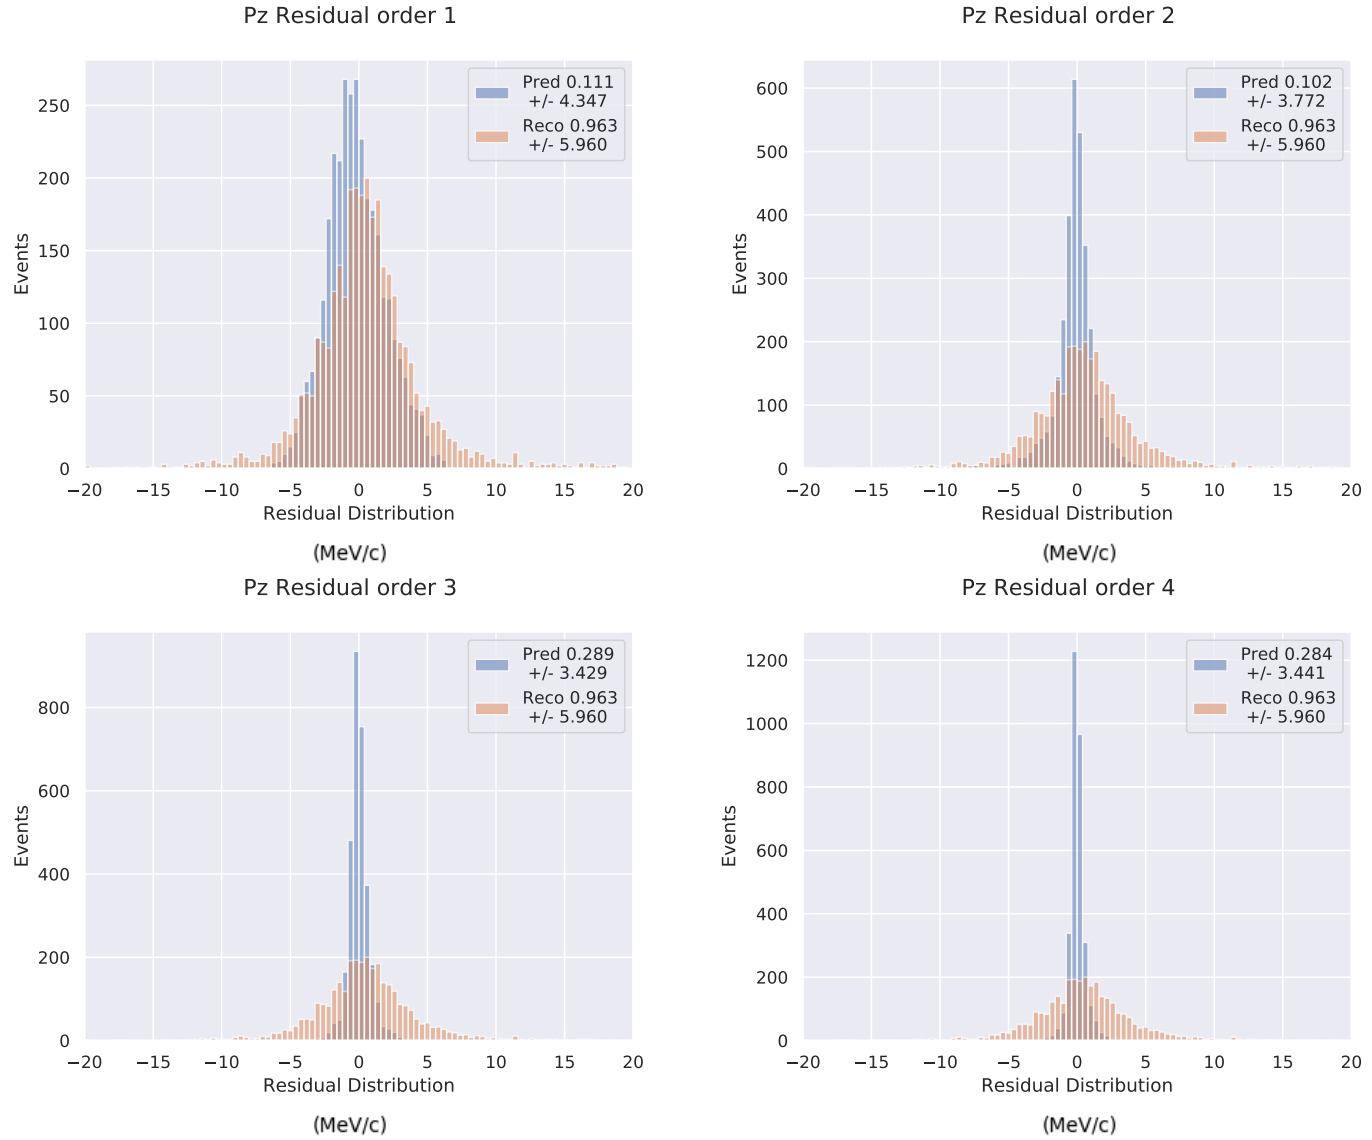


Figure 6.10: Using the same transfer matrices applied to the same independent data set as for Fig. 6.9, the plots show the p_z momentum residuals between the true and predicted (Pred) p_z momentum, and between the true and MICE reconstructed (Reco) p_z momentum at the reference plane of the downstream tracker in increasing order of the applied transfer matrix. The legend shows the corresponding mean \pm the RMS for each residual distribution.



Emittance, amplitude and density

The complete state of a particle can be described by its position and momentum coordinates¹⁹⁵. For a three-dimensional Euclidean space (formed by the x , y and z axes), this forms a six-dimensional position and momentum phase space. When all of the particles in a particle beam are considered simultaneously, the distribution of the particles in the beam can be described by the phase space density $\rho(x, y, z, p_x, p_y, p_z)$. The number of particles in a given phase-space volume can then be found by integrating the phase-space density distribution across the phase-space volume^{195,196}, such that

$$N = \int \rho(x, y, z, p_x, p_y, p_z) dx dy dz dp_x dp_y dp_z = \int \rho dV \quad (7.1)$$

Joseph Liouville^{74,195–198} showed that for conservative forces, the phase-space density remains constant, i.e. $d\rho/dt = 0$. This similarly means that the phase-space volume is preserved. This even applies when the particle beam passes through a magnetic lens system such as the MICE cooling channel, as the magnets act as conservative forces. Non-conservative forces alter the energy of the particles in the beam and cause the phase-space volume of the beam to change. Examples of non-conservative forces acting on the beam include multiple coulomb scattering and energy loss when the particles of a beam pass through an absorber material (such as lithium hydride or liquid hydrogen), or when the particles of a beam pass through an RF cavity and where the particles of the beam either gain or lose energy as they are accelerated or decelerated.

While the phase-space volume of a particle beam remains constant when subjected to a conservative force or no force at all, the shape of the phase-space volume can change as seen in Fig. 7.1. The two cases are equivalent to a particle beam travelling through a drift space (left), or being subjected to a uniform magnetic field (right).

For both cases, the shape of the phase-space volume changes predictably. If the magnetic field is non-uniform, however, the change in the shape of the phase-space volume can become more difficult to predict. This is the case for MICE, where there is a significant uncertainty in the shape and magnitude of the magnetic field in the MICE cooling channel. This poses a problem for two of the three different types of cooling measurement MICE uses, and will be discussed in the next section.




Figure 7.1: The phase-space volume of an infinitesimal hypercube¹⁹⁶ will remain constant when it is either subjected to no force (left) or when it is subjected to a constant force (right). The shape of the infinitesimal hypercube can change, however.

7.1 Three types of cooling measurement

Extensive use is made of Liouville's theorem for the demonstration of ionization cooling (and emittance exchange) at MICE. The demonstration of ionization cooling at MICE can be performed by showing a change in the beam emittance, amplitude (single particle emittance) or phase-space density.

The following subsections will consider each type of cooling measurement in more detail and will rely heavily on the descriptions by Drielsma⁷⁴ and Wiedemann¹⁹⁹. The subsections will also consider the assumptions used for each type of cooling measurement, as well as their limitations.

7.1.1 Beam emittance

In the Gaussian approximation⁷⁴, the 2D trace space density can be given in terms of the covariance matrix $\Sigma_{\mathbf{q}}$, such that

$$\rho(\mathbf{q}) = \frac{1}{(2\pi)^2 |\Sigma_{\mathbf{q}}|^{\frac{1}{2}}} \exp\left\{-\frac{1}{2} \mathbf{q}^T \Sigma_{\mathbf{q}}^{-1} \mathbf{q}\right\} \quad (7.2)$$

where the trace space vector is given by $\mathbf{q} = (q, q')$ and the covariance matrix is given by

$$\Sigma_q = \begin{pmatrix} \sigma_{qq} & \sigma_{qq'} \\ \sigma_{qq'} & \sigma_{q'q'} \end{pmatrix} \quad (7.3)$$

The geometric emittance is then defined in terms of the covariance matrix as

$$\epsilon = |\Sigma_q|^{\frac{1}{2}} \quad (7.4)$$

For a six-dimensional phase-space, where the three 2D phase-spaces can be completely decoupled, the 6D emittance is given by

$$\epsilon_{6D} = \sqrt{\epsilon_x \epsilon_y \epsilon_z} = |\Sigma_{6D}|^{\frac{1}{6}} = \begin{pmatrix} \sigma_{xx} & \sigma_{xx'} & 0 & 0 & 0 & 0 \\ \sigma_{xx'} & \sigma_{x'x'} & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_{yy} & \sigma_{yy'} & 0 & 0 \\ 0 & 0 & \sigma_{yy'} & \sigma_{y'y'} & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_{zz} & \sigma_{zz'} \\ 0 & 0 & 0 & 0 & \sigma_{zz'} & \sigma_{z'z'} \end{pmatrix}^{\frac{1}{6}} \quad (7.5)$$

Such a description can be particularly useful in linear optics, where the motion of the particles through a transport line can be described by a series of transfer matrices. The transfer matrices describing the magnets in the transport line are typically chosen so that the transport is symplectic, that is $\mathbf{M}^T \mathbf{M} = \mathbf{I}$, where \mathbf{I} is the identity matrix and \mathbf{M} is a transfer matrix. In linear optics, the transfer matrices only perform linear translations and rotations on the covariance matrix,

$$\Sigma' = \mathbf{M} \Sigma \mathbf{M}^T \quad (7.6)$$

For transport lines where the linear transfer matrices satisfy $\mathbf{M}^T \mathbf{M} = \mathbf{I}$, the emittance of the beam remains conserved. If the individual 2D phase-spaces are completely decoupled from one another, then the individual 2D emittances also remain conserved.

For a solenoid, the two transverse phase-spaces are coupled to each other due to the angular momentum of the beam⁷⁴. MICE however assumes that the transverse phase-space can still be completely decoupled from the longitudinal phase-space⁷⁸, with the linear transverse solenoid transfer matrix described by Eq. 1.32. This transfer matrix has a dependence on the overall momentum of the particle. For most solenoid transport lines, little or no energy loss occurs in the transport line. If the spread in longitudinal components of the beam is kept small, and therefore the spread in the overall momentum, then each individual particle in the beam will experience approximately the same linear transverse solenoid transfer matrix. If the same transfer matrix acts on all of the particles in the beam, then the transverse emittance of the beam will remain conserved. Fig 5.1 already showed that the magnetic field

in the tracker regions of the MICE cooling channel was not uniform, however. This means a transfer between the transverse and longitudinal beam components can occur.

The small momentum spread in the MICE beam is only achieved by selecting particles in a narrow momentum range at the upstream tracker reference plane, e.g. 140 ± 5 MeV. The beam is therefore non-Gaussian. When there is no absorber material present in the MICE cooling channel, MICE assumes that each particle will experience approximately the same transfer matrix, and therefore the emittance of the beam will remain approximately conserved.

When an absorber material is present in the MICE cooling channel, however, the particles will lose energy at different rates. The particles will therefore have a greater variation in the phase advances they make between the absorber material and the downstream reference plane. This is due to the greater variation in the transfer matrices experienced by each particle. This results in some particles travelling further along their helical trajectory in the same longitudinal distance. The variations become even larger when a non-uniform absorber material is placed into the MICE cooling channel, such as the polyethylene wedge. As the spread of the transverse beam components is changed by the differences in the longitudinal momentum, this causes a continuous change in the transverse emittance of the beam as a function of longitudinal distance (see section 7.3).

7.1.2 Amplitude - single particle emittance

The amplitude of a particle is known as the single particle emittance. As the name suggests, it uses the emittance for a distribution of particles (i.e. the particle beam) to calculate an amplitude value for each individual particle. The amplitude of a particle is given by the emittance of the beam multiplied by the Mahalanobis distance ($\mathbf{q}^T \Sigma^{-1} \mathbf{q}$) between the particle and the beam centroid⁷⁴, and is given by

$$A = \epsilon^* \mathbf{q}^T \Sigma^{-1} \mathbf{q} \quad (7.7)$$

where $\mathbf{q} = q - \langle q \rangle$, q is the phase-space vector e.g. $q = (x, p_x, y, p_y, z, p_z)$ and Σ is the covariance matrix. In d dimensions, the Mahalanobis distance follows a χ^2 distribution with d degrees of freedom. This means the amplitudes of the particles in the beam will be described by the following distribution

$$A = \epsilon_d^* \chi_d^2 \quad (7.8)$$

For a linear optics system, where there is only a small spread in the longitudinal components of the beam, the emittance of the beam is approximately conserved. This

means the differences in the amplitudes of the particles are solely due to differences in the Mahalanobis distance $\mathbf{q}^T \boldsymbol{\Sigma}^{-1} \mathbf{q}$. For a linear optics system where there are only linear translations and rotations, the Mahalanobis distance remains conserved. This means that the amplitude of a particle in a beam (with a small spread in the longitudinal components of the beam) is similarly preserved in a linear optics system.

However, as the amplitude of the particle depends on the emittance of the beam, this means that for the MICE cooling channel, the amplitude of a particle will face the exact same challenges the emittance of the beam did in the previous subsection. These included the non-uniformity of the solenoid field within the tracker regions and the longitudinal spread of the beam when the particles in the beam encountered an absorber material. As the spread of the beam has now changed (i.e. non-linear translations and rotations), so will the calculated Mahalanobis distance.

7.1.3 Phase-space density

The phase-space density considered by Liouville can be calculated directly. The two methods MICE used to calculate the phase-space density were the k-nearest neighbour (KNN) approach investigated by François Drielsma⁷⁴ and the kernel density estimation (KDE) approach investigated by Tanaz Mohayai²⁰⁰. For the parameters MICE considered, they should produce the exact same results.

Kernel Density Estimation

Phase-space density estimation is a non-parametric technique used to estimate the probability that a particle will be realized at a particular phase space density. As a non-parametric technique, it makes no assumptions about the underlying probability distribution. In Kernel Density Estimation, each individual point in the distribution is represented by a kernel (e.g a Gaussian, uniform, epanechnikov²⁰¹, etc. distribution). The underlying probability distribution is then obtained by summing over all individual kernels. Fig. 7.2 shows an example of how KDE was used to estimate the underlying probability distribution in one dimension.

The kernel density estimate for the distribution is obtained through the formula

$$\rho(\mathbf{q}) = \frac{1}{n} \sum_{i=1}^n K_{\mathbf{H}}(\mathbf{q}) = \frac{1}{nh^d} \sum_{i=1}^n K\left(\frac{\mathbf{q}}{h}\right) \quad (7.9)$$

where K is the kernel choice, h is the bandwidth of the kernel, d is the number of dimensions and n is the sample size. For higher dimensions, the covariance matrix is used to equally weight the bandwidth choice for each dimension. The kernel choice only has a minor effect on the probability distribution. For most situations a Gaussian kernel is chosen as this has the added benefit that the probability distribution remains




Figure 7.2: Example of how kernel density estimation can be used to estimate the underlying probability distribution²⁰⁰.

differentiable. The Gaussian kernel is given by

$$K(q) = (2\pi)^{-d/2} \exp\left\{-\frac{1}{2}\mathbf{q}^T \mathbf{q}\right\} \quad (7.10)$$

This results in Eq. 7.9 simplifying to

$$\rho(q) = \frac{\sum_{i=1}^n \exp\left\{-\frac{1}{2}\mathbf{q}^T \Sigma^{-1} \mathbf{q}\right\}}{n(2\pi)^{d/2} |\Sigma|^{1/2}} \quad (7.11)$$

where Σ is the covariance matrix of the distribution and $|\Sigma|$ is the determinant of the covariance matrix.

K-Nearest Neighbour

For a distribution of sample size n , the k^{th} nearest neighbour is found for each point in the distribution, with k typically chosen to be equal to \sqrt{n} , where n is the sample size. For each point in the distribution, a line is drawn to its k -nearest neighbour (Fig. 7.3), with the line acting as the radius for a unit d -ball, where d is the dimension of the phase-space. The volume of the unit d -ball is inversely proportional to the phase space density estimate for that particle.

The KNN density estimate is obtained via the formula

$$\rho(q) = \frac{k}{n\kappa_d R_k^d} = \frac{k\Gamma(\frac{d}{2} + 1)}{n\pi^{\frac{d}{2}} R_k^d} \quad (7.12)$$

where κ_d is the volume of a unit d -ball, R_k^d is the Euclidean radius of the unit d -ball, n is the sample size and $\Gamma(\frac{d}{2} + 1)$ is Euler's gamma function.




Figure 7.3: Example of the k -nearest neighbour technique applied to two different points where $k = 5$. In the denser region of the distribution, neighbouring points are closer together, resulting in smaller radii circles. The area of the circles is inversely proportional to the density estimate⁷⁴.

Phase space density distribution

Fig. 7.4 shows the expected probability distributions for a particle sampled from a Gaussian distribution to be realized at a particular phase space density in four different dimensions. The distributions can be better understood when one considers the volumes in each dimension. The 1-dimensional volume is given by a line, the 2-dimensional volume is given by a circle, and the three-dimensional volume is given by a sphere. For a 1-d Gaussian distribution along a line, most points will lie close to the centre of the distribution, i.e. the densest part of the distribution. This results in the rightmost peak of the distribution in Fig. 7.4.

For the two-dimensional case, each particle is formed of two coordinates. This means that the particle can have one of its coordinates at a very dense part of the distribution for that dimension, while the other coordinate is at a less dense part of the distribution for that dimension. As the overall density is formed by both coordinates of the distribution, this means that for a Gaussian distribution, all phase space densities are equally likely.

For the three-dimensional case, each particle is formed of three coordinates. Only one of the three coordinates of the particle needs to be in a less dense part of that coordinate's distribution to drag the overall particle to a less dense part of the sphere. This means the number of opportunities for one of the coordinates of the particle to be in a less dense part of the distribution is far greater, which is reflected by the leftmost peaks of the probability distribution in Fig. 7.4. The effect becomes even more pronounced for higher dimensions.




Figure 7.4: The expected probability distributions for a particle sampled from a Gaussian distribution to be realized at a particular phase space density in four dimension⁷⁴. The density has been normalized to the maximum density of the distribution.

7.2 Missing data

For a given experiment, the results of the experiment can be described by a probability distribution. Sometimes, not all the results of an experiment can be recorded, however, with some of the results going missing. This can bias the measured probability distribution. Missing data can be classified in three different ways^{202–205}: Missing Completely At Random (MCAR), Missing At Random (MAR) and Missing Not At Random (MNAR).

7.2.1 Missing Completely At Random

Data that is Missing Completely At Random (MCAR) doesn't bias the measured probability distribution. For example, if a coin flip experiment was performed 200 times, but the first 100 results of the experiment then went missing. The results of the experiment would not become biased, as the experimental results were recorded randomly and only depended on the coin flip. The missing data is simply a random subsample of the complete data set. The error on the reported experimental results will increase, however, as there are now fewer experimental results as a result of the missing data.

7.2.2 Missing At Random

Data that is Missing At Random (MAR) can bias the experimental results, though there are correction procedures that can be used to account for the missing data such as imputation techniques, maximum likelihood estimation, interpolation or

partial deletion²⁰². For example, in the coin flip experiment, the data would be MAR if the heads and tails results of the experiment were recorded in different parts of a notebook and one of the pages then went missing. The missing data could be estimated if it was known how many pages went missing and how many results were recorded on average on a page. The missing data has been imputed. If the number of experimental results recorded on a given page fluctuated significantly, then it may be more appropriate to use partial deletion. For example, if the heads data set had one missing page, then for the tails data set the average number of tail results recorded on a full page could be removed from the complete tails data set.

7.2.3 Missing Not At Random

A data set where the data is Missing Not At Random (MNAR) can lead to very biased results, as the missingness of the data is directly related to the experiment or study in question. This can significantly alter the interpretation of the results.

For example, consider a medical study that calculates the survival rate for a particular type of cancer as a function of time. The study sends out a survey to each patient at regular time intervals. However, each time the survey is sent out to the same group of patients in the study, fewer and fewer survey responses are returned. The missing data could be ignored and the survival rates as a function of time could be calculated using only the complete data sets. It is clear however that the non-response of a patient may be directly correlated with their non-survival, as they will be unable to complete the survey. If this direct correlation isn't accounted for, then the reported survival rates for this particular type of cancer would become biased. For current medical studies, the use of multiple imputation is becoming increasingly more common for dealing with missing data points, instead of only performing complete-case analyses^{206–209}.

7.2.4 Missing data at MICE

At MICE, significant transmission losses occur in the MICE cooling channel. The transmission losses are Missing Not At Random as the aperture of the experiment removes many of the high radius and high transverse momenta particles. The MICE beam is collimated by the aperture of the cooling channel. The loss of momentum in the absorber material can provide an additional bias. Due to their lower momentum, certain particles will be transmitted through the MICE cooling channel that previously wouldn't have been.

Fig. 7.5 shows the radial position (top) and transverse momentum (bottom) distributions at the upstream tracker reference plane when no absorber material (left) was present in MICE beamline and when the polyethylene wedge was present. The

plots consider the full distribution at the upstream tracker reference plane (blue), the distribution that is transmitted (red) and the distribution that goes missing (green). Even when no absorber material is present in the beamline, the muon beam is significantly collimated of its high radius particles. When the polyethylene wedge is present in the cooling channel, however, more of the high radius particles are transmitted. The polyethylene wedge has reduced the momentum of some particles, meaning they will no longer be collimated by the aperture of the cooling channel.

As the demonstration of ionisation cooling and emittance exchange depends on comparing the spread of the beam in position and momentum space at the upstream and downstream tracker reference planes, it is easy to see how those results may be become biased by the missing data. An obvious bias is created if the MICE cooling channel preferentially retains cooled particles and removes heated particles due to the aperture of the MICE cooling channel, i.e. the heated particles are never measured by the downstream tracker reference plane.

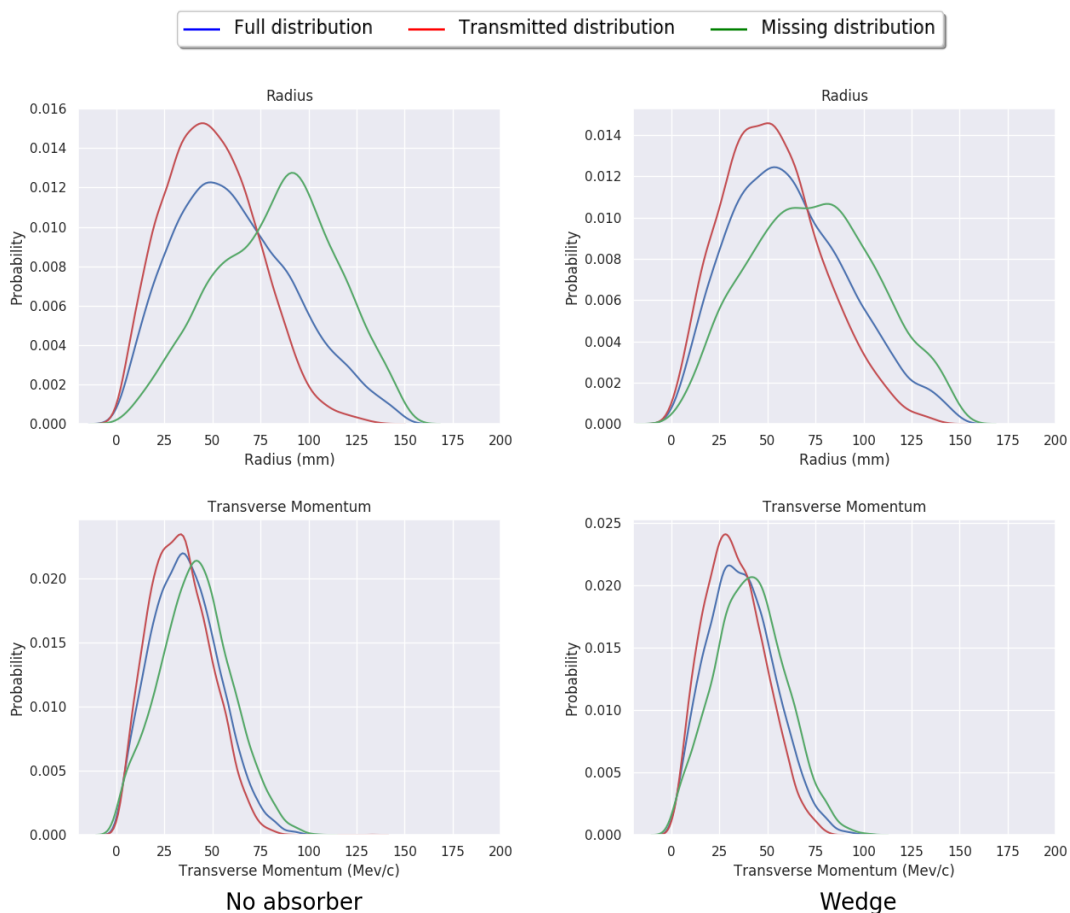


Figure 7.5: The radius and transverse momentum probability distributions at the upstream tracker reference plane when no absorber material or the polyethylene wedge was present in the MICE cooling channel. From the full probability distribution, the plots show the probability distributions of the particles that were transmitted or went missing.

7.2.5 Missing data and the covariance matrix

The emittance (Eq. 7.4), amplitude (Eq. 7.7) and density (Eq. 7.11) measurements all show a dependence on the covariance matrix. However, when transmission losses occur, the distributions in each dimension are altered, resulting in a change of the covariance matrix. If the change in the probability distributions are equal for all dimensions, then the correlations between dimensions do not change, and the change in the covariance matrix can be described by a simple scaling i.e. linearly. If the change in the probability distributions are not equal for all dimensions, then the correlations between dimensions do change, and the change in the covariance matrix becomes more complicated i.e. non-linear.

Let the covariance matrix for the full upstream sample be denoted by

$$\Sigma_{up} = \sum_i^{n_{up}} \frac{(q_i - \bar{q}_{up})^2}{n_{up}} \quad (7.13)$$

where q_i is the phase-space vector of each individual particle in the sample of size n_{up} , and \bar{q}_{up} is the mean phase-space vector for the distribution. The covariance matrices for the upstream sample which makes it downstream and the sample which does not make it downstream (i.e. it goes missing) are respectively denoted by

$$\Sigma_{down} = \sum_i^{n_{down}} \frac{(q_i - \bar{q}_{down})^2}{n_{down}} \quad (7.14)$$

$$\Sigma_m = \sum_i^{n_m} \frac{(q_i - \bar{q}_m)^2}{n_m} \quad (7.15)$$

where

$$n_{up} = n_{down} + n_m, \quad \bar{q}_{up} = \frac{n_{down}\bar{q}_{down} + n_m\bar{q}_m}{n_{down} + n_m} \quad (7.16)$$

The covariance matrix for the full upstream sample can then be rewritten as

$$\Sigma_{up} = \sum_i^{n_{up}} \left(q_i - \frac{n_{down}\bar{q}_{down} + n_m\bar{q}_m}{n_{down} + n_m} \right)^2 / (n_{down} + n_m) \quad (7.17)$$

The right-hand side of the Eq. 7.17 can then be multiplied by $\frac{(n_{down}+n_m)^2}{(n_{down}+n_m)^2}$, resulting in

$$\Sigma_{up} = \sum_i^{n_{up}} ((n_{down} + n_m) q_i - n_{down}\bar{q}_{down} - n_m\bar{q}_m)^2 / (n_{down} + n_m)^3 \quad (7.18)$$

Multiplying out the terms in Eq. 7.18 and then reordering the terms into sums over n_{down} and n_m , results in

$$\begin{aligned}
 \Sigma_{up} &= \frac{n_{down}^3}{(n_{down} + n_m)^3} \Sigma_{down} + \sum_i^{n_{down}} \left(n_{down} n_m (q_i - \bar{q}_{down})(q_i - \bar{q}_m) \right. \\
 &\quad \left. + n_{down} n_m (q_i - \bar{q}_m)(q_i - \bar{q}_{down}) + n_m^2 (q_i - \bar{q}_m)^2 \right) / (n_{down} + n_m)^3 \\
 &\quad + \frac{n_m^3}{(n_{down} + n_m)^3} \Sigma_m + \sum_i^{n_m} \left(n_{down} n_m (q_i - \bar{q}_{down})(q_i - \bar{q}_m) \right. \\
 &\quad \left. + n_{down} n_m (q_i - \bar{q}_m)(q_i - \bar{q}_{down}) + n_m^2 (q_i - \bar{q}_{down})^2 \right) / (n_{down} + n_m)^3
 \end{aligned} \tag{7.19}$$

This is the most generic solution. Simplifications can be made when symmetry considerations are taken into account. For example, consider the mean of the transverse beam components (x, y, p_x, p_y) . Providing there are no misalignments in the MICE cooling channel, the means of each of these components should coincide with zero. If the beam were to be collimated or affected by multiple Coulomb scattering, then these effects should act symmetrically, leaving the mean of the distributions unchanged, i.e. still zero. This means that for the no absorber, lithium hydride and liquid hydrogen scenarios, the mean of the x, y, p_x and p_y components of the beam for the full upstream sample, the sample which makes it downstream and the sample which goes missing should be the same i.e. zero. The simplification $\bar{q}_{up} = \bar{q}_{down} = \bar{q}_m$ can then be made for Eq. 7.19 resulting in

$$\begin{aligned}
 \Sigma_{up} &= \frac{n_{down}^3}{(n_{down} + n_m)^3} \Sigma_{down} + \frac{n_m^3}{(n_{down} + n_m)^3} \Sigma_m \\
 &\quad + \sum_i^{n_{down}} \left(2n_{down} n_m (q_i - \bar{q}_{down})^2 + n_m^2 (q_i - \bar{q}_{down})^2 \right) / (n_{down} + n_m)^3 \\
 &\quad + \sum_i^{n_m} \left(2n_{down} n_m (q_i - \bar{q}_m)^2 + n_{down}^2 (q_i - \bar{q}_m)^2 \right) / (n_{down} + n_m)^3
 \end{aligned} \tag{7.20}$$

Using Equations 7.14 and 7.15, then

$$\Sigma_{up} = \Sigma_{down} \left(\frac{n_{down}^3 + 2n_{down}^2 n_m + n_{down} n_m^2}{(n_{down} + n_m)^3} \right) + \Sigma_m \left(\frac{n_m^3 + n_{down}^2 n_m + 2n_{down} n_m^2}{(n_{down} + n_m)^3} \right) \tag{7.21}$$

$$\Sigma_{up} = \Sigma_{down} \left(\frac{n_{down}}{n_{down} + n_m} \right) + \Sigma_m \left(\frac{n_m}{n_{down} + n_m} \right) \tag{7.22}$$

Therefore

$$n_{up} \Sigma_{up} = n_{down} \Sigma_{down} + n_m \Sigma_m \tag{7.23}$$

For a radially symmetric absorber, the upstream distribution can be separated

into the covariance matrix of the sample which makes it downstream and the missing sample, weighted by their respective sample sizes.

It becomes immediately apparent that when the emittance of a beam is compared at the upstream and downstream tracker reference planes, transmission losses will significantly alter the emittance of the beam, if the missing distribution doesn't match the parent distribution (the emittance only depends on the d^{th} -root of the determinant of the covariance matrix). As the amplitude of a particle depends on the emittance of the beam, the calculated amplitude of a particle is similarly vulnerable to transmission losses. Despite the MICE collaboration being aware of these issues (these results had been presented at a MICE collaboration meeting²¹⁰), they were ignored for the publication of the MICE Nature paper⁸. For the phase-space density calculations, the transmission losses couldn't be ignored. To deal with the transmission losses for the phase-space density calculations, MICE normalized the downstream density to the upstream density by the ratio of the sample sizes⁸:

$$n_{\text{norm}} = n_{\text{up}}/n_{\text{down}} \quad (7.24)$$

where n_{up} and n_{down} are the number of particles in the upstream and downstream samples. However, Eq. 7.23 already showed that scaling the phase-space densities by their sample sizes was inadequate, as they also needed to be scaled by their changing covariance matrices. The emittance, amplitude and density plots will be considered in the following section.

7.3 Emittance, amplitude and density plots

In the following subsections, 4D plots refer to the transverse components of the beam (x, p_x, y, p_y), 2D plots refer to the longitudinal components of the beam (z, p_z) and 6D plots use all six components of the beam (x, p_x, y, p_y, z, p_z). The z and p_z coordinates of the beam are used to conserve the Liouville phase space volume.

The time component of the beam can be directly related to the z component of the beam. For the analysis in this chapter, the z component of each particle at the stations of the trackers is calculated using the time difference between the mean arrival time of the beam at a tracker station and the actual arrival time of a particle at a station. The longitudinal momentum of the particle is then used to calculate a z offset distance using that time difference. That z offset is then added or subtracted to the longitudinal position of each particle at a tracker station, giving a longitudinal position distribution at each tracker station. For the downstream tracker stations, the time component can only be reconstructed for particles that made it all the way to the TOF2 station.

The treatment described above is not ideal, as it ignores the rotation of the transverse components of the particle in that time difference. The treatment is therefore only approximately valid when the spread in the longitudinal momentum is small. As the MICE beam has been artificially created by assembling a beam from all of the individual measurements in the MICE cooling channel, the beam has no longitudinal position/time spread. Therefore, a Gaussian spread is applied to the assembled beam i.e. at the TOF1 station, a small time component is added or subtracted to each particle so that at the TOF1 station, the time distribution of the beam resembles a Gaussian distribution.

While MICE assumes that the transverse and longitudinal components of the beam can be separated, such an assumption is not made in this section. The 6D plots in this section therefore also account for any correlations between the transverse and longitudinal components of the beam.

The analysis in this section will be performed for eight different Monte Carlo (MC) simulations (summarised in Table 7.1). The different MC simulations allow a comparison between different absorber materials (no absorber, lithium hydride and the polyethylene wedge), different nominal input beam emittances (3, 6 and 10 mm) and between different magnetic field configurations (“flip” vs “solenoid”). For each simulation, a number of cuts have been applied. A cut has been applied on the time-of-flight between the TOF0 and TOF1 stations to ensure the time-of-flight is consistent with that of a muon (Fig. 1.7). At the end of the upstream tracker, the particle had to have a reconstructed momentum of 140 ± 10 MeV. For all tracker stations, the particle had to hit the tracker stations at a radius of less than 150 mm. Only tracks that had a single track in the upstream tracker, or had a single track in both trackers, were accepted.

Table 7.1: The analysis in this section considers the following Monte Carlo simulations. Shown are also the corresponding run geometries (MICE run number), the MICE magnetic field setting (either 2017-02-7 for “flip” mode or 2017-02-6 for “solenoid” mode), the absorber setting (no, lithium hydride or polyethylene wedge) and the nominal beam emittance (3, 6 or 10 mm) and momentum (140 MeV).

MC Run Number	MICE Run Number	Field Setting	Absorber	Emittance and Momentum
167	10444	Flip	No	3-140
168	10446	Flip	No	6-140
169	10447	Flip	No	10-140
171	10481	Flip	LiH	6-140
175	10534	Flip	Wedge	6-140
247	10318	Solenoid	No	6-140
251	10581	Solenoid	Wedge	6-140
258	10509	Solenoid	LiH	6-140

7.3.1 Emittance

The transverse, longitudinal and 6D emittances for the eight different MC simulations were calculated at the ten tracker stations and subsequently plotted in Figures 7.6, 7.7 and 7.8. The transverse emittance plots consider three different distributions that passed the radial and momentum cuts: the full upstream sample (Full upstream) where particles were reconstructed in the upstream tracker, the sample which made it downstream (Downstream cut) where particles were reconstructed in both the upstream and downstream trackers and the sample which made it to TOF2 (TOF2 cut) where particles were also reconstructed in both trackers. The longitudinal and 6D emittance plots only consider the sample which made it all the way to TOF2, as they require a downstream time co-ordinate. Additionally, each plot considers the MC reconstruction sample and the corresponding MC truth sample.

The transverse emittance plots (Fig 7.6) show that the emittance for the full upstream distribution bears no resemblance to the emittance of the downstream distribution. This is because the beam has been collimated by the aperture of the MICE cooling channel. The beam emittance has been artificially reduced (except for the 3-140 no absorber beam, which didn't pass through the diffuser material). At the tracker reference planes, the beam emittances can only be compared for the beams that contained the same sample of particles, i.e. by applying a downstream cut or a TOF2 cut. For the 6-140 and 10-140 no absorber settings, the emittance grew between the tracker reference planes when only the particles of the downstream cut or TOF2 cut were considered. Both cuts were biased, however, as they only included particles that stayed within the MICE cooling channel, i.e. they excluded many of the heated particles.

For the 6-140 “flip” no absorber, lithium hydride and polyethylene wedge scenarios, the emittance of the full upstream distributions were similar. However, when one considers the upstream samples which made it downstream (or to TOF2), differences between the samples became apparent. The different absorber materials scattered different particles beyond the aperture of the MICE cooling channel that previously would have been transmitted. Similarly, some particles remained within the aperture of the MICE cooling channel as a result of the energy loss they experienced in the differing absorber materials.

The emittance changes within the downstream tracker were the most interesting feature of the plots. For the no absorber scenario, the emittance of the beam showed some minor changes within the downstream tracker. This indicates that a transfer between the transverse and longitudinal components of the beam may have taken place. For the lithium hydride and polyethylene wedge scenarios, the emittance of the beam grew within the downstream tracker. The change in the beam emittance

can be explained when the energy loss of the beam is considered.

For the no absorber scenario, no energy loss occurs, resulting in no appreciable change in the momentum spread. For the lithium hydride scenario, nearly all of the particles will lose a similar amount of energy. The energy loss distribution will still have a small spread, however. This will broaden the spread of the longitudinal momentum distribution. As the solenoid transfer matrix has a dependence on the momentum of the particle going through the solenoid (Eq. 1.32), each particle in the beam will experience a slightly different transfer matrix depending on the momentum of the particle. This will affect the helical trajectory a particle makes within a given longitudinal distance. The beam therefore appears to filament in the transverse position and momentum plane, causing the transverse emittance of the beam to grow. The effect becomes even more pronounced for the polyethylene wedge scenario, as the polyethylene wedge induces a far greater spread in the longitudinal momentum distribution of the beam.

The corresponding growth in the longitudinal emittances can be seen in Fig. 7.7. The MC reconstruction sample underestimated the longitudinal emittance of the beam for the upstream tracker and overestimated the longitudinal emittance of the beam for the downstream tracker due to the biases of the MICE reconstruction when it calculated the longitudinal momentum of a particle. Only the Monte Carlo truth samples should therefore be compared.

Similar plots for the 6D emittance can be seen in Fig. 7.8. In each scenario, the 6D emittance of the beam has grown between the tracker reference planes, including crucially for the no absorber material scenario. This indicates that the MICE cooling channel does not in fact preserve emittance and challenges some of the assumptions MICE has used. These include the assumption of using linear optics (where emittance would be preserved), that the longitudinal and transverse phase-spaces can therefore be treated independently and that the demonstration of ionization cooling can be performed using only the transverse components of the beam^{8,78}.

For all eight scenarios, the transverse emittances of a selection of particles grew to be larger at the end of the downstream tracker than at any point in the upstream tracker. Due to the high transmission losses in the MICE cooling channel, it is therefore unclear whether any ionisation cooling has taken place (without the use of an RF cavity) or whether only a selection effect is being observed. An alternative approach is to show ionization cooling for only the remaining fraction of the beam, i.e. by showing heating and cooling for individual particles. This can be done by either calculating the amplitude or the phase-space density of each particle within a beam. However, as the amplitude of a particle depends on the emittance of the beam, an obvious issue immediately arises for the amplitude calculation. This will be investigated in the following subsection.

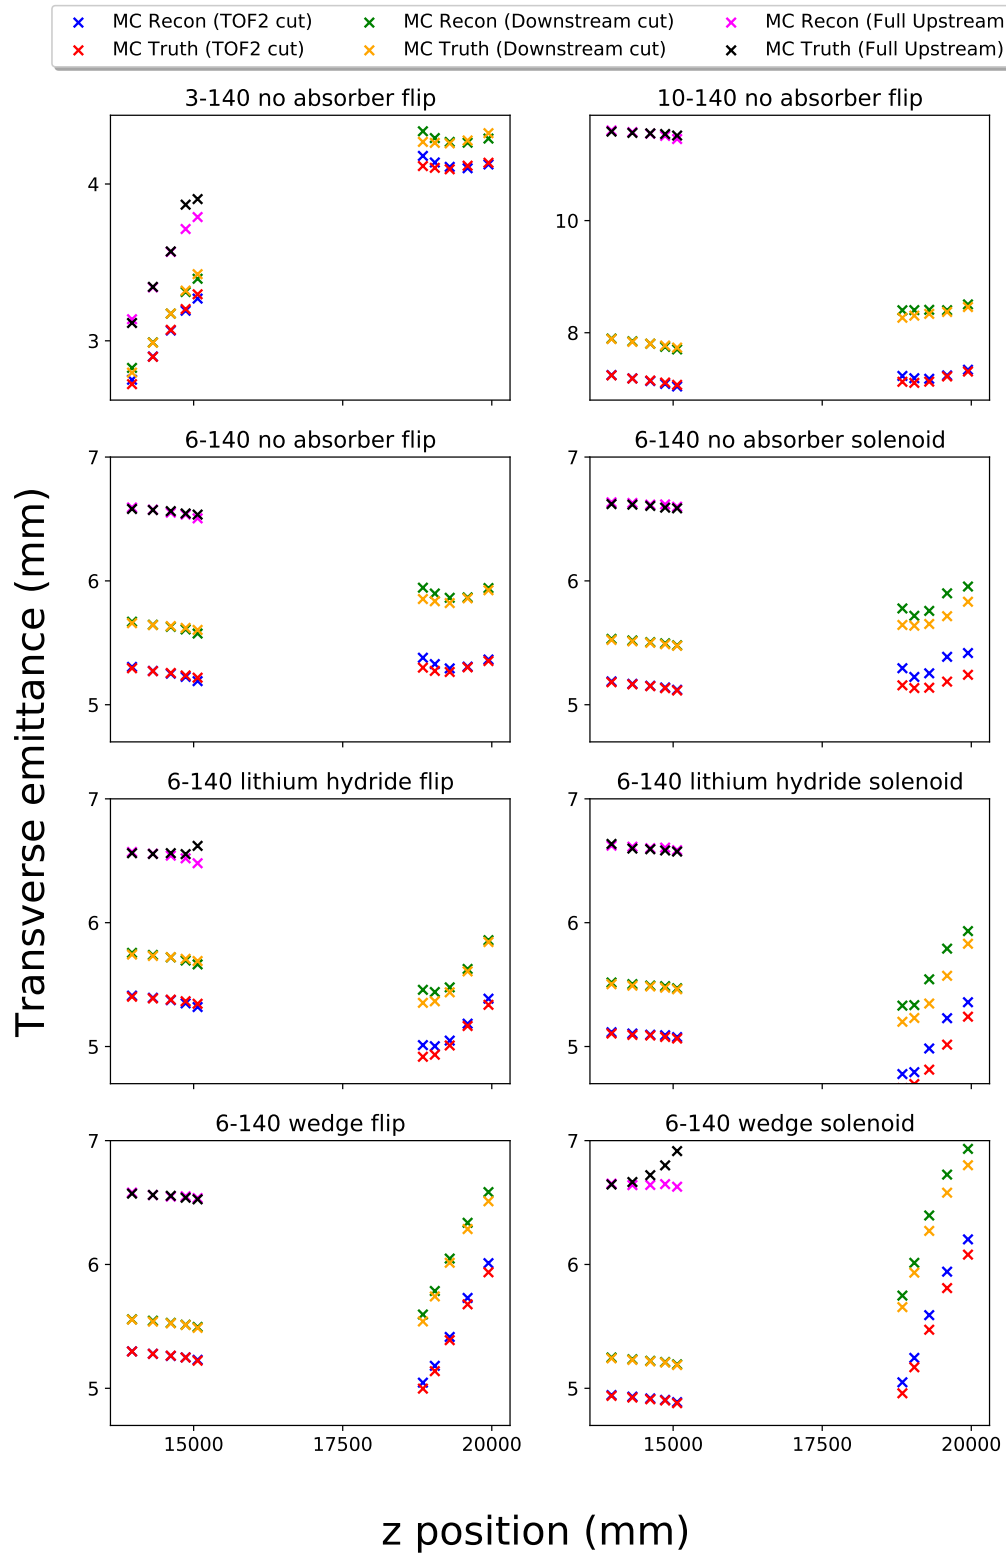


Figure 7.6: The transverse emittance for various beams and for various beam samples at the ten tracker stations of the MICE cooling channel.




Figure 7.7: The longitudinal emittance for various beams and for various beam samples at the ten tracker stations of the MICE cooling channel.

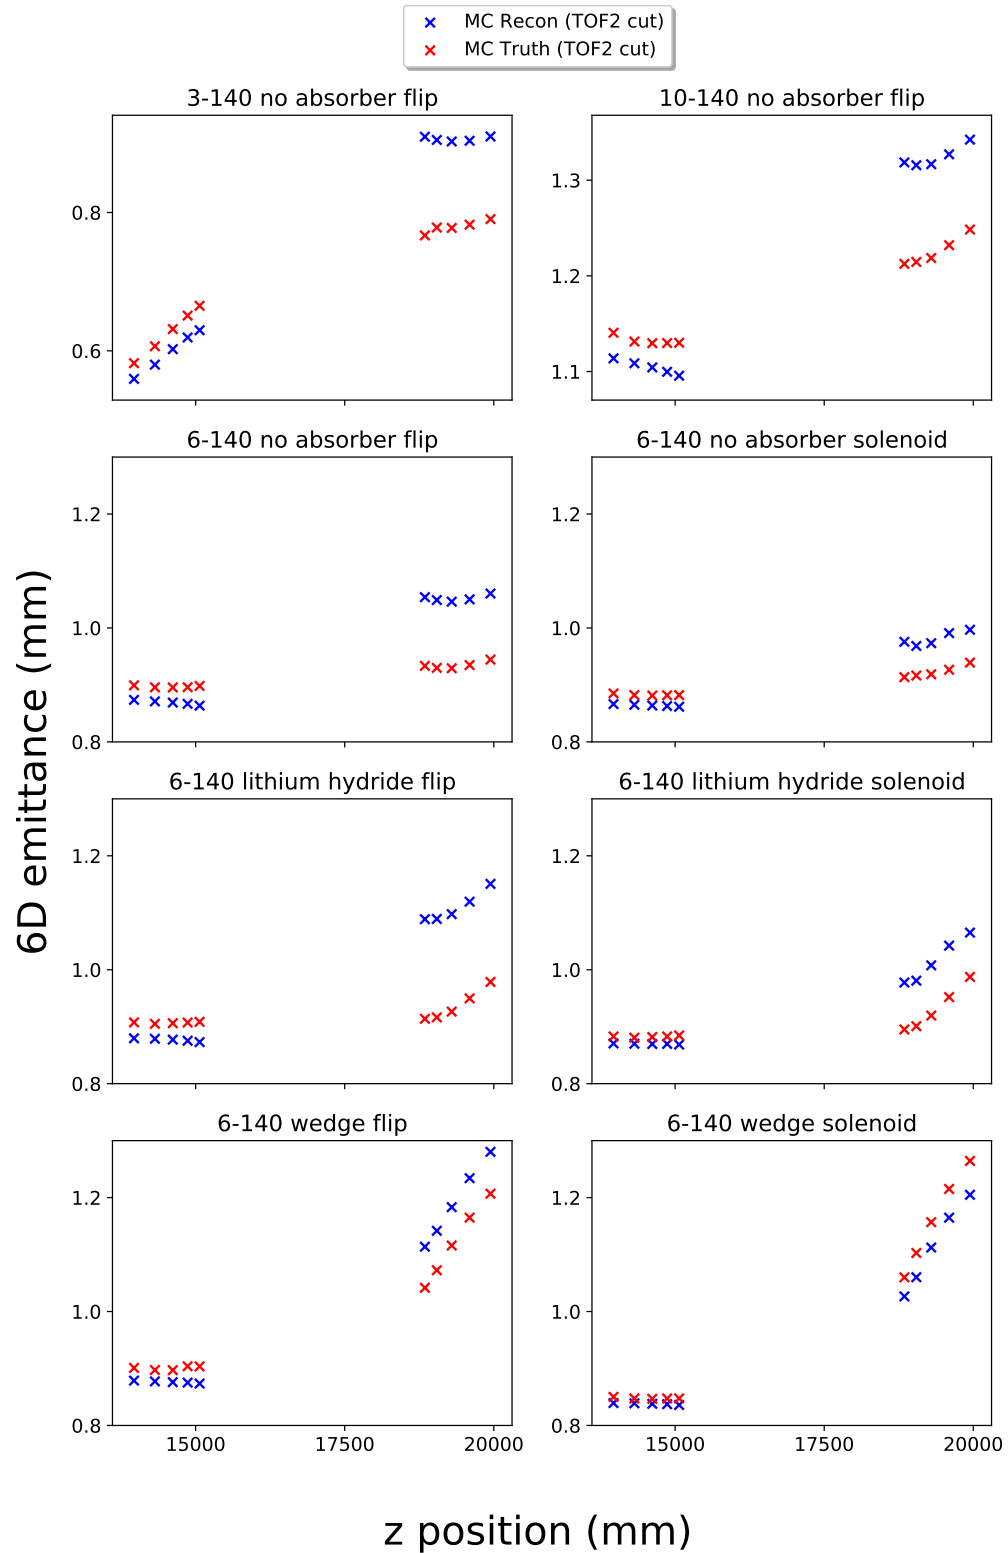


Figure 7.8: The 6D emittance for various beams and for various beam samples at the ten tracker stations of the MICE cooling channel.

7.3.2 Amplitude

The amplitude of a particle was given by Eq. 7.7 and is simply the emittance of the beam multiplied by the Mahalanobis distance of a particle to the beam centroid, with the emittance of the beam simply acting as a scaling factor. If the emittance of the beam has changed, however, then the covariance matrix has also changed. This will affect the calculation of the Mahalanobis distance if either the beam centroid or the correlations between the dimensions have changed. The effect of the bias on the calculated amplitude can therefore be far more subtle if the bias on the calculated emittance due to transmission losses is partially balanced out by the changing Mahalanobis distance.

This can be seen in Fig. 7.9 which considers the MC truth transverse amplitude distribution for a beam that has passed through the polyethylene wedge and made it to TOF2. Only minor variations in the amplitude distribution can be seen for the five upstream tracker stations. However, for the five downstream tracker stations, the amplitude distribution changes significantly. This is similar to the corresponding emittance plot (Fig. 7.6). The plot indicates that the beam has initially been cooled (lower amplitudes) as a result of passing through the wedge, and then begins to heat (higher amplitudes) as it passes through the five downstream tracker stations. The growth in the transverse amplitude distribution is due to the wedge creating a position-momentum correlation in the beam (dispersion). Similarly, as for the emittance plots, it is not clear whether cooling has occurred or whether a selection effect has been observed.

To deal with the effects of optical aberrations and transmission losses affecting the covariance matrix, MICE applied an algorithm (without proof) that continuously recalculates the amplitudes of particles closer to the low amplitude beam core for the MICE Nature paper⁸: “To expose the behaviour in the beam core, independently of aberrations affecting the beam tail, V [the covariance matrix] . . . [is] recalculated for each amplitude bin, including particles that are in lower-amplitude bins and excluding particles that are in higher-amplitude bins. This results in a distribution that, in the core of the beam, is independent of scraping effects and spherical aberrations”.

The algorithm was tested by Drielsma⁷⁴ on a Gaussian beam (top plots of Fig. 7.10), and a Gaussian beam passing through a nonlinear lens with an aberration coefficient of $C_\alpha = -10^{-4} \text{ mm}^{-2}$ and collimated by a 150 mm radius aperture (bottom plots of Fig. 7.10). The algorithm performs well on the Gaussian beam. No transmission losses occur, and the correlations between the dimensions of the covariance matrix are left unchanged. For the Gaussian beam that passes through the non-linear lens, however, the algorithm reproduces the expected χ^2 distribution for a Gaussian beam near the core of the beam, but deviates at higher amplitudes. This




Figure 7.9: The MC truth transverse amplitude distribution at the ten tracker stations for a “flip” mode beam that has passed through the polyethylene wedge and made it to TOF2.

is not too surprising, as the non-linear lens will mainly affect particles at the edge of the distribution. This similarly applies to the aperture cut. For very high amplitudes, the algorithm appears to leave the amplitude unchanged. There is however one significant problem with the MICE argument, even at the low amplitude beam-core, particles go missing (Fig. 7.11).

The main issue with the MICE approach is that it offers no mathematical proof that the algorithm is a valid approach to compare the continuously recalculated amplitude distributions at the upstream and downstream tracker reference planes. As Drielsma⁷⁴ noted, the amplitude “assumes a single underlying covariance matrix to describe the full ensemble. Unless the transmission losses are isotropic in amplitude, the shape of the covariance matrix is not preserved and neither is the amplitude measurement”. Once the covariance matrix begins to change as a result of the algorithm, the amplitudes of the particles within a distribution, calculated over the different covariance matrices, are no longer comparable. To see this, consider Fig. 7.12. For the same set of particles at the upstream tracker reference plane it has calculated the amplitude of the particles using two different covariance matrices, the covariance matrix of the full upstream sample and the covariance matrix of only the particles that made it to TOF2 when no absorber material was present in the

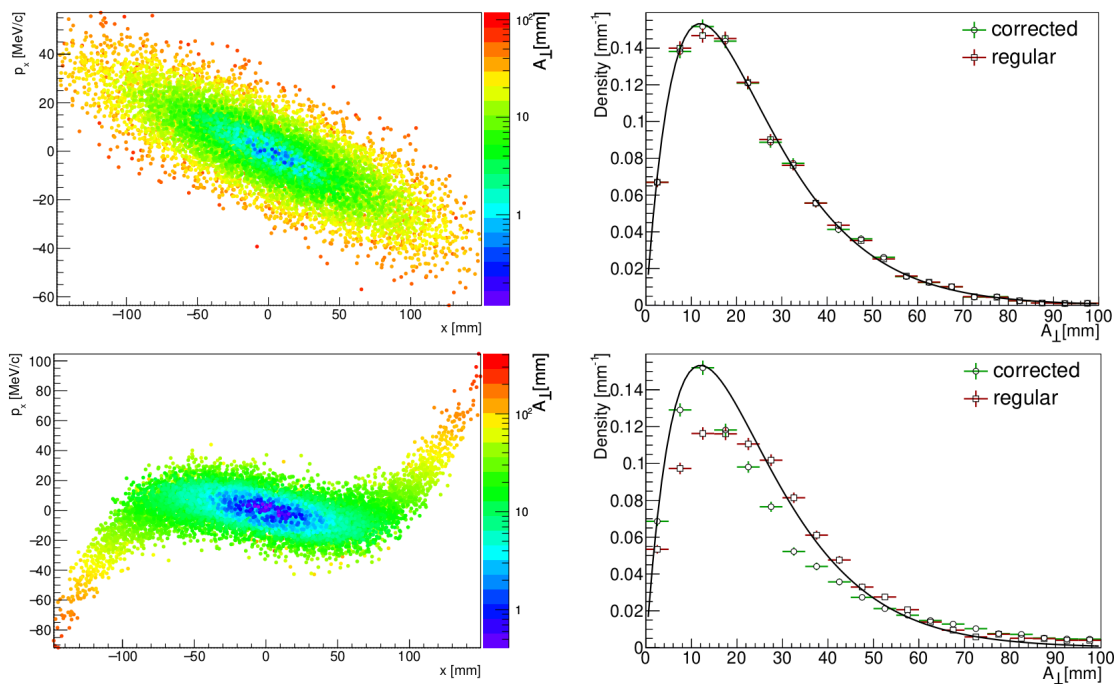


Figure 7.10: MICE uses an algorithm to continuously recalculate the amplitudes of particles closer to the low amplitude beam core. This algorithm was tested by Drielmsa⁷⁴ on a Gaussian beam (top plots), and a Gaussian beam passing through a nonlinear lens with an aberration coefficient of $C_\alpha = -10^{-4} \text{ mm}^{-2}$ with a radius aperture of 150 mm (bottom plots). The algorithm performs well at amplitudes unaffected by non-linearities or transmission losses.

MICE cooling channel. Fig. 7.12 then shows the residual between the calculated amplitudes, with Table 7.2 considering the mean and RMS residual amplitudes for a number of different beams. For all scenarios, the mean amplitude is lower when it is calculated over its own distribution rather than over the full upstream distribution, with the mean residual amplitude changing depending on the beam size and absorber material considered. For the no absorber scenario, as the nominal beam emittance beam grows, so does the mean residual amplitude. The effect is less noticeable for the lithium hydride and polyethylene wedge scenarios.

The difference between the no absorber and absorber cases can be explained when the acceptance of the MICE cooling channel is considered. Fig. 7.5 showed that it was primarily high radius particles rather than high transverse momentum particles that caused transmission losses. As the transmission losses between the phase-space dimensions are asymmetrical, the correlations of the covariance matrix begin to change. However, when an absorber material is present in the MICE cooling channel, the acceptance of high radius particles increases and reduces for low radius particles as a result of multiple Coulomb scattering and energy loss. This means the asymmetry of the transmission losses between the position and momentum space is




Figure 7.11: The amplitudes of the particles that go missing (red) from the full upstream amplitude distribution (blue) for the no absorber 6-140 “flip” mode beam after the TOF2 cut has been made. Transmission losses can even occur for low amplitude particles.




Figure 7.12: The residual true amplitude distribution between the amplitudes calculated over their own sample sizes and as if part of the full upstream distribution for the no absorber 6-140 “flip” mode beam after the TOF2 cut has been made.

less severe for the absorber scenarios, resulting in the correlations of the covariance matrix changing less. This in turn leads to a smaller change in the calculated amplitude.

However, the fundamental issue remains that due to the non-linearities of the MICE cooling channel, the emittance of the beam and therefore the amplitudes of the particles are not conserved in the MICE cooling channel. These emittances and amplitudes then can be biased by any transmission losses. For this reason, phase-space density will be considered in the next subsection.

Table 7.2: The MC truth mean amplitude of the distribution and the MC truth mean and RMS residual amplitude between the amplitudes calculated over their own distribution and the amplitudes calculated as if they were part of the full upstream distribution.

MC Run	Absorber	Beam	Distribution Mean amplitude	Mean Residual	RMS Residual
167	No	3-140	13.21	0.13	1.13
168	No	6-140	21.08	0.25	1.98
169	No	10-140	28.35	0.75	4.18
171	LiH	6-140	21.60	0.03	0.76
175	Wedge	6-140	21.31	0.04	0.77

7.3.3 Density

The main benefit of using phase-space density plots over emittance or amplitude plots is that the 6D position-momentum phase-space remains conserved even in non-linear systems as a result of Liouville's theorem. However, transfers between the longitudinal and transverse phase spaces can still take place. Additionally, a large momentum spread in the beam will produce similar issues as for the emittance and amplitude plots. This can be seen in Fig. 7.13 which considered the phase-space densities at the ten tracker stations when the polyethylene wedge was present in the MICE cooling channel. The figure shows that the transverse phase space density has increased as a result of the wedge. Transmission losses will however introduce a selection bias, i.e. the removal of low density particles that would heat the beam.

Similarly, as for the emittance and amplitude plots, the dispersion of the beam results in a reduction of the transverse phase-space density of the beam through the downstream tracker stations. The transverse phase-spaces at the downstream tracker stations are not directly comparable, however, when the beam has a large longitudinal momentum spread. A particle's individual Hamiltonian will remain conserved unless it encounters a dissipative force. However, collective effects need to be treated more carefully. From Liouville's theorem, it is the phase-space density that remains constant as a function of time. However, over that time period, the




Figure 7.13: The MC truth transverse density distribution at the ten tracker stations for the 6-140 “flip” mode beam for the distribution that made it to TOF2 when the polyethylene wedge was present in the MICE cooling channel.

phase-space volume can change shape. Therefore, if the time-of-flight differences for a selection of particles between two tracker stations is large, the particles arriving at the downstream tracker reference plane will be part of a continuously evolving phase-space volume passing through the downstream tracker (i.e. a helically rotating phase-space volume), with the measured phase-space volume not reflective of any of the particles. A possible correction procedure could involve using the mean arrival time at a tracker station for a collection of particles and correcting the transverse components of the beam by the difference in the mean arrival time and the actual arrival time. This can only be done if the magnetic field is sufficiently well known. For the polyethylene wedge, the magnitude of the time-of-flight differences between the particles continues to grow through the downstream tracker as a result of the dispersion in the beam, resulting in the calculated transverse phase-space density continuing to change.

The 6D phase-space density appears to be approximately conserved for the 6-140 no absorber scenario (Fig. 7.14), while the 10-140 no absorber scenario shows a slight reduction in the 6D density. This may be due to a transfer between the transverse and longitudinal beam components. In this case, the transverse beam components may require a correction, as high radius particles with a larger transverse momentum will have longer time-of-flights paths than low radius particles of the same momentum.

For the polyethylene wedge case, the 6D density has decreased, while for the lithium hydride case, the 6D density has slightly increased. The distributions are biased however by transmission losses.

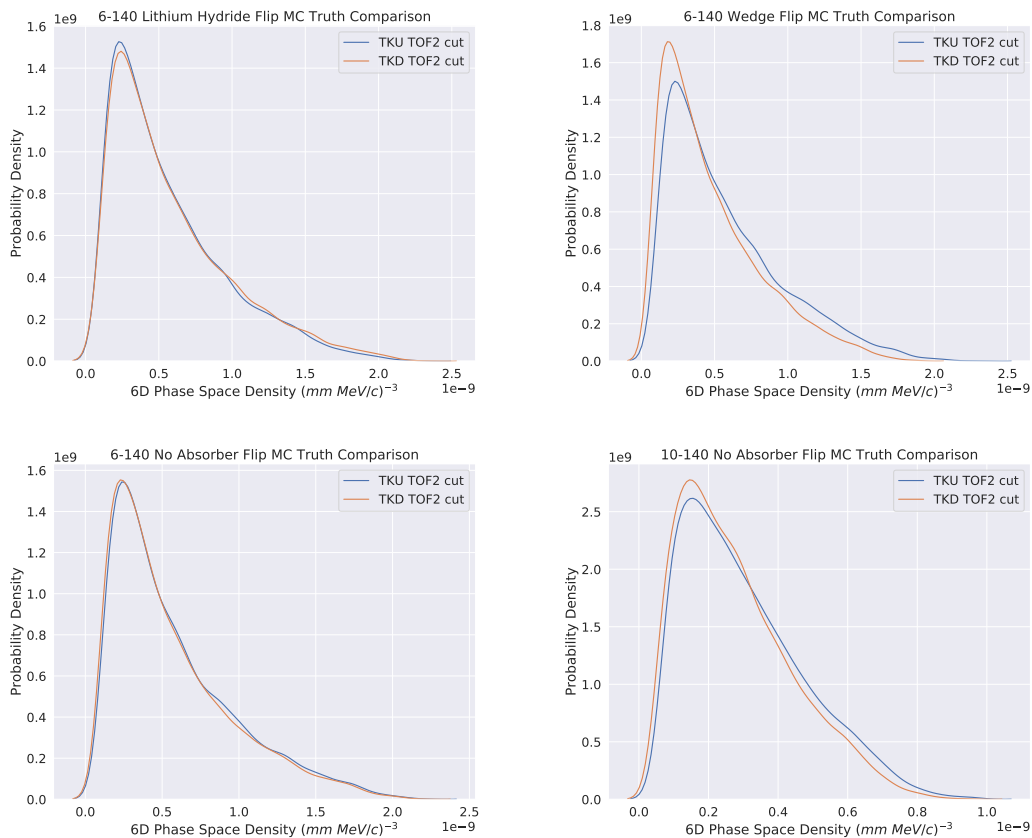


Figure 7.14: The true 6D density distribution in “flip” mode for the distribution that made it to TOF2 at the tracker reference planes. *Top left:* 6-140 lithium hydride, *top right:* 6-140 polyethylene wedge, *bottom left:* 6-140 no absorber, *bottom right:* 10-140 no absorber.

In an ideal scenario, the full upstream sample could be compared to the measured downstream sample. However, transmission losses mean this is not possible. This can be seen in Fig. 7.15 which considers the MC truth transverse density distribution for the 6-140 no absorber “flip” mode beam for the full upstream distribution, the distribution that made it downstream and the distribution that made it to TOF2 at the tracker reference planes. The transmission losses bias the downstream sample. To get around the issues of transmission losses, MICE normalizes the downstream phase-space density by the ratio of the sample sizes between the upstream and downstream distributions. However, as was previously shown, such a normalization is incorrect if the particle distribution of the missing sample doesn’t match that of the full upstream sample (as is the case in MICE). In this case, the normalization procedure must also account for the changing covariance matrix to be able to compare

the full upstream distribution to the remaining downstream distribution (Eq. 7.23). Such a correction procedure will be attempted in the next section.




Figure 7.15: The MC truth transverse density distribution for the 6-140 no absorber “flip” mode beam showing the full upstream distribution, the distribution that made it downstream and the distribution that made it to TOF2 at the tracker reference planes.

7.4 Unbiased cooling

The aim of this section is to show ionization cooling unaffected by transmission losses. However, this is not possible, as the undetected particles can’t be measured. Phase-space density plots can therefore be used to show cooling for a fraction of the beam within a given phase-space volume. This means the downstream particles need to be calculated over the same phase-space volume as the upstream sample, after that phase-space volume has been transported downstream. For the remaining downstream sample, the densities of the particles in that phase-space volume can then be shown. This means that for a fraction of the beam, an unbiased cooling signal can be shown unaffected by transmission losses. The plots can also show at what fraction of the beam transmission losses begin to spoil the cooling signal.

The full downstream phase-space volume (and therefore the covariance matrix for that beam sample) can be calculated using an imputation technique. For the particles that were transported downstream from the upstream tracker reference

plane when no absorber was present in the MICE cooling channel, a transfer matrix can be found. If the Residuals between the true position (and momentum) and the position (and momentum) predicted by the transfer matrix is small enough, the position (and momentum) of where the un-transmitted particles should have ended up can be determined to a reasonable accuracy. The transfer matrix can then be used to transport all the particles of the full upstream sample downstream. The transported particles can then be used to calculate the covariance matrix at the downstream tracker reference plane that should be used to calculate the densities for the particles that were actually transmitted. The particles at the upstream and downstream tracker reference planes can now be calculated over the same phase-space volumes. Therefore, for the remaining fraction of the transmitted beam, an unbiased change in the phase-space density can be shown.

To summarise, six different distributions will be considered:

- The Transmitted Upstream Sample, i.e. after a downstream cut or a TOF2 cut has been applied.
- The Transmitted Downstream Sample. The density distribution for only the measured sample.
- The Full Upstream Sample.
- The Predicted Downstream Sample. The transfer matrix is applied to the Full Upstream Sample, to predict what the downstream sample should look like without any absorber or acceptance effects.
- The Full Downstream Sample. It uses the predicted downstream sample, but substitutes in all of the actual measurements at the downstream tracker reference plane if they exist.
- The Unbiased Downstream Sample. It takes the full downstream sample, but removes any particles whose density was predicted, i.e. it only includes the downstream measurements, but with their density calculated over a phase-space volume that didn't have any transmission losses.

Fig. 7.16 shows the six distributions for the 6-140 no absorber scenario when a downstream cut has been applied to the appropriate distributions. The plot shows that the transverse phase-space density for the full upstream sample (blue) is nearly identical to the density of the predicted downstream sample (green). This means the two distributions occupy the same phase-space volume, with the covariance matrix describing the samples having only been changed by the magnetic field transporting the particles, and not by any transmission losses.




Figure 7.16: The MC truth transverse phase space density distribution for the six beam samples considered in this section, when a 6-140 no absorber “flip” mode beam passes through the MICE cooling channel.

The cumulative survival plots of the six different density distributions can be seen in Fig. 7.17 for a 6-140 no absorber beam (top) and a 6-140 LiH beam (bottom). In the left plots, the full upstream sample can be compared to the predicted downstream sample, and shows a small amount of heating has occurred as a result of the transfer matrix. In the middle plots, the transmitted samples have been replaced by the transmitted upstream sample calculated over the full upstream distribution and the predicted downstream sample. As the differences between those two samples are similar to the differences between the full upstream sample and the predicted downstream sample for both absorber scenarios, the transfer matrix approach has found an appropriate downstream phase-space volume for both absorber scenarios. Finally, in the right plots, the predicted downstream distribution is replaced by the full downstream distribution, while the transmitted predicted distribution is replaced by the unbiased downstream sample. The unbiased downstream sample can now be compared directly to the full upstream distribution. If the unbiased downstream distribution (red) is above the full upstream distribution (blue), then the beam has been cooled for that fraction of the beam. This can be more clearly seen in the ratio plots.

Figures 7.18, 7.19 and 7.20 show the ratios of various beam samples for the 6-140 “flip” mode beam when no absorber material (left), lithium hydride (middle) or a polyethylene wedge (right) was present in the MICE cooling channel. A downstream

cut has been applied for the top plots, while a TOF2 cut has been applied for the bottom plots. In the following ratio plots, a value above one indicates heating, where the phase-space volume has increased, while a value below one indicates cooling, where the phase-space volume has decreased.

Fig. 7.18 shows the ratio of the transmitted upstream sample calculated over the full upstream covariance matrix versus the predicted transmitted downstream sample. The plots show that the transfer matrix assumes a similar amount of heating for all three absorber scenarios. This means a direct comparison between the three absorbers scenarios can be made when the actual downstream samples are considered.

Fig. 7.19 shows the ratio of the full upstream sample versus the transmitted upstream sample calculated over the covariance matrix of the full upstream sample. The plots show at what fraction of the beam transmission losses kick in that could spoil the ionization cooling result.

Finally, Fig. 7.20 shows the ratio of the full upstream sample over the unbiased downstream sample. This is the ionization cooling plot, unaffected by transmission losses (though not necessarily unaffected by any transfers between the longitudinal and transverse beam components). For the no absorber scenario, heating of the beam can be seen, while for the lithium hydride Scenario, cooling of the beam can be seen for approximately 70% of the beam before transmission losses spoil the result. The cooling result is spoiled even earlier when a stricter TOF2 cut is applied. The transmission losses pose a significant issue for the polyethylene wedge scenario where the TOF2 cut is required for the longitudinal components of the beam. The curves in the polyethylene wedge plots are due to the dispersion in the beam that needs to be corrected for, however such corrections are beyond the scope of this thesis. As the 6D density remains approximately conserved for the no absorber scenario (7.14), the heating seen for the transverse phase-space volume is due to a transfer between the longitudinal and transverse components of the beam.

The difference between the no absorber and lithium hydride plots (Fig. 7.20) gives an indication of how much the lithium hydride absorber has reduced the transverse phase-space volume of the majority of the beam. The reduction in the phase-space volume appears far more modest in the 6D density plots (Fig. 7.14) as the lithium hydride absorber has also increased the longitudinal phase-space volume of the beam. To decrease the longitudinal phase-space volume of a beam requires an RF cavity. For MICE, the RF cavity can be used to ensure that the reduction in the transverse phase-space volume can be maintained as the longitudinal phase-space volume of the beam is reduced.

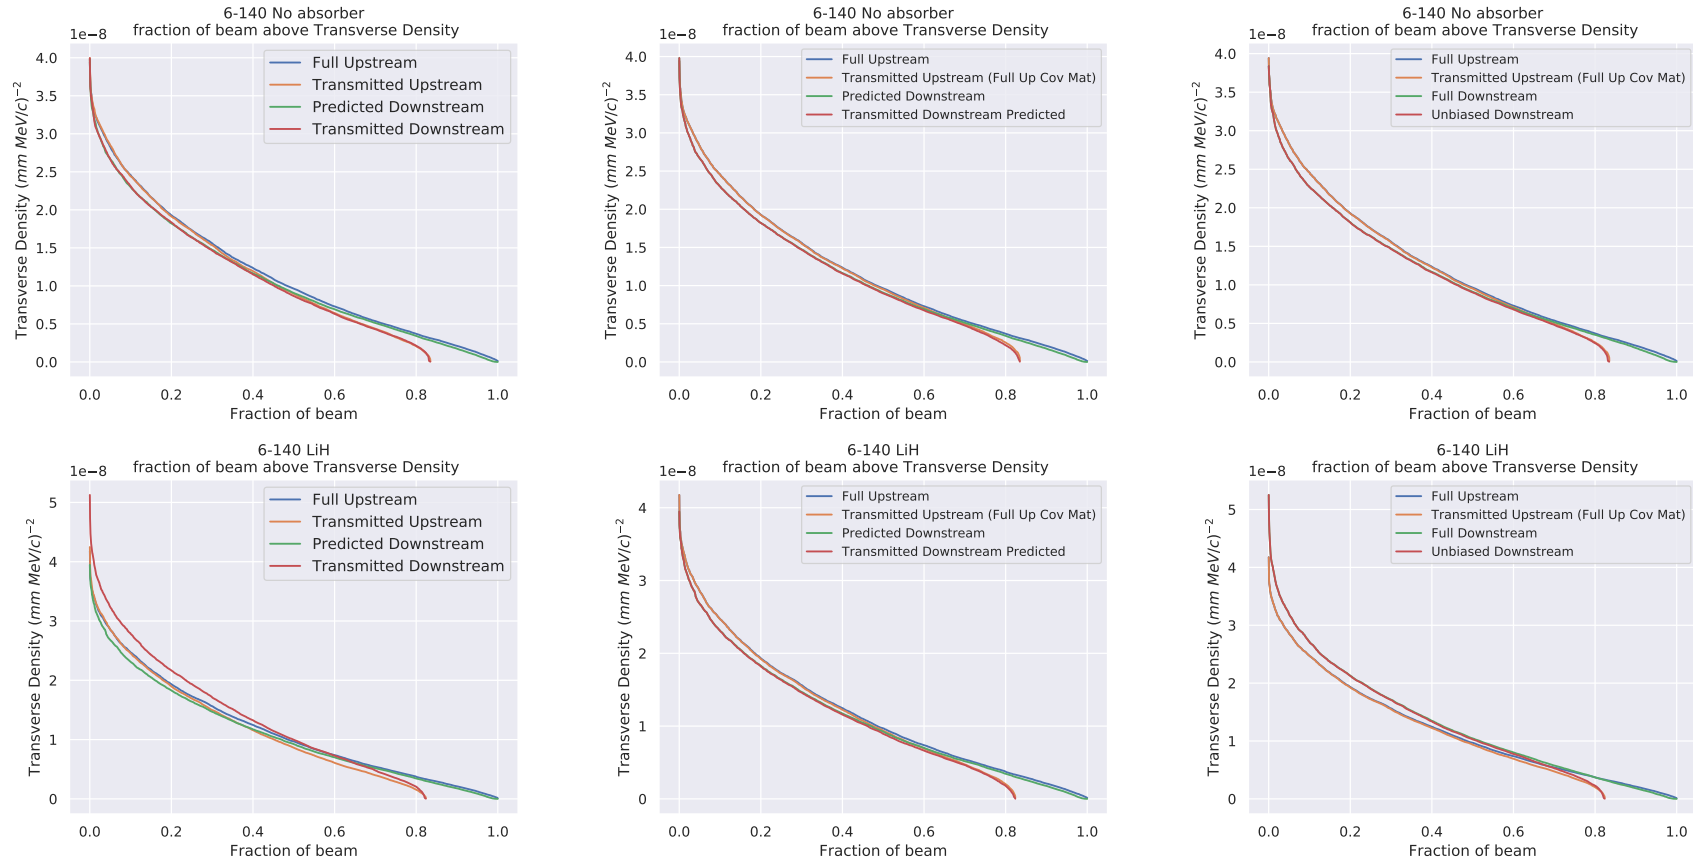


Figure 7.17: The MC truth transverse phase space density distributions for the six different beam samples considered in this section. The distributions are considered across three plots for the 6-140 no absorber beam (*top*) and the 6-140 lithium hydride beam (*bottom*).

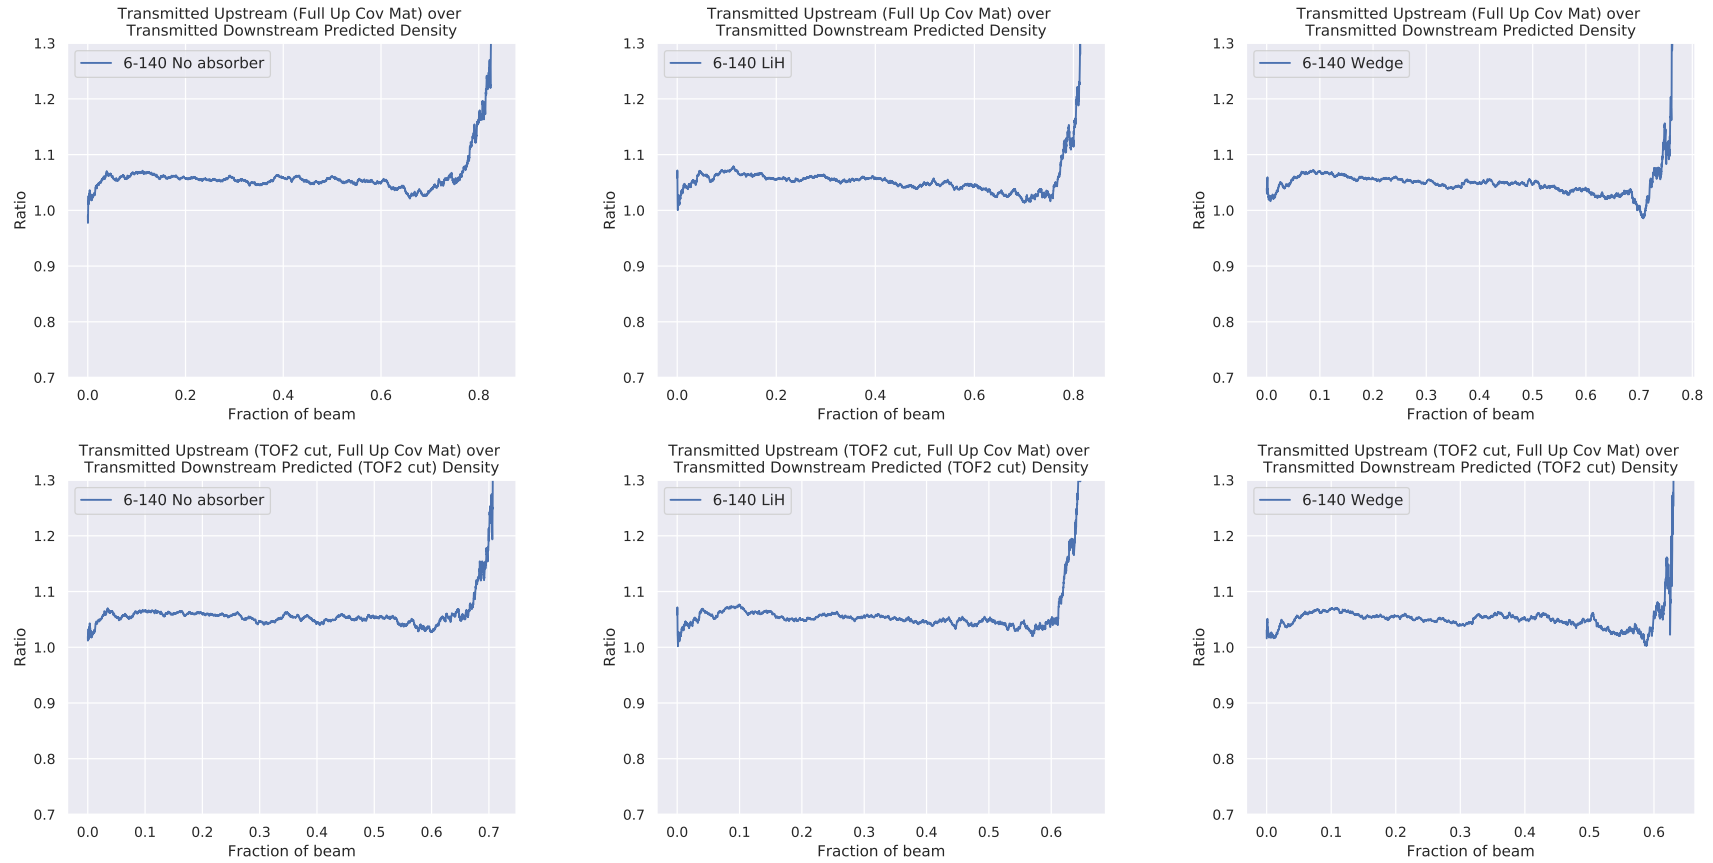


Figure 7.18: The ratio of the transmitted upstream sample calculated over the full upstream covariance matrix versus the predicted transmitted downstream sample. *Left:* 6-140 no absorber, *Middle:* 6-140 lithium hydride, *Right:* 6-140 polyethylene wedge, *top:* Downstream cut, *bottom:* TOF2 cut.

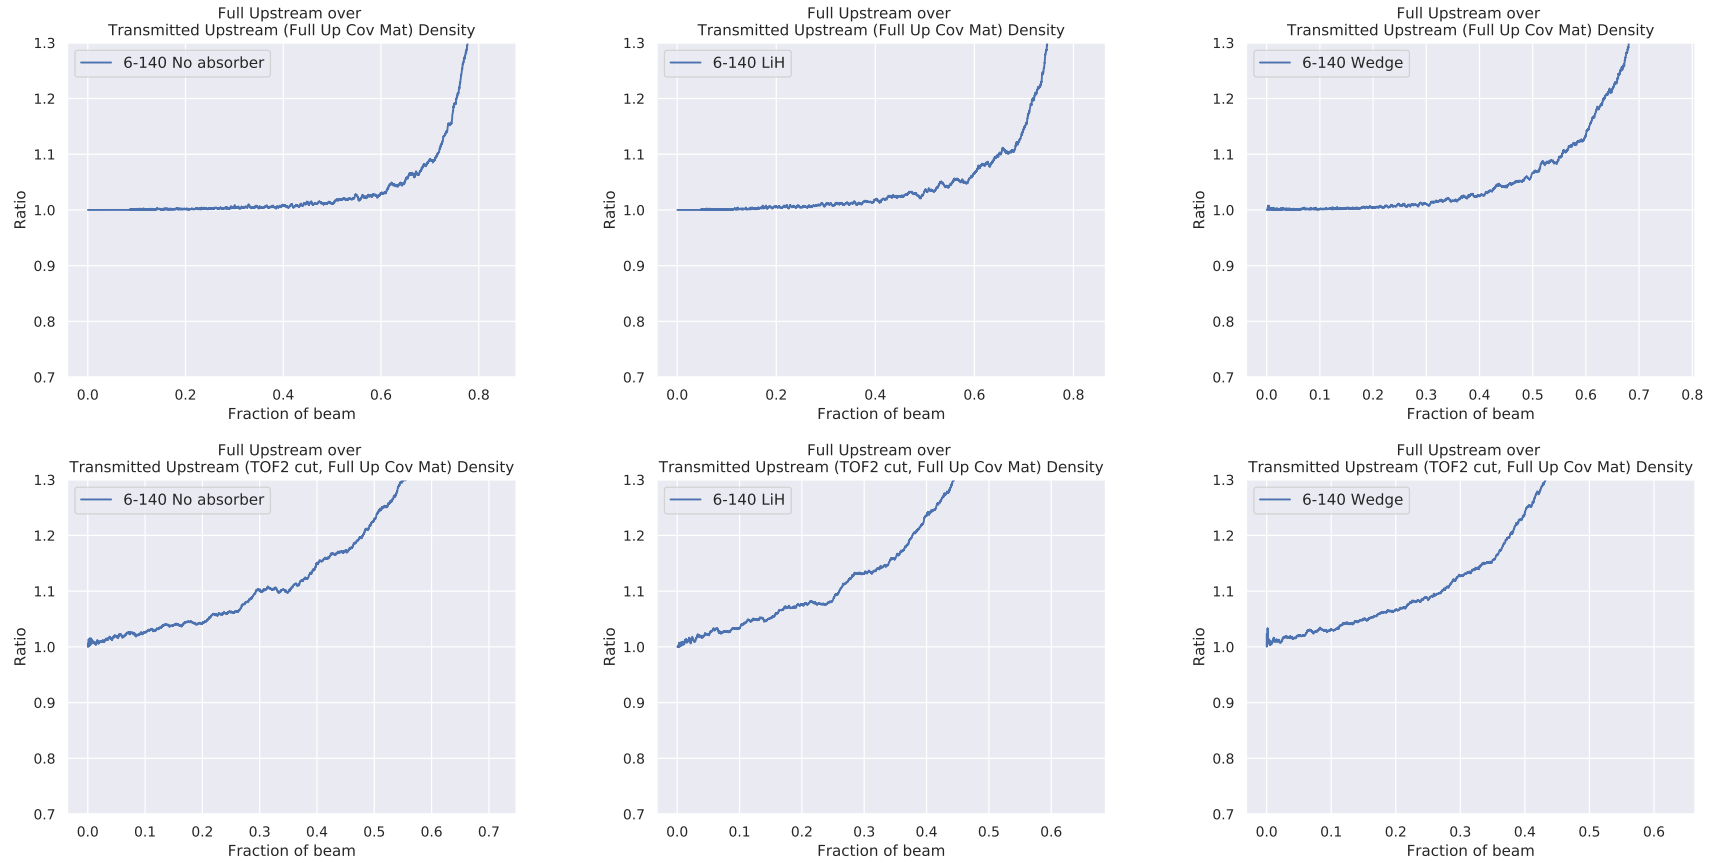


Figure 7.19: The ratio of the full upstream sample versus the transmitted upstream sample, calculated over the covariance matrix of the full upstream sample. *Left:* 6-140 no absorber, *Middle:* 6-140 lithium hydride, *Right:* 6-140 polyethylene wedge, *top:* Downstream cut, *bottom:* TOF2 cut.

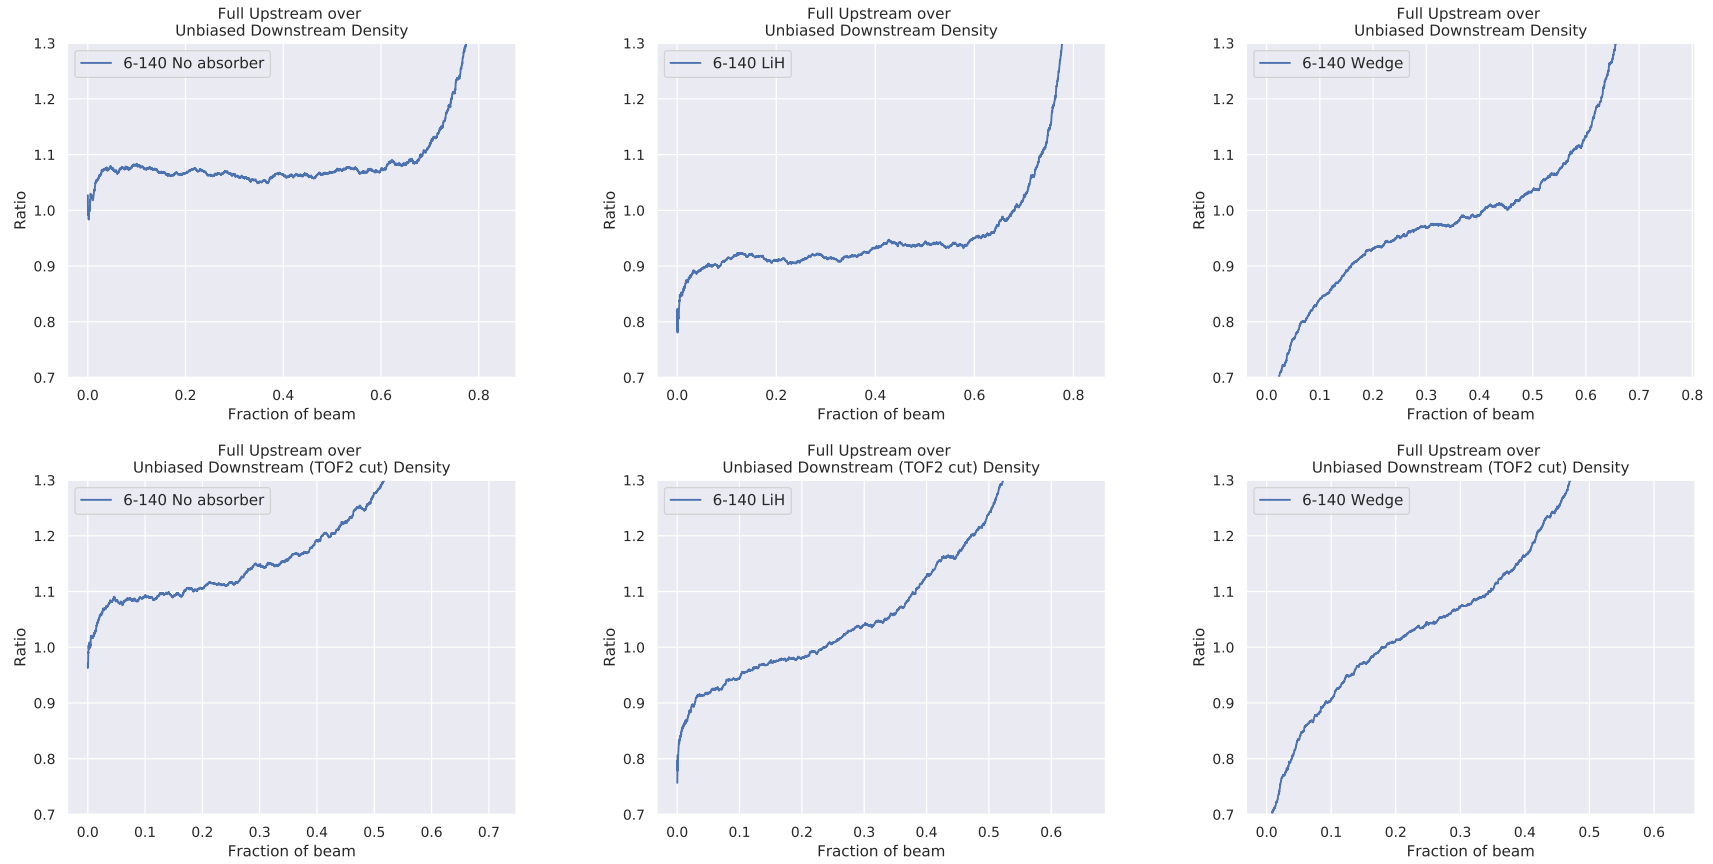


Figure 7.20: The ratio of the full upstream sample over the unbiased downstream sample *Left:* 6-140 no absorber, *Middle:* 6-140 lithium hydride, *Right:* 6-140 polyethylene wedge, *top:* Downstream cut, *bottom:* TOF2 cut.

Conclusions

A charge exchange experiment was performed at KURNS to investigate electron detachment cross-sections as a function of projectile energy. Using an 11 MeV H^- beam striking a carbon foil, the σ_{-10} electron detachment cross-section (Fig. 2.15) was found to be $1.896_{-0.536}^{+1.859} \times 10^{-17}$ cm²/atom, while the σ_{01} electron detachment cross-section was found to be $0.749_{-0.212}^{+0.734} \times 10^{-17}$ cm²/atom. The limitations of the analysis in this thesis should be kept in mind, however.

The Muon Ionization Cooling Experiment (MICE) investigated ionization cooling by passing a muon beam through several types of absorber materials, including liquid hydrogen and lithium hydride. The systematic uncertainties associated with the liquid hydrogen absorber were investigated, finding three main contributions. The contraction of the absorber vessel and deflection of the absorber vessel windows due to internal pressure reduced the central warm bore length of the MICE absorber vessel by 0.4 ± 0.2 mm. The combined absorber vessel window thickness at the centre of the absorber was found to be 785 ± 13 microns. The temperature during the steady state period of the MICE experiment when the pressure remained constant at 1085 ± 5 mbar was found to be 20.5 ± 0.2 K for each sensor. Combining the eight sensors, an average temperature of 20.51 ± 0.06 K (Fig. 3.8) was obtained during the steady-state period of the experiment, corresponding to a liquid hydrogen density of 70.55 ± 0.07 kg/m³.

The MICE reconstruction software was investigated, finding small biases for the reconstructed transverse momentum and large biases for the reconstructed longitudinal momentum. The biases showed various dependencies including the magnetic field used, the tracker considered, as well as the transverse and longitudinal momentum of the particle. The biases were primarily driven by an assumption of a uniform magnetic field in the tracker regions of the MICE experiment. The modelled

magnetic field however varied by 2-3% from the mean magnetic field (Fig. 5.1).

The two magnetic field mapping studies performed by MICE could only measure the magnetic fields of the individual spectrometers and focus coils, as it was not possible to perform the field mapping study with all of the solenoid coils powered simultaneously. Of the two field mapping studies performed, only the latter was performed with the steel Partial Return Yoke installed. However, for the MICE analyses, the magnetic field map from the first study was used. This means there are significant uncertainties associated with the magnetic field as well as the alignment of the solenoid, tracker and magnetic axes. Simulations were performed varying the magnetic field and the alignment of the solenoid, tracker and magnetic axes to observe how the bias on the MICE reconstruction changed. The results were summarised in Tables 5.6, 5.7 and 5.8.

When there are symmetric transmission losses, the covariance matrix of a beam sample can be separated into the covariance matrix of the transmitted sample and the covariance matrix of the missing sample, normalized by their respective sample sizes (Eq. 7.23). This becomes important when comparing different beam samples as the quantities used to compare them (emittance, amplitude and density), depend on the covariance matrix of the beam sample.

Transmission losses meant the particles measured at the downstream tracker reference plane couldn't be compared directly to the particles at the upstream tracker reference plane. As the missing particles are inaccessible, a transfer matrix approach was used to determine the covariance matrix the particles at the downstream tracker reference plane should be calculated over, so that a direct comparison could be made with particles at the upstream tracker reference plane, even when there were transmission losses i.e. they were calculated over the same phase-space volumes.

Phase-space density plots of Monte Carlo simulations were presented accounting for transmission losses when no absorber material, the lithium hydride absorber or the polyethylene wedge were present in the MICE cooling channel (Fig. 7.20). The curves in the polyethylene wedge plots are due to the dispersion in the beam and need to be corrected for, however such corrections are beyond the scope of this thesis. The no absorber material scenario showed heating of the beam, while the lithium hydride scenario showed cooling for a fraction of the beam. However, for the lithium hydride scenario the longitudinal phase-space volume has also increased, meaning acceptance effects of the MICE cooling channel can't be ruled out. The longitudinal phase-space volume of a beam can be reduced using an RF cavity, though this wasn't possible at MICE. Restoring the longitudinal phase-space volume of a beam would allow for the direct comparison of the transverse phase-space volumes of a beam in the MICE cooling channel when no absorber material was present or when the lithium hydride absorber was present.

“For a moment there, I thought we were in trouble”

— Butch Cassidy

Bibliography

- [1] P. Woit. “Macroscopic: Is String Theory Even Wrong?” *American Scientist* 90.2 (2002), pp. 110–112. URL: <https://www.jstor.org/stable/27857620>.
- [2] C. Brown and C. Rogers. “Systematic uncertainties in the liquid Hydrogen absorber”. *MICE note* 524 (2019), pp. 1–33. URL: <http://mice.iit.edu/micenotes/public/pdf/MICE0524/MICE0524.pdf>. [Accessed 5 October 2021].
- [3] M. Bogomilov et al. “Performance of the MICE diagnostic system”. *Journal of Instrumentation* 16.P08046 (2021), pp. 1–24. DOI: <https://doi.org/10.1088/1748-0221/16/08/p08046>.
- [4] C. Brown. “Monte Carlo simulation of reverse emittance exchange in MICE”. *Neutrino 2020* (2020), p. 1. URL: <https://indico.fnal.gov/event/19348/contributions/186725/>. [Accessed 5 October 2021].
- [5] C. Brown. “Monte Carlo simulation of reverse emittance exchange in MICE”. *ICHEP 2020 WEPTS108* (2020), pp. 1–3. DOI: <https://doi.org/10.22323/1.390.0707>.
- [6] C. Brown and C. Whyte. “Emittance Exchange in MICE”. *IPAC 2019* (2019), pp. 3378–3381. DOI: <https://doi.org/10.18429/jacow-ipac2019-wepts108>.
- [7] MICE collaboration. “First particle-by-particle measurement of emittance in the Muon Ionization Cooling Experiment”. *The European Physical Journal C* 79.257 (2019), pp. 1–15. DOI: <https://doi.org/10.1140/epjc/s10052-019-6674-y>.
- [8] MICE collaboration. “Demonstration of cooling by the Muon Ionization Cooling Experiment”. *Nature* 578 (2020), pp. 53–59. DOI: <https://doi.org/10.1038/s41586-020-1958-9>.
- [9] M. Syphers and F. Zimmermann. “31. Accelerator Physics of Colliders”. *Particle Data Group* (2020), pp. 1–16. URL: <https://pdg.lbl.gov/2021/reviews/rpp2020-rev-accel-phys-colliders.pdf>. [Accessed 1 November 2021].
- [10] F. Zimmermann. “LHC/FCC-based muon colliders”. *Journal of Physics: Conference Series* 1067.022017 (2018), pp. 1–7. DOI: <https://doi.org/10.1088/1742-6596/1067/2/022017>.
- [11] D. Alesini et al. “Positron driven muon source for a muon collider” (2019). arXiv: 1905.05747 [physics.acc-ph].

- [12] Camilla Curatolo and Luca Serafini. “GeV muon beams with picometer-class emittance from electron-photon collisions”. *Applied Sciences* 12.6 (Mar. 2022), p. 3149. DOI: [10.3390/app12063149](https://doi.org/10.3390/app12063149). URL: <https://doi.org/10.3390%2Fapp12063149>.
- [13] Manuela Boscolo, Jean-Pierre Delahaye, and Mark Palmer. “The Future Prospects of Muon Colliders and Neutrino Factories”. In: *Reviews of Accelerator Science and Technology*. World Scientific, Sept. 2019, pp. 189–214. DOI: [10.1142/9789811209604_0010](https://doi.org/10.1142/9789811209604_0010). URL: https://doi.org/10.1142%2F9789811209604_0010.
- [14] Yuri Alexahin et al. “Muon Collider Higgs Factory for Snowmass 2013” (2013). arXiv: 1308.2143 [hep-ph].
- [15] M. Battaglia et al. “31. Energy Frontier Lepton and Photon Colliders”. *Community Planning Study: Snowmass 2013* (2013), pp. 1–35. URL: <https://inspirehep.net/files/9940336a93a576159c673d83fbc90342>. [Accessed 4 November 2021].
- [16] “Enabling Intensity and Energy Frontier Science with a Muon Accelerator Facility in the U.S.: A White Paper Submitted to the 2013 U.S. Community Summer Study of the Division of Particles and Fields of the American Physical Society”. *U.S. Muon Accelerator Program* (2013), pp. 1–56. URL: <https://lss.fnal.gov/archive/2013/conf/fermilab-conf-13-307-apc.pdf>. [Accessed 4 November 2021].
- [17] L. D’Alessi. “The ESSnuSB/HIFI Design Study”. *EPS-HEP 2021* (2021). URL: <https://indico.desy.de/event/28202/contributions/105916/>. [1 November 2021].
- [18] W. Chou, C. Ankenbrandt, and E. Malamud. “The Proton Driver Design Study”. *FERMILAB-TM-2136* (2000). URL: <https://lss.fnal.gov/archive/test-tm/2000/fermilab-tm-2136.pdf>. [Accessed 1 November 2021].
- [19] R. Schmidt. “Introduction to Machine Protection”. *Proceedings of the 2014 Joint International Accelerator School: Beam Loss and Accelerator Protection 2* (2016), pp. 1–20. DOI: <https://doi.org/10.5170/CERN-2016-002.1>.
- [20] D. Faircloth and S. Lawrie. “An overview of negative hydrogen ion sources for accelerators”. *New Journal of Physics* 20.025007 (2018), pp. 1–19. DOI: <https://doi.org/10.1088/1367-2630/aaa39e>.
- [21] A. Borisov, D. Teillet-Billy, and J. Gauyacq. “H⁻ formation by electron capture in hydrogen-Al(111) collisions: perturbative and nonperturbative approaches”. *Surface Science* 278.1-2 (1992), pp. 99–110. DOI: [https://doi.org/10.1016/0039-6028\(92\)90586-U](https://doi.org/10.1016/0039-6028(92)90586-U).
- [22] V. Dudnikov. “Forty years of surface plasma source development”. *Review of Scientific Instruments* 83.02A708 (2012), pp. 1–5. DOI: <https://doi.org/10.1063/1.3670600>.

- [23] R. Johnson and V. Dudnikov. “H- Ion Sources for High Intensity Proton Drivers”. *U.S. Department of Energy Office of Scientific and Technical Information* (2015), pp. 1–72. DOI: <https://doi.org/10.2172/1170411>.
- [24] K. T. McDonald. “The MERIT high-power target experiment at the CERN PS”. *PAC09 TU4GRI03* (2009), pp. 795–799. URL: <https://cds.cern.ch/record/1370100/files/tu4gri03.pdf>. [Accessed 5 November 2021].
- [25] M. et al Catanesi. “Large-angle production of charged pions with 3–12.9 GeV/c incident protons on nuclear targets”. *Physical Review C* 77.055207 (2008), pp. 1–49. DOI: <https://doi.org/10.1103/PhysRevC.77.055207>.
- [26] A. Blondel. “HARP (i)”. *Agenda of the 83rd Meeting of the SPSC* (2007). URL: <https://indico.cern.ch/event/20347/contributions/391115/attachments/307282/429051/Blondel-SPSC-oct-2007-HARP-v2.pdf>. [Accessed 5 November 2021].
- [27] E. Cartlidge. “Particle feud goes public”. *Physics World* (2009). URL: <https://physicsworld.com/a/particle-feud-goes-public/>. [Accessed 5 November 2021].
- [28] European Commission. “Mercury”. *Environment* (2023). URL: https://environment.ec.europa.eu/topics/chemicals/mercury_en. [Accessed 4 November 2023].
- [29] M. Calviani et al. “Considerations on Muon Collider Targetry”. *Workshop on Muon Collider Testing Opportunities* (2021). URL: https://indico.cern.ch/event/1016248/contributions/4282384/attachments/2215324/3752155/MCa__MUC_Targetry__25Mar2021_v1.pdf. [Accessed 5 November 2021].
- [30] O. Caretta et al. “Preliminary experiments on a fluidised powder target”. *EPAC08* (2008), pp. 2862–2864. URL: <https://accelconf.web.cern.ch/e08/papers/wepp161.pdf>. [Accessed 5 November 2021].
- [31] H. J. Cai et al. “New target solution for a muon collider or a muon-decay neutrino beam facility: The granular waterfall target”. *Physical Review Accelerators and Beams* 20.023401 (2017), pp. 1–11. DOI: <https://doi.org/10.1103/PhysRevAccelBeams.20.023401>.
- [32] D. Neuffer, P. Snopok, and Y. Alexahin. “Front End for a neutrino factory or muon collider”. *Journal of Instrumentation* 12.T11007 (2017), pp. 1–13. DOI: <https://doi.org/10.1088/1748-0221/12/11/T11007>.
- [33] M. Appollonio. “Accelerator design concept for future neutrino facilities”. *International scoping study of a future Neutrino Factory and super-beam facility* RAL-TR-2007-23 (2009), pp. 1–76. URL: <https://www.cap.bnl.gov/mumu/project/ISS/ISS-AcceleratorWG-final-submitted-R1.pdf>. [Accessed 5 November 2021].

- [34] M. Palmer. “A Muon Collider Based on a Proton Source”. *Muon Collider - Preparatory Meeting* (2019). URL: https://indico.cern.ch/event/801616/contributions/3358920/attachments/1827644/2993010/ProtonDriverMC_CERN_2019.pdf. [Accessed 5 November 2021].
- [35] A. Kolomenski and A. Lebedev. “The Effect Of Radiation On The Motion Of Relativistic Electrons In A Synchrotron”. *CERN Symposium on High-Energy Accelerators and Pion Physics (HEACC 56)* (1956), pp. 447–455. URL: <https://inspirehep.net/files/aef19e69cd9e29ac43c00c4185f67168>. [Accessed 17 November 2021].
- [36] D. Möhl and A. Sessler. “Beam cooling: Principles and achievements”. *U.S. Department of Energy Office of Scientific and Technical Information* (2003), pp. 1–9. URL: <https://www.osti.gov/biblio/861300>. [Accessed 15 November 2021].
- [37] I. Meshkov. “Beam Cooling Techniques”. *Advanced School on Accelerator Optimization* (2014). URL: <https://indico.cern.ch/event/297045/contributions/1658342/>. [17 November 2021].
- [38] A. Sessler. “The Cooling of Particle Beams”. *U.S. Department of Energy Office of Scientific and Technical Information* (1994), pp. 1–17. URL: <https://www.osti.gov/biblio/10107696-cooling-particle-beams>. [Accessed 17 November 2021].
- [39] G. Budker. “An effective method of damping of particle oscillations in proton and antiproton storage rings (translated)”. *Atomnaya Énergiya* 22.5 (1967), pp. 346–348. URL: https://www.hep.princeton.edu/mumu/physics/budker_sae_22_438_67.pdf. [Accessed 17 November 2021].
- [40] H. Poth. “Electron cooling: Theory, experiment, application”. *Physics Reports* 196.3-4 (1990), pp. 135–297. DOI: [https://doi.org/10.1016/0370-1573\(90\)90040-9](https://doi.org/10.1016/0370-1573(90)90040-9).
- [41] V. Parkhomchuk and N. Skrinskii. “Electron cooling: 35 years of development”. *Physics-Uspekhi* 43.5 (2000), pp. 433–452. DOI: <https://doi.org/10.1070/PU2000v043n05ABEH000741>.
- [42] V. Kamerzhiev. “Introduction to electron cooling”. *Belgian Dutch German graduate school in particle physics 2015* (2015). URL: <https://indico.cern.ch/event/357886/contributions/849347/>. [Accessed 16 November 2021].
- [43] J. R. Hunt. “Beam Quality Characterisation and the Optimisation of Next Generation Antimatter Facilities”. *Doctoral Thesis, University of Liverpool* (2019). DOI: <http://doi.org/10.17638/03039086>.
- [44] D. Möhl et al. “Physics and technique of stochastic cooling”. *Physics Reports* 58.2 (1980), pp. 73–102. DOI: [https://doi.org/10.1016/0370-1573\(80\)90140-4](https://doi.org/10.1016/0370-1573(80)90140-4).
- [45] J. Marriner. “Stochastic cooling overview”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 532.1-2 (2004), pp. 11–18. DOI: <https://doi.org/10.1016/j.nima.2004.06.025>.

- [46] D. Neuffer. “Principles and applications of muon cooling”. *Particle Accelerators* 14 (1983), pp. 75–90. URL: <https://cds.cern.ch/record/142710/files/p75.pdf>. [Accessed 17 November 2021].
- [47] D. Greenwald and A. Caldwell. “Frictional cooling of positively charged particles”. *Physical Review Special Topics Accelerators and Beams* 15.024003 (2012), pp. 1–7. DOI: <https://doi.org/10.1103/PhysRevSTAB.15.024003>.
- [48] H. Abramovicz et al. “A Muon Collider scheme based on Frictional Cooling”. *Nuclear Instruments and Methods in Physics Research A: Accelerators, Spectrometers, Detectors and Associated Equipment* 546.3 (2005), pp. 356–375. DOI: <https://doi.org/10.1016/j.nima.2005.03.125>.
- [49] M. Mühlbauer, H. Daniel, and F. J. Hartmann. “Beam cooling by frictional forces: Monte Carlo calculations and design studies”. *Hyperfine Interactions* 82 (1993), pp. 459–467. DOI: <https://doi.org/10.1007/BF01027980>.
- [50] M. Mühlbauer et al. “Frictional cooling: Experimental results”. *Hyperfine Interactions* 119 (1999), pp. 305–310. DOI: <https://doi.org/10.1023/A:1012624501134>.
- [51] M. Mühlbauer et al. “Frictional cooling: Latest experimental results and first application”. *Nuclear Physics B - Proceedings Supplements* 51.1 (1996), pp. 135–142. DOI: [https://doi.org/10.1016/0920-5632\(96\)00424-0](https://doi.org/10.1016/0920-5632(96)00424-0).
- [52] MICE collaboration. “Repository of images and public information”. *Muon Ionization Cooling Experiment* (2019). URL: https://micewww.pp.rl.ac.uk/projects/mice/wiki/_Repository_of_Images. [Accessed 17 November 2021].
- [53] M. Tanabashi et al. (Particle Data Group). “33. Passage of Particles Through Matter”. *Physical Review D* 98.030001 (2018), pp. 1–36. URL: <https://pdg.lbl.gov/2019/reviews/rpp2018-rev-passage-particles-matter.pdf>. [Accessed 17 November 2021].
- [54] T. Mohayai, C. Rogers, and P. Snopok. “Simulated measurements of cooling in muon ionization cooling experiment”. *IPAC 2016* (2016). URL: <http://accelconf.web.cern.ch/accelconf/ipac2016/papers/tupmy011.pdf>.
- [55] Y. Derbenev and R. Johnson. “Six-dimensional muon beam cooling using a homogeneous absorber: Concepts, beam dynamics, cooling decrements, and equilibrium emittances in a helical dipole channel”. *Physical Review Special Topics Accelerators and Beams* 8.041002 (2005), pp. 1–20. DOI: <https://doi.org/10.1103/PhysRevSTAB.8.041002>.
- [56] Y. Derbenev and R. Johnson. “Parameters for absorber-based reverse emittance exchange of muon beams”. *EPAC* (2006), pp. 1–3. URL: <https://www.osti.gov/biblio/890578-parameters-absorber-based-reverse-emittance-exchange-muon-beams>. [Accessed 16 November 2021].

- [57] J. S. Berg et al. “Acceleration Stages for a Muon Collider”. *PAC99* (1999), pp. 3152–3154. URL: <https://cds.cern.ch/record/552797/files/thp86.pdf>. [Accessed 17 November 2021].
- [58] D. Summers et al. “Acceleration for the $\mu^+\mu^-$ Collider”. *PAC97* (1997), pp. 624–626. URL: <https://accelconf.web.cern.ch/pac97/papers/pdf/3PC004.PDF>. [Accessed 17 November 2021].
- [59] M. Palmer. “Chapter 41: Fast cooling, muon acceleration and the prospect of muon colliders”. *Challenges and Goals for Accelerators in the XXI Century* (2016), pp. 781–798. DOI: https://doi.org/10.1142/9789814436403_0041.
- [60] R. Edgecock et al. “EMMA - The World’s first non-scaling FFA”. *EPAC08* (2008), pp. 3380–3382. URL: <http://accelconf.web.cern.ch/e08/papers/thpp004.pdf>. [Accessed 17 November 2021].
- [61] J. S. Berg. “Minimizing longitudinal distortion in a nearly isochronous linear nonscaling fixed-field alternating gradient accelerator”. *Physical Review Special Topics Accelerators and Beams* 9.034001 (2006), pp. 1–11. DOI: <https://doi.org/10.1103/PhysRevSTAB.9.034001>.
- [62] S. Machida et al. “Acceleration in the linear non-scaling fixed-field alternating-gradient accelerator EMMA”. *Nature Physics* 8 (2012), pp. 243–247. DOI: <https://doi.org/10.1038/nphys2179>.
- [63] S. Brooks. “Vertical orbit excursion fixed field alternating gradient accelerators”. *Physical Review Special Topics Accelerators and Beams* 16.084001 (2013), pp. 1–11. DOI: <https://doi.org/10.1103/PhysRevSTAB.16.084001>.
- [64] S. Machida et al. “Optics design of vertical excursion fixed-field alternating gradient accelerators”. *Physical Review Accelerators and Beams* 24.021601 (2021), pp. 1–11. DOI: <https://doi.org/10.1103/PhysRevAccelBeams.24.021601>.
- [65] S. Machida. “FFA Options for ISIS Upgrade and the Feasibility Study”. *Proceedings of the 3rd J-PARC Symposium (J-PARC2019)* 33.011005 (2021), pp. 1–8. DOI: <https://doi.org/10.7566/JPSCP.33.011005>.
- [66] D. Schulte. “The Muon Collider”. *Muon Collider KEK--PH lectures and workshops* (2021). URL: <https://indico.cern.ch/event/1056654/>. [Accessed 17 November 2021].
- [67] M. Zisman. “Technical Challenges and Scientific Payoffs of Muon Beam Accelerators for Particle Physics”. *The 20th Biennial Conference on Magnet Technology, Philadelphia, PA, August 27-31, 2007* (2007), pp. 1–10. URL: <https://www.osti.gov/biblio/932527-technical-challenges-scientific-payoffs-muon-beamaccelerators-particle-physics>. [Accessed 18 November 2021].

- [68] M. Bogomilov et al. “The MICE Muon Beam on ISIS and the beam-line instrumentation of the Muon Ionization Cooling Experiment”. *Journal of Instrumentation* 7.P05009 (2012), pp. 1–37. DOI: <https://doi.org/10.1088/1748-0221/7/05/P05009>.
- [69] M. Bogomilov (on behalf of the MICE Collaboration). “The MICE Particle Identification System”. *Nuclear Physics B - Proceedings Supplements* 215.1 (2011), pp. 316–318. DOI: <https://doi.org/10.1016/j.nuclphysbps.2011.04.042>.
- [70] R. Bertoni et al. “The design and commissioning of the MICE upstream time-of-flight system”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 615.1 (2010), pp. 14–26. DOI: <https://doi.org/10.1016/j.nima.2009.12.065>.
- [71] D. Lietti et al. “The prototype of the MICE Electron–Muon Ranger: Design, construction and test”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 604.1–2 (2009), pp. 314–318. DOI: <https://doi.org/10.1016/j.nima.2009.01.105>.
- [72] M. Ellis et al. “The design, construction and performance of the MICE scintillating fibre trackers”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1 (2011), pp. 136–153. DOI: <https://doi.org/10.1016/j.nima.2011.04.041>.
- [73] R. Asfandiyarov et al. “MAUS: the MICE analysis user software”. *Journal of Instrumentation* 14.T04005 (2019), pp. 1–21. DOI: <https://doi.org/10.1088/1748-0221/14/04/T04005>.
- [74] F. Drielsma. “Measurement of the increase in phase space density of a muon beam through ionization cooling”. *Doctoral Thesis, Université de Genève* (2018). DOI: <https://doi.org/10.13097/archive-ouverte/unige:114100>.
- [75] C. Hunt. “High precision track reconstruction and first emittance measurements in the MICE step IV cooling channel”. *Doctoral Thesis, Imperial College London* (2016). DOI: <https://doi.org/10.25560/44974>.
- [76] J. Langlands. “Measuring and Modelling the Magnetic Field of the MICE Spectrometer Solenoids”. *Doctoral Thesis, University of Sheffield* (2018). URL: <https://etheses.whiterose.ac.uk/23398/>. [Accessed 19 November 2021].
- [77] S. Middleton. “Characterisation of the MICE experiment”. *Doctoral Thesis, Imperial College London* (2018). DOI: <https://doi.org/10.25560/63880>.
- [78] C. Rogers. “Beam Dynamics in an Ionisation Cooling Channel”. *Doctoral Thesis, Imperial College London* (2007). URL: http://mice.iit.edu/phd/rogers_thesis.pdf. [Accessed 5 May 2022].
- [79] M. Calkin. *Lagrangian and Hamiltonian mechanics*. Singapore: World Scientific, 1996.

- [80] T. Helgaker, P. Jørgensen, and J. Olsen. “The electronic Hamiltonian in an electromagnetic field” (2014). URL: <https://trygvehelgaker.no/Presentations/Hamiltonian.pdf>. [Accessed 6 May 2022].
- [81] S. Malham. “An introduction to Lagrangian and Hamiltonian mechanics” (2016). URL: <http://www.macs.hw.ac.uk/~simonm/mechanics.pdf>. [Accessed 6 May 2022].
- [82] M. Morii. “Mechanics 151 Lecture 18” (2018). URL: <http://users.physics.harvard.edu/~morii/phys151/lectures/Lecture18.pdf>. [Accessed 26 February 2018].
- [83] G. Stupakov. “Lecture notes on Classical Mechanics and Electromagnetism in Accelerator Physics”. *The US Particle Accelerator School* (2011). URL: <https://jseldredphysics.files.wordpress.com/2018/03/stupakov-notes-2011.pdf>. [Accessed 6 May 2022].
- [84] D. Tong. “4. The Hamiltonian Formalism” (2004). URL: <http://www.damtp.cam.ac.uk/user/tong/dynamics/four.pdf>. [Accessed 6 May 2022].
- [85] G. Vilasi. *Hamiltonian Mechanics*. Singapore: World Scientific, 2001.
- [86] S. Y. Lee. *Accelerator Physics*. Hackensack, New Jersey, USA: World Scientific, 2004.
- [87] M. Reiser. “Theory and Design of Charged Particle Beams” (2008), pp. 243–248. URL: http://www-eng.lbl.gov/~dleitner/USPAS_2016_Fundamental_Of_Ion_Sources/Reference%20Material/Theory_and_Design_of_Charged_Pa.pdf. [Accessed 24 January 2022].
- [88] A. Dragt. “Surface Methods for the Computation of Realistic Symplectic Transfer Maps from Numerical Field Data on a Grid” (2017). URL: http://helper.ipam.ucla.edu/publications/bd2017/bd2017_13799.pdf.
- [89] A. Dobbs et al. “The reconstruction software for the MICE scintillating fibre trackers”. *Journal of Instrumentation* 11.T12001 (2016), pp. 1–17. DOI: <https://doi.org/10.1088/1748-0221/11/12/T12001>.
- [90] J. Stedeford and J. Hasted. “Further Investigations of Charge Exchange and Electron Detachment. I. Ion Energies 3 to 40 keV. II. Ion Energies 100 to 4000 eV”. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* 227.1171 (1955), pp. 466–486. DOI: <https://doi.org/10.1098/rspa.1955.0024>.
- [91] A. Wishart. “The bound-free photodetachment cross section of H^- ”. *Journal of Physics B: Atomic and Molecular Physics* 12 (1979), pp. 3511–3519. DOI: <https://doi.org/10.1088/0022-3700/12/21/009>.
- [92] S. Smith and D. Burch. “Relative Measurement of the Photodetachment Cross Section for H^- ”. *Physical Review* 116 (1959), pp. 1125–1131. DOI: <https://doi.org/10.1103/PhysRev.116.1125>.

- [93] B. Folsom et al. “Stripping mechanism and remediation for H^- beams”. *Physical Review Accelerators and Beams* 24.074201 (2021), pp. 1–9. DOI: <https://doi.org/10.1103/PhysRevAccelBeams.24.074201>.
- [94] L. Scherk. “An improved value for the electron affinity of the negative hydrogen ion”. *Canadian Journal of Physics* 57.4 (1979), pp. 558–563. DOI: <https://doi.org/10.1139/p79-077>.
- [95] J. Donahue et al. “Photodetachment of Relativistic Ions”. *IEEE Transactions on Nuclear Science* 28.2 (1981), pp. 1203–1206. DOI: <https://doi.org/10.1109/TNS.1981.4331379>.
- [96] V. Danilov et al. “Proof-of-principle demonstration of high efficiency laser-assisted H^- beam conversion to protons”. *Physical Review Special Topics Accelerators and Beams* 10.053501 (2007), pp. 1–6. DOI: <https://doi.org/10.1103/PhysRevSTAB.10.053501>.
- [97] S. Cousineau et al. “First Demonstration of Laser-Assisted Charge Exchange for Microsecond Duration H^- Beams”. *Physical Review Letters* 118.074801 (2017), pp. 1–5. DOI: <https://doi.org/10.1103/PhysRevLett.118.074801>.
- [98] A. Shishlo et al. “First Observation of Intrabeam Stripping of Negative Hydrogen in a Superconducting Linear Accelerator”. *Physical Review Letters* 108.114801 (2012), pp. 1–4. DOI: <https://doi.org/10.1103/PhysRevLett.108.114801>.
- [99] M. Chanel et al. “Measurements of H^- intra-beam stripping cross section by observing a stored beam in LEAR”. *Physics Letters B* 192.3-4 (1987), pp. 475–477. DOI: [https://doi.org/10.1016/0370-2693\(87\)90141-9](https://doi.org/10.1016/0370-2693(87)90141-9).
- [100] G. Gillespie. “Double closure calculation of the electron-loss cross section for H- in high-energy collisions with H and He”. *Physical Review A* 15.2 (1977), pp. 563–573. DOI: <https://doi.org/10.1103/PhysRevA.15.563>.
- [101] G. Gillespie. “Excitation and ionization contributions to sum-rule Born cross sections for collisions of one-electron ions with atoms”. *Physical Review A* 18.5 (1978), pp. 1967–1989. DOI: <https://doi.org/10.1103/PhysRevA.18.1967>.
- [102] G. Gillespie. “High-energy cross sections for H- ions incident on intermediate and high-Z atoms”. *Physical Review A* 16.3 (1977), pp. 943–950. DOI: <https://doi.org/10.1103/PhysRevA.16.943>.
- [103] L. Alvarez. “Energy Doubling in dc Accelerators”. *Review of Scientific Instruments* 22 (1951), pp. 705–706. DOI: <https://doi.org/10.1063/1.1746050>.
- [104] A. Mohagheghi et al. “Interaction of relativistic H- ions with thin foils”. *Physical Review A* 43.3 (1991), pp. 1345–1365. DOI: <https://doi.org/10.1103/physreva.43.1345>.

- [105] R. Webber and C. Hojvat. “Measurement of the Electron Loss Cross Sections for Negative Hydrogen Ions on Carbon at 200 MeV”. *IEEE Transactions on Nuclear Science* 26.3 (1979), pp. 4012–4014. DOI: <https://doi.org/10.1109/TNS.1979.4330681>.
- [106] H. Funsten, D. McComas, and B. Barraclough. “Thickness uniformity and pinhole density analysis of thin carbon foils using incident keV ions”. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 66.4 (1992), pp. 470–478. DOI: [https://doi.org/10.1016/0168-583X\(92\)95421-M](https://doi.org/10.1016/0168-583X(92)95421-M).
- [107] F. Allegrini, R. Ebert, and H. Funsten. “Carbon foils for space plasma instrumentation”. *Journal of Geophysical Research: Space Physics* 121.5 (2016), pp. 3391–3350. DOI: <https://doi.org/10.1002/2016JA022570>.
- [108] M. Holl et al. “Preparation of self-supporting carbon foils of uniform thickness below $2 \mu\text{g}/\text{cm}^2$ density”. *Journal of Applied Physics* 45.7 (1974), pp. 3069–3071. DOI: <https://doi.org/10.1063/1.1663725>.
- [109] M. Plum et al. “Stripper foil failure modes and cures at the Oak Ridge Spallation Neutron Source”. *Physical Review Special Topics Accelerators and Beams* 14.030102 (2011), pp. 1–6. DOI: <https://doi.org/10.1103/PhysRevSTAB.14.030102>.
- [110] C. Liaw et al. “Calculation of the maximum temperature on the Carbon Stripping Foil of the Spallation Neutron Source”. *Proceedings of the 1999 Particle Accelerator Conference, New York, 1999* (1999), pp. 3300–3302. URL: <https://accelconf.web.cern.ch/p99/PAPERS/THP143.PDF>. [Accessed 4 February 2022].
- [111] C. Liaw, Y. Lee, and J. Tuozzolo. “Life Time of Carbon Stripping Foils for the Spallation Neutron Source”. *Proceedings of the 2001 Particle Accelerator Conference, Chicago* (2001), pp. 1538–1540. URL: <https://accelconf.web.cern.ch/p01/PAPERS/TPAH138.PDF>. [Accessed 4 February 2022].
- [112] A. Drozhdin et al. “Modeling multiturn stripping injection and foil heating for high intensity proton drivers”. *Physical Review Special Topics Accelerators and Beams* 15.011002 (2012), pp. 1–10. DOI: <https://doi.org/10.1103/PhysRevSTAB.15.011002>.
- [113] S. G. Lebedev and A. S. Lebedev. “Calculation of the lifetimes of thin stripper targets under bombardment of intense pulsed ions”. *Physical Review Special Topics Accelerators and Beams* 11.020401 (2008), pp. 1–7. DOI: <https://doi.org/10.1103/PhysRevSTAB.11.020401>.
- [114] W. Chou, M. Kostin, and Z. Tang. “Stripping efficiency and lifetime of carbon foils”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 590.1-3 (2008), pp. 1–12. DOI: <https://doi.org/10.1016/j.nima.2008.02.060>.

- [115] T. Gorlov et al. “Sequential excitation scheme for laser stripping for a H^- beam”. *Physical Review Accelerators and Beams* 22.121601 (2019), pp. 1–8. DOI: <https://doi.org/10.1103/PhysRevAccelBeams.22.121601>.
- [116] J. P. Carneiro, B. Mustapha, and P. N. Ostroumov. “Numerical simulations of stripping effects in high-intensity hydrogen ion linacs”. *Physical Review Special Topics Accelerators and Beams* 12.040102 (2009), pp. 1–12. DOI: <https://doi.org/10.1103/PhysRevSTAB.12.040102>.
- [117] B. Goddard et al. “Stripping foil issues for H^- injection into the CERN PSB at 160 MeV”. *Proceedings of IPAC’10, Kyoto, Japan* (2010), pp. 3951–3953. URL: <https://accelconf.web.cern.ch/IPAC10/papers/thpeb030.pdf>. [Accessed 1 February 2022].
- [118] Y. Chae and Y. Cho. “Study of field ionization in the charge exchange injection for the IPNS Upgrade”. *Proceedings Particle Accelerator Conference 95* (1995), pp. 3412–3414. DOI: <https://doi.org/10.1109/PAC.1995.505899>.
- [119] D. Faircloth et al. “The front end test stand high performance H^- ion source at Rutherford Appleton Laboratory”. *Review of Scientific Instruments* 81.02A721 (2010), pp. 1–5. DOI: <https://doi.org/10.1063/1.3271169>.
- [120] P. Saha et al. “Measurement of 181 MeV H- ions stripping cross-sections by carbon stripper foil”. *Nuclear Instruments and Methods in Physics Research A* 776 (2015), pp. 87–93. DOI: <https://doi.org/10.1016/j.nima.2014.12.068>.
- [121] J. Heinemeier, P. Hvelplund, and F. Simpson. “Collisional detachment cross sections for H- and He- at high energies”. *Journal of Physics B: Atomic and Molecular Physics* 9.15 (1976), pp. 2669–2684. URL: <https://iopscience.iop.org/article/10.1088/0022-3700/9/15/017>. [Accessed 4 August 2021].
- [122] S. Allison. “Experimental results on charge changing collisions of Hydrogen and Helium atoms and ions at kinetic energies above 0.2 keV”. *Reviews of Modern Physics* 30.4 (1958), pp. 1137–1168. DOI: <https://doi.org/10.1103/RevModPhys.30.1137>.
- [123] I. Kovacs. “One and two electron loss cross sections of H- ions in CO₂ or N₂ gas”. *Nuclear Instruments and Methods* 51.2 (1967), pp. 224–228. DOI: [https://doi.org/10.1016/0029-554X\(67\)90005-5](https://doi.org/10.1016/0029-554X(67)90005-5).
- [124] I. Dmitriev, Y. Teplova, and Y. Fainberg. “An experimental study of single electron loss cross sections for hydrogen atoms and negative hydrogen ions in various media”. *Journal of Experimental and Theoretical Physics* 80.1 (1995), pp. 28–31. URL: http://www.jetp.ac.ru/cgi-bin/dn/e%5C_080%5C_01%5C_0028.pdf. [Accessed 14 January 2019].
- [125] R. Smythe and J. Toevs. “Collisional Electron Detachment from Hydrogen Atoms and Negative Hydrogen Ions between 4 and 18 MeV”. *Physical Review* 139 (1965), A15–A18. DOI: <https://doi.org/10.1103/PhysRev.139.A15>.

- [126] E. Acerbi et al. “Collisional electron detachment from hydrogen atoms between 22 and 45 MeV”. *Lettere Al Nuovo Cimento Series 2* 10.14 (1974), pp. 598–602. DOI: <https://doi.org/10.1007/BF02782633>.
- [127] L. Welsh et al. “Cross Sections for Electron Capture by Fast Protons in H_2 , He, N_2 , and Ar”. *Physical Review* 158 (1967), pp. 85–92. DOI: <https://doi.org/10.1103/PhysRev.158.85>.
- [128] A. Weiss. “Hartree-Fock multiplet strengths for K I, Ca II, and Sc III”. *Journal of Research of the National Bureau of Standards Section A: Physics and Chemistry* 71A.2 (1967), p. 157. DOI: <http://dx.doi.org/10.6028/jres.071A.022>. Accessed 4 August 2021.
- [129] K. Banyard and G. Taylor. “Exposition of atomic pair correlation functions and associated energies obtained by analysis of correlated wave functions”. *Physical Review A* 10.6 (1974), pp. 1972–1980. DOI: <https://doi.org/10.1103/PhysRevA.10.1972>.
- [130] M. Inokuti. “Inelastic Collisions of Fast Charged Particles with Atoms and Molecules — The Bethe Theory Revisited”. *Reviews of Modern Physics* 43.3 (1971), pp. 297–347. DOI: <https://doi.org/10.1103/RevModPhys.43.297>.
- [131] P. Kürpick et al. “Excited-state subshell population of hydrogen atoms after transmission of relativistic H^- ions through thin foils”. *Physical Review A* 58 (1998), pp. 2183–2190. DOI: <https://doi.org/10.1103/PhysRevA.58.2183>.
- [132] M. Gulley et al. “Measurement of H-, H0, and H+ yields produced by foil stripping of 800-MeV H- ions”. *Physical Review A* 53.5 (1996), pp. 3201–3210. DOI: <https://doi.org/10.1103/PhysRevA.53.3201>.
- [133] Arizona Carbon Foil Company. “Arc Evaporated Carbon Foils - The Workhorse!” (2022). URL: <https://www.acf-metals.com/arc-evaporated-carbon-the-workhorse>. [Accessed 22 February 2022].
- [134] FFA Collaboration. “Files” (2019). URL: <https://sites.google.com/site/kurrifag/file-cabinet>. [Accessed 21 February 2022].
- [135] I. Newton. *De analysi per æquationes numero terminorum infinitas*. 1669.
- [136] I. Newton and (edited by) W. Jones. *Analysis per Quantitatum Series, Fluxiones, ac Differentias: cum Enumeratione Linearum Tertii Ordinis*. London, UK: Pearson, 1711.
- [137] J. Raphson. *Analysis Æquationum Universalis, seu Ad Æquationes Algebraicas Resolvendas Methodus Generalis, & Expedita, Ex nova Infinitarum Serierum Methodo, Deducta ac Demonstrata*. London, UK: Abel Swall, 1690.
- [138] T. Simpson. *Essays on several Curious and Useful Subjects In Speculative and mix'd Mathematicks*. London, UK: John Nourse, 1740.

- [139] M. Green and S. Yang. “Does One Know the Properties of a MICE Solid or Liquid Absorber to Better than 0.3 Percent?”. *MICE Note* 155 (2006), pp. 1–18. URL: <http://mice.iit.edu/micenotes/public/pdf/MICE0155/MICE0155.pdf>. [Accessed 2 October 2018].
- [140] V. Bayliss et al. “The liquid-hydrogen absorber for MICE”. *Journal of Instrumentation* 13 (2018), T09008, 1–26. DOI: <https://doi.org/10.1088/1748-0221/13/09/t09008>. [Accessed 3 October 2018].
- [141] S. Ishimoto et al. “Liquid Hyrdogen absorber for MICE”. *IPAC 2010* (2010), pp. 421–423. URL: <http://accelconf.web.cern.ch/AccelConf/IPAC10/papers/mopeb065.pdf>. [Accessed 2 October 2018].
- [142] G. Hardin. *NIST, Aluminum 6061-T6 (UNS AA96061)*. 2018. URL: <https://www.nist.gov/mml/acmd/aluminum-6061-t6-uns-aa96061>. [Accessed 3 October 2018].
- [143] MICE collaboration. *Index of /analysis/processvariables/v0/201702*. 2018. URL: https://micewww.pp.rl.ac.uk/analysis/process_variables/v0/. [Accessed 15 March 2022].
- [144] Lakeshore. *Cernox RTDs*. 2018. URL: https://www.lakeshore.com/docs/default-source/product-downloads/catalog/lstc_cernox_1.pdf?sfvrsn=41b96c23_6. [Accessed 15 March 2022].
- [145] S. Courts. “Cernox TM Resistance Temperature Sensors for High Energy Physics Applications” (2018). URL: https://www.lakeshore.com/docs/default-source/product-downloads/application-notes/cernoxforhighenergy.pdf?sfvrsn=756b8bbf_4. [Accessed 15 March 2022].
- [146] Lakeshore. “Temperature Measurement and Control Catalog” (2018). URL: https://www.lakeshore.com/docs/default-source/product-downloads/lakeshore_tc_1.pdf. [Accessed 15 March 2022].
- [147] B. Brandt, D. Liu, and L. Rubin. “Low temperature thermometry in high magnetic fields. VII. CernoxTM sensors to 32 T”. *Review of Scientific Instruments* 70.1 (1999), pp. 104–110. DOI: <https://doi.org/10.1063/1.1149549>.
- [148] G. Heine and W. Lang. “Magnetoresistance of the new ceramic “Cernox” thermometer from 4.2K to 300K in magnetic fields up to 13T”. *Cryogenics* 38.4 (1998), pp. 377–379. DOI: [https://doi.org/10.1016/S0011-2275\(97\)00130-6](https://doi.org/10.1016/S0011-2275(97)00130-6).
- [149] S. Ishimoto and S. Suzuki. *Absorber R and D*. 2006. URL: http://mice.iit.edu/cm/cm16/cm16_agenda.html. [Accessed 20 July 2019].
- [150] G. Buntkowsky et al. “Mechanism of nuclear spin initiated para-H₂ to ortho-H₂ conversion”. *Physical Chemistry Chemical Physics* 8.16 (2006), pp. 1929–1035. DOI: <https://doi.org/10.1039/B601594H>.

- [151] D. Weitzel et al. “Ortho-Para Catalysis In Liquid-Hydrogen Production”. *Journal of Research of the National Bureau of Standards* 60.3 (1958), pp. 221–227. DOI: <http://dx.doi.org/10.6028/jres.060.026>. [Accessed 2 October 2018].
- [152] J. Hust and R. Stewart. “A compilation of the property differences of Ortho and Para Hydrogen or mixtures of Ortho and Para Hydrogen”. *National Bureau of Standards Report 8812* (1965), pp. 1–39. URL: <https://nvlpubs.nist.gov/nistpubs/Legacy/RPT/nbsreport8812.pdf>. [Accessed 15 March 2022].
- [153] Particle Data Group. “Atomic and nuclear properties of Aluminium” (2018). URL: http://pdg.lbl.gov/2017/AtomicNuclearProperties/HTML/aluminum_Al.html. [Accessed 2 October 2018].
- [154] Particle Data Group. “Atomic and nuclear properties of liquid Hydrogen” (2018). URL: http://pdg.lbl.gov/2017/AtomicNuclearProperties/HTML/liquid_hydrogen.html. [Accessed 2 October 2018].
- [155] Particle Data Group. “Muon in Aluminium” (2018). URL: http://pdg.lbl.gov/2017/AtomicNuclearProperties/MUE/muE_aluminum_Al.pdf. [Accessed 2 October 2018].
- [156] Particle Data Group. “Muons in liquid Hydrogen” (2018). URL: http://pdg.lbl.gov/2017/AtomicNuclearProperties/MUE/muE_liquid_hydrogen.pdf. [Accessed 2 October 2018].
- [157] S. Virostek. “MICE Spectrometer Solenoid Design and Assembly”. *MICE Spectrometer Solenoid Magnets Review* (2015). URL: https://indico.fnal.gov/event/10651/contributions/4590/attachments/3180/3828/MICE_SS_Des_Assy.pdf. [Accessed 8 June 2022].
- [158] J. Tarrant. “Step IV Engineering”. *MICE Collaboration Meeting 46* (2016). URL: https://indico.cern.ch/event/563953/contributions/2284360/attachments/1349237/2036849/037_06_October_-_PRY_Stand-Offs.pdf. [Accessed 21 May 2022].
- [159] H. Witte and S. Plate. “Partial Return Yoke for MICE”. *Technical Report BNL-100819-2013-IR* (2013). DOI: <https://doi.org/10.2172/1108553>.
- [160] V. Bayliss et al. “Mitigation of stray magnetic fields in MICE”. *MICE Note 419* (2013). URL: <http://mice.iit.edu/mnp/MICE0419.pdf>. [Accessed 30 May 2022].
- [161] H. Witte et al. “Partial return yoke for MICE step IV and final step”. *IPAC2015 WEPJE027* (2015), pp. 2732–2734. URL: <https://lss.fnal.gov/archive/2015/conf/fermilab-conf-15-476-apc.pdf>. [Accessed 21 May 2022].
- [162] H. Witte et al. “Analysis of the Training Behaviour of the MICE Spectrometer Solenoid”. *MT25 - 25th International Conference on Magnet Technology THU-AF-PO4* (2017), pp. 1–5. URL: <https://www.bnl.gov/isd/documents/95569.pdf>. [Accessed 20 May 2022].

- [163] H. Witte et al. “Analysis of the Training Behaviour of the MICE Spectrometer Solenoid (Poster)”. *MT25 - 25th International Conference on Magnet Technology* (2017). URL: <https://indico.cern.ch/event/445667/contributions/2563130/attachments/1512069/2358354/Thu-Af-Po4Poster.pdf>. [Accessed 20 May 2022].
- [164] C. Rogers. “Emittance Evolution”. *MICE Collaboration Meeting 50* (2018). URL: <https://indico.cern.ch/event/678517/>. [Accessed 30 May 2022].
- [165] C. T. Rogers. “First Observation of Normalised Emittance Reduction through Ionization Cooling”. *DRAFT - MICE Internal* (2019), pp. 1–113. URL: https://micewww.pp.rl.ac.uk/attachments/download/10879/MICE_Note.pdf. [Accessed 23 May 2022].
- [166] C. Rogers. “Emittance Evolution”. *MICE Analysis Workshop - Glasgow* (2019). URL: <https://indico.cern.ch/event/803896/>. [Accessed 30 May 2022].
- [167] C. Rogers. “Emittance Evolution”. *MICE Collaboration Meeting 53* (2019). URL: <https://indico.cern.ch/event/786717/>. [Accessed 30 May 2022].
- [168] Paolo Franchini. “Field measurements of the MICE beamline dipoles”. *MICE Note 520* (2018), pp. 1–5. URL: <http://mice.iit.edu/mnp/MICE0520.pdf>. [Accessed 27 May 2022].
- [169] P. Franchini. “Dipole Model”. *MICE Collaboration Meeting 51* (2018). URL: <https://indico.cern.ch/event/727201/>. [Accessed 30 May 2022].
- [170] F. Drielsma. “Beam-based detector alignment in the MICE Muon Beam line”. *MICE Note 500* (2016), pp. 1–22. URL: <http://mice.iit.edu/mnp/MICE0500.pdf>. [Accessed 8 June 2022].
- [171] T. Millington. “MICE Hall Solenoid1, Solenoid2 & Focus Coil - Installed Positions”. *MICE Note 469* (2015), pp. 1–3. URL: <http://mice.iit.edu/mnp/MICE0469.pdf>. [Accessed 8 June 2022].
- [172] T. Millington. “MICE Hall Solenoid1 & Solenoid2 Current Positions 31/10/2016”. *MICE Note 501* (2016), pp. 1–3. URL: <http://mice.iit.edu/mnp/MICE0501.pdf>. [Accessed 8 June 2022].
- [173] T. Millington. “MICE Focus Coil Re-positioning Survey Report”. *MICE Note 509* (2017), pp. 1–3. URL: <http://mice.iit.edu/mnp/MICE0509.pdf>. [Accessed 8 June 2022].
- [174] T. Millington. “Survey Report”. *MICE Note 510* (2017), pp. 1–7. URL: <http://mice.iit.edu/mnp/MICE0510.pdf>. [Accessed 8 June 2022].
- [175] T. Millington. “MICE Focus Coil Re-positioning Survey Report”. *MICE Note 514* (2017), pp. 1–5. URL: <http://mice.iit.edu/mnp/MICE0514.pdf>. [Accessed 8 June 2022].

- [176] T. Millington. “MICE Pre-run Solenoid Positioning Survey Report”. *MICE Note 515* (2017), pp. 1–7. URL: <http://mice.iit.edu/mnp/MICE0515.pdf>. [Accessed 8 June 2022].
- [177] T. Millington. “MICE Pre-run Solenoid Positioning Survey Report”. *MICE Note 517* (2017), pp. 1–8. URL: <http://mice.iit.edu/mnp/MICE0517.pdf>. [Accessed 8 June 2022].
- [178] T. Millington. “MICE Pre-run Solenoid Positioning Survey Report”. *MICE Note 518* (2017), pp. 1–8. URL: <http://mice.iit.edu/mnp/MICE0518.pdf>. [Accessed 8 June 2022].
- [179] T. Millington and A. Hooper. “MICE Hall TOF 1”. *MICE Note 513* (2017), pp. 1–13. URL: <http://mice.iit.edu/mnp/MICE0513.pdf>. [Accessed 8 June 2022].
- [180] T. Millington and A. Hooper. “MICE Hall TOF 2”. *MICE Note 494* (2016), pp. 1–8. URL: <http://mice.iit.edu/mnp/MICE0494.pdf>. [Accessed 8 June 2022].
- [181] F. Drielsma. “Beam-based alignment corrections for the 2015/04 ISIS user cycle”. *MICE Note 502* (2016), p. 1. URL: <http://mice.iit.edu/mnp/MICE0502.pdf>. [Accessed 8 June 2022].
- [182] F. Drielsma. “Beam-based alignment corrections for the 2016/02 ISIS user cycle”. *MICE Note 503* (2016), pp. 1–2. URL: <http://mice.iit.edu/mnp/MICE0503.pdf>. [Accessed 8 June 2022].
- [183] F. Drielsma. “Beam-based alignment corrections for the 2016/03 ISIS user cycle”. *MICE Note 504* (2016), pp. 1–2. URL: <http://mice.iit.edu/mnp/MICE0504.pdf>. [Accessed 8 June 2022].
- [184] F. Drielsma. “Beam-based alignment corrections for the 2016/04 ISIS user cycle”. *MICE Note 505* (2016), pp. 1–2. URL: <http://mice.iit.edu/mnp/MICE0505.pdf>. [Accessed 8 June 2022].
- [185] F. Drielsma. “Beam-based alignment corrections for the 2017/01 ISIS user cycle”. *MICE Note 512* (2017), pp. 1–2. URL: <http://mice.iit.edu/mnp/MICE0512.pdf>. [Accessed 8 June 2022].
- [186] F. Drielsma. “Beam-based alignment corrections for the 2017/02 ISIS user cycle”. *MICE Note 516* (2017), p. 1. URL: <http://mice.iit.edu/mnp/MICE0516.pdf>. [Accessed 8 June 2022].
- [187] F. Drielsma. “Beam-based alignment corrections for the 2017/03 ISIS user cycle”. *MICE Note 519* (2017), p. 1. URL: <http://mice.iit.edu/mnp/MICE0519.pdf>. [Accessed 8 June 2022].
- [188] V. Blackmore. “Field Mapping: Magnetic Axis of FC2, SSU & SSD”. *MICE Collaboration Meeting 41* (2015). URL: https://indico.cern.ch/event/360388/contributions/1773899/attachments/717944/985502/09Feb15_MagneticAxis.pdf. [Accessed 8 June 2022].

-
- [189] V. Blackmore and J. Cobb. “Magnetic axes of the Spectrometer Solenoids”. *MICE Note 481* (2016), pp. 1–4. URL: <http://mice.iit.edu/mnp/MICE0481.pdf>. [Accessed 8 June 2022].
- [190] J. Cobb and V. Blackmore. “Magnet Mapping”. *MICE Collaboration Meeting 42* (2015). URL: <https://indico.cern.ch/event/374187/contributions/888672/attachments/745674/1022924/Mapping-CM42s.pdf>. [Accessed 21 May 2022].
- [191] J. Cobb. “The Alignment of the MICE Magnet Modules”. *MICE Note 445* (2014), pp. 1–26. URL: <http://mice.iit.edu/mnp/MICE0445.pdf>. [Accessed 8 June 2022].
- [192] D. Maletic. “Monte Carlo Production” (2021). URL: <https://micewww.pp.rl.ac.uk/projects/analysis/wiki/MCProduction>. [Accessed 26 May 2022].
- [193] R. Klie. “Reaching a new resolution standard with electron microscopy”. *Physics* 2.85 (2009). DOI: <https://doi.org/10.1103/Physics.2.85>.
- [194] J. Bullinaria. “Bias and Variance, Under-Fitting and Over-Fitting”. *Neural Computation: Lecture 9* (2015). URL: <https://www.cs.bham.ac.uk/~jxb/INC/19.pdf>. [Accessed 23 June 2022].
- [195] J. Branson. “Phase Space and Liouville’s Theorem” (2013). URL: https://hepweb.ucsd.edu/ph110b/110b_notes/node93.html. [Accessed 24 July 2022].
- [196] M. Dvorak. “Liouville’s Theorem”. *PHGN 505 Report* (2011), pp. 1–12. URL: https://inside.mines.edu/~tohno/teaching/PH505_2011/liouville_dvorak.pdf. [Accessed 24 July 2022].
- [197] J. Liouville. “Note Sur la Théorie de la Variation des constantes arbitraires”. *Journal de mathématiques pures et appliquées 1^{re} série* 3 (1838), pp. 342–349. URL: http://portail.mathdoc.fr/JMPA/PDF/JMPA_1838_1_3_A26_0.pdf. [Accessed 24 July 2022].
- [198] M. Fowler. “11. Introduction to Liouville’s Theorem” (2009). URL: https://galileoandstein.phys.virginia.edu/7010/CM_11_Introduction_to_Liouville.html. [Accessed 24 July 2022].
- [199] H. Wiedemann. *Particle Accelerator Physics, 3rd edition*. Berlin Heidelberg, Germany: Springer, 2007.
- [200] T. Mohayai. “A Novel Non-Parametric Density Estimation Approach to Measuring Muon Ionization Cooling and Reverse Emittance Exchange in the MICE Experiment”. *Doctoral Thesis, Illinois Institute of Technology* (2018). DOI: <https://doi.org/10.2172/1515052>.
- [201] V. A. Epanechnikov. “Non-Parametric Estimation of a Multivariate Probability Density”. *Theory of Probability & Its Applications* 14.1 (1969), pp. 153–158. DOI: [10.1137/1114019](https://doi.org/10.1137/1114019).

- [202] A. Baraldi and C. Enders. “An introduction to modern missing data analyses”. *Journal of School Psychology* 48.1 (2010), pp. 5–37. DOI: <https://doi.org/10.1016/j.jsp.2009.10.001>.
- [203] J. Peugh and C. Enders. “Missing Data in Educational Research: A Review of Reporting Practices and Suggestions for Improvement”. *Review of Educational Research* 74.4 (2004), pp. 525–556. DOI: <https://doi.org/10.3102/00346543074004525>.
- [204] D. Rubin. “Inference and missing data”. *Biometrika* 63.3 (1976), pp. 581–592. DOI: <https://doi.org/10.1093/biomet/63.3.581>.
- [205] J. Nissen, R. Donatello, and B. Van Busen. “Missing data and bias in physics education research: A case for using multiple imputation”. *Physical Review Physics Education Research* 15.020106 (2019), pp. 1–15. DOI: <https://doi.org/10.1103/PhysRevPhysEducRes.15.020106>.
- [206] M. Baneshi and A. Talei. “Multiple Imputation in Survival Models: Applied on Breast Cancer Data”. *Iranian Red Crescent Medical Journal* 13.8 (2011), pp. 544–549. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3371994/>. [Accessed 25 July 2022].
- [207] Q. Luo et al. “Validity of using multiple imputation for "unknown" stage at diagnosis in population-based cancer registry data”. *PLOS ONE* 12.6 (2017), pp. 1–16. DOI: <https://doi.org/10.1371/journal.pone.0180033>.
- [208] N. Eisemann, A. Waldmann, and A. Katalinic. “Imputation of missing values of tumour stage in population-based cancer registration”. *BMC Medical Research Methodology* 11.129 (2011), pp. 1–13. DOI: <https://doi.org/10.1186/1471-2288-11-129>.
- [209] R. Stannard et al. “Obtaining long-term stage-specific relative survival estimates in the presence of incomplete historical stage information”. *British Journal of Cancer* (2022), pp. 1–8. DOI: <https://doi.org/10.1038/s41416-022-01866-8>.
- [210] C. Brown. “Emittance Exchange in MICE”. *MICE Collaboration Meeting 55* (2019). URL: <https://indico.cern.ch/event/841655/>. Accessed 29 July 2022.