Beam Dynamics in an Ionisation Cooling Channel

Chris Rogers

Imperial College, London

A thesis submitted for the degree of
Doctor of Philosophy
of the University of London
and the Diploma of Imperial College.

Abstract

The Neutrino Factory has been proposed as a facility to provide an intense source of neutrinos suitable for the measurement of neutrino oscillation parameters and a possible CP violating phase to unprecedented precision. In the Neutrino Factory, neutrinos are produced by the decay of a muon beam with 20-50 GeV per muon. Initially, the muon beam occupies a large volume in phase space, which must be reduced before the beam can be accelerated. The proposed method to achieve this is to use a solenoidal ionisation cooling channel.

In this thesis the physics of the cooling channel is described including the cooling effects of different materials and the beam optics of solenoidal focussing and radio frequency cavities. The Muon Ionisation Cooling Experiment is described, in which a section of a long cooling channel will be constructed at the Rutherford Appleton Laboratory (RAL) in Oxfordshire, the cooling performance measured and compared with computer simulation. The simulation of the cooling channel in the G4MICE code is detailed and the accuracy of such simulations is investigated. The beam dynamics of the simulated cooling channel is examined and the cooling performance of the channel is studied. The cooling measurement is simulated and the simulated measurement is compared with the simulated true cooling performance.

Contents

List of Tables

List of Figures

Acknowledgements

I am indebted to many people for their help and advice with this thesis. Thanks are due to the following especially for useful discussion and support:
Ken Long
John Cobb
Malcolm Ellis
Pip Crosby

Chapter 1

Introduction

Three major facilities have been proposed for the precision study of neutrino oscillation parameters, the Neutrino Factory, the Betabeam and the Superbeam. Of these, the Neutrino Factory offers the finest precision measurement of oscillation parameters [1]. The Neutrino Factory generates neutrinos from an intense muon beam that has been accelerated to high energy.

Muons are generated by firing a multi-megawatt proton beam onto a target to generate pions. The pions decay to a highly dispersed muon beam. The phase space density of the muon beam must be increased in order that enough muons can be accelerated within a reasonable accelerator acceptance, such that the neutrino yield is sufficient to give good statistics in the Neutrino Factory detector systems. The only technique that can increase the muon beam density on the timescale of the muon lifetime is ionisation cooling.

1.1 Neutrino Oscillations

Neutrino oscillations were discovered at Super-Kamiokande in 1998 by examining the angular distribution of muon neutrinos generated as decay products of atmospheric hadronic showers. A deficit of upward travelling muon neutrinos was observed without a corresponding decrease in upward travelling electron neutrinos, as compared with downward travelling neutrinos. Later analysis indicated a surplus of tau decay products, indicating mixing between muon and tau neutrinos [2]. Strong support was given by observations at the Sudbury Neutrino Observatory (SNO) published in 2001 and 2002 [3]. SNO found a deficit of electron neutrinos produced by solar nuclear fusion as compared with the number predicted by the Standard Solar Model. Later, KamLAND observed oscillations from neutrinos produced by nuclear reactors [4]. In addition, CNGS, MINOS and K2K have yielded observations of neutrino oscillations from neutrino beams

[5] [6] [8].

Current observations can be explained by the existence of three neutrino mass eigenstates that mix to form the weak eigenstates. A neutrino is created and annihilated in the weak interaction as a weak eigenstate, but travels through space as a mass eigenstate. The mass eigenstates interfere with each other, possibly leading to an oscillation into another weak eigenstate. In the case of Dirac neutrinos the oscillations are parametrised by three angles, θ_{13}, θ_{23}, θ_{12}, the neutrino mass squared differences $\Delta m^2_{ij} = m^2_i - m^2_j$ and one phase, δ [10]. There is particular interest in the measurement of the phase δ, as non-zero values would imply CP violation in the leptonic sector.

1.2 Measurement of Neutrino Oscillations

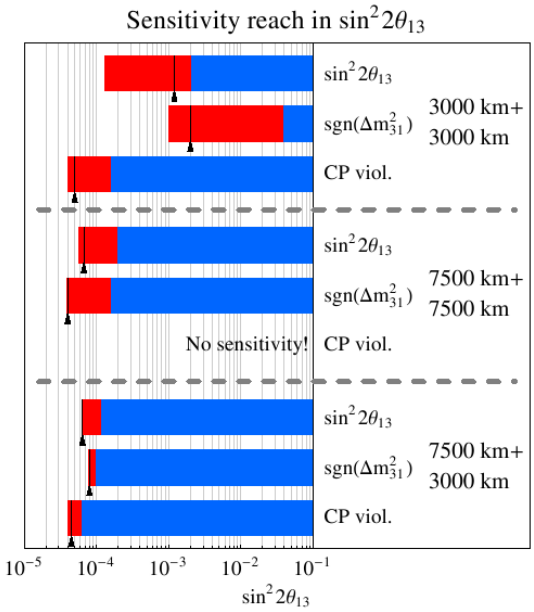

Figure 1.1: Sensitivity of the Neutrino Factory to various parameters as a function of $\sin^2 2\theta_{13}$ for different choices of accelerator-detector distance [1].

The Neutrino Factory has been shown to provide the most accurate measurements in much of parameter space compared to other next-generation neutrino beam facilities as it provides an intense, high energy and well characterised beam. The presence of muon antineutrinos and electron neutrinos for μ^- decay and muon neutrinos and electron antineutrinos for μ^+ decay allows independent experiments on both flavours and both matter and antimatter. In particular, the Neutrino Factory has been shown to offer at least an order of magnitude greater sensitivity to the value of $\sin\theta_{13}$, sensitivity to the sign of Δm^2_{13} and more sensi-

tivity to δ for order-of-magnitude smaller values of $\sin\theta_{13}$ as compared with other neutrino facilities. The sensitivity of the Neutrino Factory to Δm^2_{13}, θ_{13} and δ is displayed in Figure 1.1. In addition, the Neutrino Factory offers a programme of muon physics and non-oscillation neutrino physics that cannot be performed at other facilities.

Measurement of the oscillation parameters is performed by firing the neutrino beam through the earth into a number of detectors. Different accelerator-detector distances have been proposed that give differing sensitivity to the neutrino oscillation parameters. It is expected that two different distances will be used with two different detectors.

1.3 The Role of Ionisation Cooling

The baseline specification of the Neutrino Factory in the International Scoping Study assumes of order 10^{21} muon decays per year. Higher fluxes will enable more accurate measurements of the Neutrino Oscillation parameters. Even with large acceptance accelerators such a flux is challenging to achieve as the muon beam occupies a large phase space volume. The phase space volume of the beam arises due to the intrinsic phase space density of the beam that occurs in pion production and decay.

An alternative and more effective method than building larger accelerators is to increase the density of the muon beam using a technique known as ionisation cooling. Here, muons are passed through some material so that the total momentum of the beam is reduced. This momentum is returned to the muons in RF cavities, but only in the longitudinal direction. Thus the overall transverse momentum of the muon beam is reduced resulting in a higher density, or cooler, beam. The technique has been shown in simulation to be quite effective. However, it has never been observed in reality.

Chapter 2

The Need For Cooling

Ionisation cooling is an important part of a Neutrino Factory design as it enables a much larger proportion of the muon beam to be accelerated, resulting in a more intense neutrino beam. There are three conceptual designs for a Neutrino Factory, presented by European [11], US [1], [12], [13] and Japanese [14] groups. The Japanese design presents and option with no ionisation cooling, instead relying on very high acceptance accelerators. At the time of writing, the four independent design studies are being combined into the International Scoping Study (ISS) [15]. The ISS design calls for only a modest amount of cooling, but this is quite dependent on the acceptance that can be achieved in the accelerators. Recent work indicates that the construction of high acceptance accelerators may be more challenging than previously thought [16].

The ISS Neutrino Factory baseline design is shown in Figure 2.1. The challenge in Neutrino Factory design is to provide a beam of sufficient intensity. All designs call for a source of protons that is as intense as possible. Beam powers of up to 4 MW have been proposed, 20 times as powerful as the highest power sources currently operational. Protons are fired at energies of order 10 GeV onto a target where pions are produced. The intensity of the beam is sufficient to destroy conventional targets; moving solid or liquid targets are under investigation. Pions are captured either in a 20 T solenoid which is adiabatically weakened along the accelerator axis; or a magnetic horn where a toroidal field with strength proportional to $1/r$ is used to make a paraxial beam. Pions decay to muons with a large energy spread, which is converted into a time spread using phase rotation. Under the ISS scheme the pions are subsequently bunched and the beam is cooled using ionisation cooling to increase the number of muons that are accepted by the accelerating system. The beam is then accelerated as fast as possible, to minimise decay losses, using a combination of recirculating linacs (RLAs) and high acceptance Fixed Field Alternating Gradient machines (FFAGs). Finally

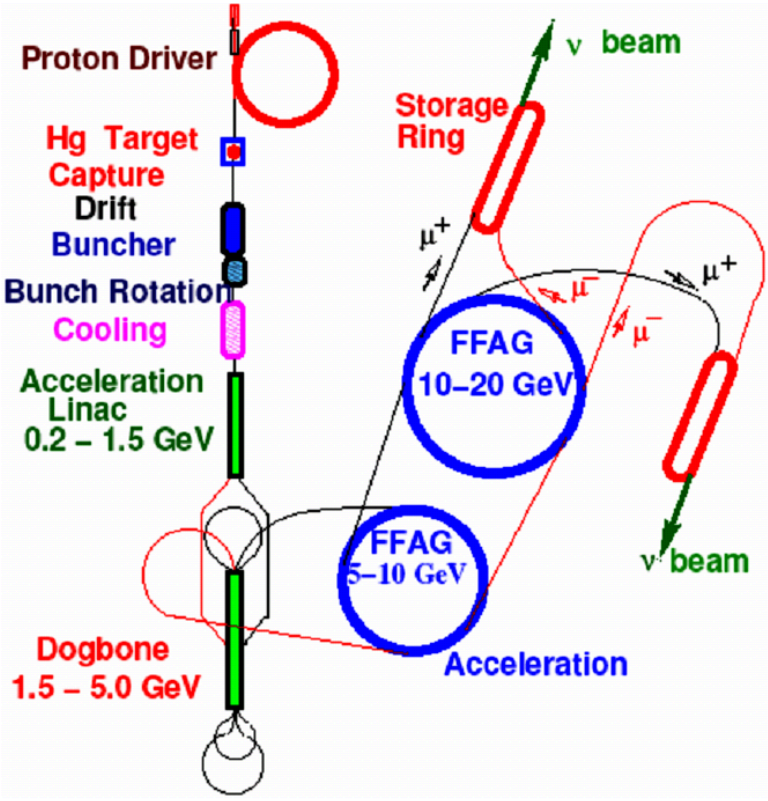

Figure 2.1: A schematic of the ISS baseline Neutrino Factory.[15].

the muons are stored in a large aperture triangular or racetrack-shaped storage ring where the muons decay to neutrinos.

The baseline muon energy in the storage ring is 50 GeV. A 20 GeV machine may be an alternative first stage followed by an upgrade to 50 GeV. The baseline specification of the machine is for 10^{21} muon decays per year aimed at one or two detectors. The challenges involved in achieving these energies and intensities are revealed by examining the technologies required in more detail.

2.1 Proton Driver

In order to maximise the number of protons that can be brought onto the target, several designs for an intense proton source has been proposed operating with beam power of up to 4 MW [18]. This should be compared to beam powers of order 0.1 MW characteristic of the most intense proton sources currently available and beam powers of order 1 MW that may be available in the next generation of proton sources. Accelerator energies of order 10-30 GeV give the best pion production yields, although proton energies of 5-10 GeV tend to give better muon yields as the pions produced are more easily captured [19]. Short bunches,

typically around 1-3 ns, are used in order to provide a structure to the muon beam that will fit into the relatively high frequency, high gradient RF buckets necessary for fast muon acceleration.

The various proton driver designs are quite different as they are usually based around an existing proton machine. Common to all designs is a high intensity H^- ion source followed by acceleration by RFQs and fast beam chopping. Subsequently, acceleration is performed by linacs, Rapid Cycling Synchrotrons (RCSs) or FFAGs depending on the particular design.

2.2 Target Station

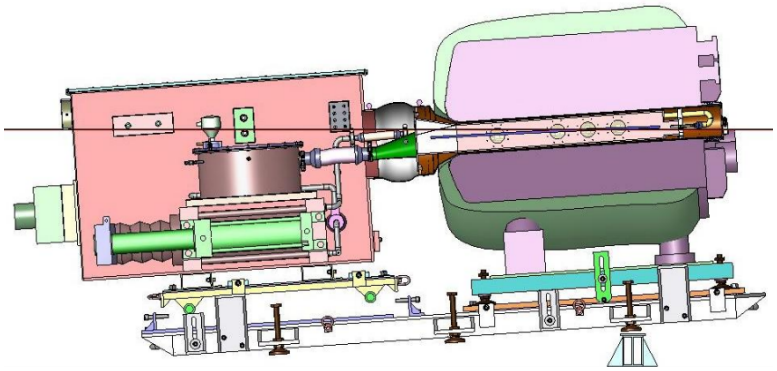

Figure 2.2: A schematic of the MERIT experiment. The mercury syringe sits in the left hand container and fires a mercury jet into the bore of the pulsed 15 T solenoid on the right hand side. Protons are fired on a horizontal path from the right edge [20].

The high beam current, fast repetition rate and short, tightly focussed bunch place great strains on the proton target. Simulations indicate that a fixed target would be destroyed within a few pulses. Several alternative solutions have been proposed involving either moving solid or liquid targets.. The preferred option is to fire a mercury jet along the beamline.

The challenges are considerable for such a scheme. Mercury is highly corrosive and toxic and will be activated by the intense radiation. The jet cannot be contained by material as thermal expansion of mercury in the proton beam is expected to cause such a pipe to explode, while cavitation may erode the pipe. The mercury jet may be focused in the intense field of the target station so that the jet will remain intact. However, if the mercury jet is destroyed mercury will be sprayed inside the solenoid bore, possibly introducing a material that will absorb pions produced from the target.

In the MERIT proof-of-principle experiment a mercury jet will be fired across a 24 GeV/c proton beam with maximum intensity of $3x10^{1}3$ protons per bunch, focussed to a spot size of 1-2 mm by a 15 T field. After passing through the beam-line, the mercury pools on the bottom of the target station and is recirculated. The MERIT experiment [20] is shown in Figure 2.2.

2.3 Phase Rotation

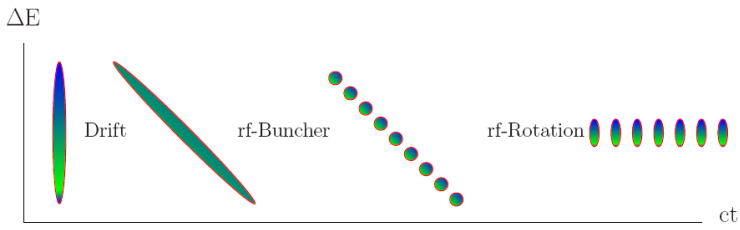

Figure 2.3: Manipulations in the longitudinal phase space of the muon front end. Muons drift in a macro-bunch and then are collected into micro-bunches. The energy of the bunches is aligned to a central momentum using Vernier phase rotationi [1].

After the target the strong solenoidal fields are gradually tapered to a few Tesla. In this distance, the majority of pions decay to muons. Under the ISS scheme the muons are adiabatically bunched into so-called *micro-bunches* with a spread of 0.5 ns RMS before being rotated in momentum-time phase space using a system known as Vernier Rotation [21]. The resulting momentum-time phase space is shown in Figure 2.4. The shorter bunch length enables a higher frequency RF voltage to be used. In general higher frequency RF systems are less prone to breakdown so that higher RF voltages may be achieved, enabling faster acceleration.

First the macro-bunch is allowed to drift for 90 m, so that an energy-time correlation arises. Then high frequency RF is applied to the bunch. The RF voltage is adiabatically increased over a distance of 60 m from 0 to 4.8 MV/m, giving rise to micro-bunches. As the micro-bunches continue to drift in the macro-bunch for this distance, the RF frequency must be altered so that the centre of each micro-bunch stays at a constant RF phase and the final RF frequency corresponds to the final beam. In the ISS, the RF frequency is initially 333 MHz and falls to 200 MHz.

Finally, the energy of the tails of the macro-bunch are manipulated so that

they are at the same energy as the centre. This is achieved by introducing a slight variation in the RF frequency so that the phase of the bunch centre varies across the macro bunch. Muons that are late see a positive phase, giving them an energy boost, while muons that are early see a negative phase reducing their energy. The RF frequency is then modulated to maintain this phase offset until the micro-bunches all have the same energy, where the appropriate frequency is calculated by firing reference particles at the central momentum and time of the off-momentum bunches so that a higher muon flux is transferred into the muon cooling channel.

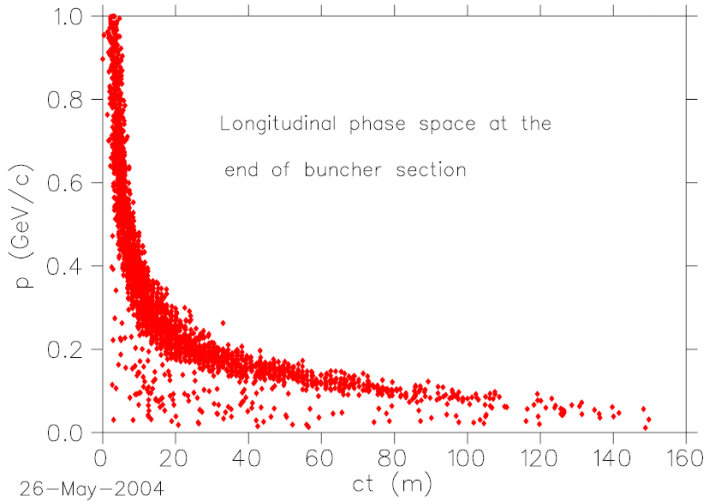

Figure 2.4: The energy-time relationship after the buncher section. Micro-bunches can be seen faintly [1].

2.4 Ionisation Cooling

After phase rotation the beam emittance is reduced such that more muons are accepted into the acceleration system and storage ring using ionisation cooling. The baseline ISS ionisation cooling scheme uses a linear cooling channel based on Lithium Hydride absorbers with a thin Beryllium coating to ease handling issues. A cartoon of the ISS scheme is shown in Figure 2.5.

The main challenge in the construction of such a cooling channel is expected to arise due to the enhanced breakdown probability in the RF cavities. 200 MHz RF is placed in intense solenoidal fields produced by superconducting coils. The RF must be run at high peak fields in order to minimise the overall length of the cooling channel to reduce decay losses. However, evidence suggests that in

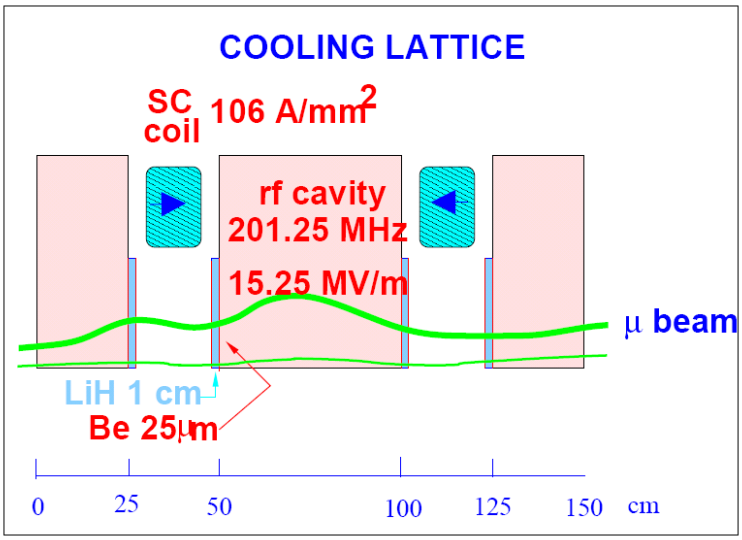

Figure 2.5: The ISS baseline cooling channel. Emittance is reduced in Lithium Hydride absorbers. The beam is focussed by superconducting solenoids and longitudinal momentum is returned by 200 MHz RF cavities [1].

strong magnetic fields breakdown occurs at lower peak electric fields [22]. This may necessitate a longer cooling channel with lower voltage cavities.

The Muon Ionisation Cooling Experiment (MICE) collaboration will construct a cell of a cooling channel [23] as a Proof-of-Principle demonstration of ionisation cooling techniques. The collaboration seeks to demonstrate cooling over a range of beam momenta, beam emittances and absorber materials. MICE is examined in more detail in Chapter 5

2.5 Acceleration

Conventional synchrotrons are unable to increase their bending magnet field quickly enough to compete with decay losses, making acceleration too slow. Several alternate acceleration schemes have been envisaged. Very Rapid Cycling Synchrotrons (VRCSs) [24], Recirculating Linacs (RLAs) [25] and Fixed Field Alternating Gradient Accelerators (FFAGs) [26] have been proposed as alternatives. The range of transverse acceptances for these machines varies between 15 and 45 mm rad, although specific values depend on the machine design.

2.5.1 VRCSs

VRCSs use very fast ramping magnets to achieve a rapid cycling time, necessitating a careful magnet design. The gaps between dipoles are minimised to reduce

the stored energy in the ring, for example using combined function magnets and using longer bending arcs and straights than in conventional synchrotrons. Eddy currents can be minimised by design of the magnets and by using a relatively low duty cycle. However, the synchrotron would have a relatively low acceptance and so would require greater beam cooling.

2.5.2 RLAs

In RLAs, one or two linacs provide acceleration of the muons for RLAs shaped in so-called 'dogbones' or 'racetracks'. Magnet arcs recirculate muons back into the linacs such that the RF cavities can be used several times on each cycle. Unlike synchrotrons, several different arcs are provided with different fixed dipole field strengths for muons of different momenta so that there is no issue with the magnet rise time. A high aperture dipole guides muons of different momenta into and out of the different arcs. In order to operate RLAs the longitudinal momentum spread must be much smaller than the beam energy to allow clean separation of beams before they enter the arcs, placing a limit on the minimum injection energy.

2.5.3 FFAGs

FFAGs have fixed field bending magnets where the magnet field is much higher at larger radii. Hence higher momentum particles see a greater bending field and can be contained within the same bending magnet as lower momentum particles, even over a relatively large range of momenta. Thus there is no need to ramp.

Several different types of FFAGs exist. In scaling FFAGs the focusing fields scale as a power law with radius. This means that the beam optics are the same at higher momenta so that the beam is not affected by resonances. However, magnets are expensive and technically challenging to design and build. Also, the flight time of the muons changes as the muons travel at higher momenta. It is not possible to modulate the frequency of the cavities fast enough, limiting the energy over which the muons can be accelerated.

In non-scaling FFAGs the fields do not scale with radius. This means muons at different energies travel through different path lengths, and with a careful choice of beam optics, the muons can be held in the rf bucket for longer, allowing greater acceleration. However, these machines have a smaller acceptance and the muon beam crosses resonances that may cause emittance growth.

2.6 Storage Ring

The final component of the Neutrino Factory is the muon storage ring. Here neutrinos are stored and decay in long straights aimed at the various detectors. The storage ring must be designed not only so that the neutrino rate at the detectors is as high as possible, but also so that the rate is well known.

The number of neutrinos that strike the detectors is optimised by ensuring that muons travel towards the detectors for as long as possible. It is only muons that are travelling towards the detector as they decay that contribute to the neutrino beam. Two basic designs have been proposed for a storage ring; based on an oval-shaped racetrack scheme for a single far detector on one baseline; and an isosceles triangle scheme for two far detectors on two baselines. The isosceles triangle scheme requires two vertically stacked rings of opposite polarity to enable neutrinos from opposite-sign muons to be fired at the same detector.

The neutrino rate at the detectors is dependent on the neutrino beam divergence as well as the number of muons in the beam. There are two contributing factors to the neutrino beam divergence; the intrinsic divergence from muon decay and the divergence of the muon beam before decay. The storage ring is designed to ensure that the muon beam divergence is much smaller than the intrinsic neutrino beam divergence, which requires soft focusing and large aperture magnets.

2.7 Cooling-Accelerator Payoff

The advantages of cooling can be offset by the use of a larger acceptance accelerator and storage ring or detector systems with a larger fiducial volume. The cost of a long cooling channel was compared with the cost of larger acceptance accelerators by Bob Palmer [27] and the result is shown in Figure 2.6.

The cost of cooling has been estimated assuming a linear cooling channel with Lithium Hydride absorbers. This cooling channel may be impractical due to the significant magnetic fields in the RF cavity volume, in which case a longer and more expensive cooling channel would be required. However, the cooling performance may be significantly improved through the use of an emittance exchange scheme.

For a scheme with less cooling, high acceptance FFAGs will be required. Recent simulation results indicate that longitudinal heating due to transverse-longitudinal correlations may make these schemes impractical [16]. There is also a need to understand the accelerator and cooling costs better, as well as more

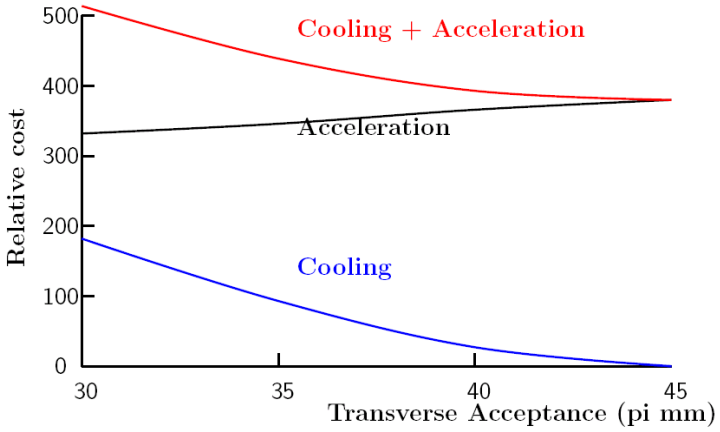

Figure 2.6: The cost payoff between higher acceptance acceleration and more extensive cooling [27].

well developed storage ring and detector designs. Further investigations are in progress.

Chapter 3

Solenoidal Beam Optics

In an accelerating structure it is necessary to ensure that particle trajectories are stable so that the beam is contained. In ionisation cooling channels this is usually achieved in transverse phase space using solenoids and in longitudinal phase space using RF cavities.

In vacuum, beam optics can be determined using a transfer map formalism, whereby operators are determined that transport particle trajectories along the beamline. Under a first order approximation in a stable lattice, particle trajectories can be decomposed into conserved amplitudes and angles, one for each decoupled phase space. The mean of the amplitudes for a beam are also conserved and are called the emittances. Instabilities are known as resonances and the conditions for linear instabilities are found to be completely dependent on the transfer map.

Terms beyond first order are also of interest in cooling channels due to the high emittance of the beams involved. However, these non-linear terms are beyond the scope of this thesis and as such will only be discussed qualitatively.

3.1 Motion in a Solenoidal Focussing System

Most cooling channels rely on solenoidal focussing systems. Typically solenoids have larger acceptances and greater focussing can be achieved due to the symmetry between x and y axes, which increases the amount of cooling achieved. Also, the focussing is stronger in solenoids than quadrupoles at the low momentum typical of cooling channels.

There are several routes to derive the beam optical equations for a solenoid. Penn [28] derives the equations of motion and uses these to derive Hill's equation. Wang and Kim [29] extend this formalism by taking a Fourier decomposition of the focusing function to reach the same beam optical equations as well as

resonance phenomena of the lattice. Berz [30] outlines a method to derive transfer maps to high precision using a Taylor expansion and this has been applied to cooling channels [31]. Dragt [32] extends this formalism to use a Lie algebraic approach that enables a faster numerical integration by choosing the order that terms in the Taylor expansion are summed with care.

In this thesis, transfer maps will be derived using a Taylor expansion of the Hamiltonian under standard perturbation theory [33]. This technical approach has the advantage that transfer maps are symplectic, as described below. This condition is required if phase space volume is to be conserved and leads to emittance conservation in a rigorous manner. It also makes a number of useful mathematical tools available and is readily extensible to non-linear terms.

First the Hamiltonian for a solenoid is derived. Subsequently the transfer map operator, that transfers particles between z and $z + dz$, is derived to first order. This is shown to be symplectic and emittance conserving in this first-order approximation. Particle transport is decomposed into a lattice-dependent β function and a particle-dependent angle. Finally stability criteria are derived for the beam.

3.1.1 Coordinate System and notation

Vectors are indicated with an arrow above a letter, as in $\vec{u}$ while matrices are indicated with bold font. $\vec{u}$ will be used to denote a general phase space vector while $\vec{q}$, $\vec{A}$ and $\vec{p}$ will be used to denote position, four vector potential and momentum respectively. Subscripts such as p_x indicate, for example, momentum along $\vec{x}$. Subscripts i, j and k will be used to indicate an arbitrary axis.

Right-handed cartesian coordinates will be used with $\vec{z}$ along the beam (solenoid) axis, $\vec{x}$ in the horizontal directoin and $\vec{y}$ pointing in the vertical direction where relevant. (r, θ) denote cylindrical polar coordinates about $\vec{z}$. t denotes time. In this chapter, $\vec{p}^{\mathrm{c}}$ denotes the canonical conjugate momentum to the position coordinates and $\vec{p}^{\mathrm{k}}$ the kinetic equivalent. Upper case will be used to denote lab coordinates while lower case will be used to denote deviation coordinates, as explained below. In subsequent chapters, lower case will be used throughout and in general only kinetic coordinates will be considered.

3.1.2 Hamiltonian expansion for a Solenoid

The aim here is to find the Hamiltonian for a charged particle in a solenoidal field and expand it as a power series. This can then be used to find operators, known as transfer maps, that transport particles down the beam line.

The Hamiltonian for a particle of charge q and mass m moving in a general electromagnetic field with canonical phase space vector $\vec{u}_T^c = (T, P_T^c, X, P_X^c, Y, P_Y^c)$, using z as the independent variable and taking the speed of light $c = 1$, is [34]

$$H_z(\vec{u}_T; z) = qA_z - \Big[(P_T^c - qA_T)^2 - m^2 - (P_X^c - qA_X)^2 - (P_Y^c - qA_Y)^2\Big]^{1/2} \quad (3.1)$$

Here the canonical momentum P_i^c is related to the kinetic momentum P_i^k by

$$P_i^k = P_i^c - qA_i, \quad (3.2)$$

P_T^k is the particle energy, T is the absolute time and A_T is the scalar potential. The 4-vector potential from an axially symmetric system of currents is given by [35]

$$\begin{pmatrix} A_T \\ A_X \\ A_Y \\ A_z \end{pmatrix} = \begin{pmatrix} 0 \\ -Yf(z, r^2) \\ Xf(z, r^2) \\ 0 \end{pmatrix}, \quad (3.3)$$

where

$$f(z, r^2) = \Sigma \frac{(-1)^n (r^2)^n}{2^{2n+1} n!(n+1)!} \frac{\partial^{2n} B_z(0, 0, z)}{\partial z^{2n}} \quad (3.4)$$

and B_z is the magnetic field in the $\vec{z}$ direction.

Then the Hamiltonian becomes

$$H_z(\vec{u}_T^c; z) = -\Big[(P_T^c)^2 - m^2 - (P_X^c - qA_X)^2 - (P_y^c - qA_Y)^2\Big]^{1/2}. \quad (3.5)$$

In order to perform a perturbation expansion of the system it is necessary to transform the Hamiltonian to deviation variables

$$\begin{aligned} t &= T - T_0, & (3.6)\\ x &= X, & (3.7)\\ y &= Y, & (3.8)\\ p_t^c &= P_T^c - p_0^c, & (3.9)\\ p_x^c &= P_X^c, & (3.10)\\ p_y^c &= P_Y^c. & (3.11) \end{aligned}$$

Here T_0 and p_0^c are the time and momentum coordinates of the reference particle that moves with speed v_0 where $T_0 = z/v_0$. In choosing a generating function for the transformation, it is necessary to ensure that the variables of the generating function are independent. Hence a generating function of the second kind, F_2, is chosen as $(T, X, Y; p_t^c, p_x^c, p_y^c; z)$ are independent.

Generating functions of the second kind obey the relations [36]

$$P_i = \frac{\partial F_2}{\partial Q_i}, \tag{3.12}$$

$$q_i = \frac{\partial F_2}{\partial p_i}, \tag{3.13}$$

which gives

$$F_2 = X p_x^c + Y p_y^c + (T - T_0)(p_t^c + p_0^c). \tag{3.14}$$

for the transformation between lab and deviation coordinates. Then the Hamiltonian with phase space coordinates $\vec{u}^c = (t, p_t^c, x, p_x^c, y, p_y^c)$ is [36]

$$\begin{aligned} K_z(\vec{u}^c; z) &= H_z(\vec{u}_T^c; z) + \frac{\partial F_2}{\partial z} \\ &= H_z(\vec{u}_T^c; z) - \frac{p_t^c + p_0^c}{v_0} \\ &= -[(p_t^c + p_0^c)^2 - m^2 - (p_x^c - qA_x)^2 - (p_y^c - qA_y)^2]^{1/2} \\ &\quad - \frac{p_t^c + p_0^c}{v_0}. \end{aligned} \tag{3.15}$$

The Hamiltonian is next expanded as a power series,

$$K_z = K_0 + K_1 + K_2 + \ldots, \tag{3.16}$$

where K_n is a homogeneous polynomial in phase space coordinates $u_{i_1} u_{i_2} \ldots u_{i_n}$ of degree n. In the linear approximation, only terms 0 through 2 are required. Using a binomial expansion on (3.15) and the expression for the vector potential (3.3) gives

$$K_0 = \frac{p_0}{v_0}(1 + v_0^2) \tag{3.17}$$

$$K_1 = 0 \tag{3.18}$$

$$\begin{aligned} K_2 &= \frac{(p_x^c)^2 + (p_y^c)^2}{2p} - \frac{B_0}{2p}(x p_y^c - y p_x^c) + \frac{B_0^2}{8p}(x^2 + y^2) \\ &\quad + \frac{(p_t^c)^2}{2p\beta_{rel}^2\gamma_{rel}^2} \end{aligned} \tag{3.19}$$

where $B_0 = qB_z(0, 0, z)$ and p, β_{rel} and γ_{rel} are the momentum and relativistic β and γ of the reference particle.

K_2 is then used to calculate the first order transfer map for a solenoid.

3.1.3 First order transfer map for a solenoid

Define a transfer map $\mathbf{M}$ as an operator that transports single particle coordinates between two planes at z_{in} and z_{fin} on the beamline such that

$$\mathbf{M}\vec{u}^c(z_{in}) = \vec{u}^c(z_{fin}). \tag{3.20}$$

The equation of motion is

$$\frac{d\vec{u}}{dz} =: \vec{u} : K = [\vec{u}, K] = -[K, \vec{u}]. \tag{3.21}$$

Here the Poisson bracket and :: operators are defined by

$$: f : g = [f, g] = \frac{\partial f}{\partial q_i}\frac{\partial g}{\partial p_i} - \frac{\partial f}{\partial p_i}\frac{\partial g}{\partial q_i}, \tag{3.22}$$

and the transfer map to transfer particles from z to $z + dz$ is defined by

$$\mathbf{M}(dz) = \mathbf{1} - : K : dz = \mathbf{1} - (: K_0 : + : K_1 : + : K_2 : + \ldots)dz. \tag{3.23}$$

Then $\mathbf{M_i}$ is defined as the transfer map where the Hamiltonian is truncated to a polynomial of i^{th} order. For example, $\mathbf{M_2}$ is given by

$$\mathbf{M_2}(dz) = \mathbf{1} - (: K_0 : + : K_1 : + : K_2 :)dz = \mathbf{1} - (\frac{\partial K_2}{\partial q_i}\frac{\partial}{\partial p_i} - \frac{\partial K_2}{\partial p_i}\frac{\partial}{\partial q_i})dz \tag{3.24}$$

In this case $\mathbf{M_i}\vec{u}$ is a linear combination of i^{th} order polynomials. $\mathbf{M_2}$ can be written as a matrix,

$$\mathbf{M_2}(dz) = \mathbf{1} + \begin{pmatrix} 0 & -1/(p\beta_{rel}^2\gamma_{rel}^2) & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/p & B_0/2p & 0 \\ 0 & 0 & -B_0^2/4p & 0 & 0 & B_0/2p \\ 0 & 0 & -B_0/2p & 0 & 0 & 1/p \\ 0 & 0 & 0 & -B_0/2p & -B_0^2/4p & 0 \end{pmatrix} dz. \tag{3.25}$$

It can be seen that x and y phase spaces are coupled and to first order longitudinal and transverse phase spaces are independent. As the first non-unitary transfer map is $\mathbf{M_2}$ this is referred to as the first order approximation.Under certain conditions higher order polynomials in the Hamiltonian become dominant and this approximation is no longer valid. These are known as resonances and are discussed in more details below.

Numerical Integration of $\mathbf{M_2}$

It will be useful to understand how $\mathbf{M_2}$ can be integrated. The transfer map at $z + dz$ is related to the transfer map at z by

$$\mathbf{M_2}(\mathbf{z} + \mathbf{dz}) = \mathbf{M_2}(dz)\mathbf{M_2}(z) = \mathbf{M_2}(z) - : K_2 : \mathbf{M_2}(z)dz. \tag{3.26}$$

Then

$$\frac{d\mathbf{M_2}(z)}{dz} = - : K_2 : \mathbf{M_2}(z) \tag{3.27}$$

where $- : K_2 :$ is $\mathbf{M_2} - \mathbf{1}$.

3.2 Moment Transport

The transfer map has been derived to transport a single particle through a solenoid. In ionisation cooling, it is important to consider the transport of the entire bunch through the solenoid. The bunch is parameterised in terms of moments of the phase space coordinates such as the second moments, variance and covariance. At this stage it is useful to consider moment transport in general, although second moment transport is of especial importance in ionisation cooling.

Consider a bunch of particles with some continuous density function in phase space $h(\vec{u})$. N^{th} order *raw* moments of the bunch are defined by

$$< u_{i_1} u_{i_2} \dots u_{i_N} >= \int d^6\vec{u}\, h(\vec{u})\, u_{i_1} u_{i_2} \dots u_{i_N}. \tag{3.28}$$

A raw moment is a moment taken about $\vec{0}$ rather than the centre of the probability distribution. In the case of beam transport, this is the reference trajectory.

Consider the density of particles in a volume element at some point z^{in} and again at some point z^{fin}. By Liouville's Theorem [36], the density in a volume element is constant under a canonical transformation. Hence

$$h^{fin}(\vec{u}^{fin}) = h^{in}(\vec{u}^{in}) \tag{3.29}$$

and by the definition of the transfer matrix (3.20)

$$h^{fin}(\vec{u}^{in}) = h(\mathbf{M^{-1}}\vec{u}^{fin}). \tag{3.30}$$

The N^{th} moments at z^{fin} are

$$\begin{aligned} < u_{i_1} u_{i_2} \dots u_{i_N} >^{fin} &= \int d^6\vec{u}^{fin} h^{fin}(\vec{u}^{fin}) u_{i_1}^{fin} u_{i_2}^{fin} \dots u_{i_N}^{fin} \\ &= \int d^6\vec{u}^{fin} h^{in}(\mathbf{M^{-1}}\vec{u}^{fin}) u_{i_1}^{fin} u_{i_2}^{fin} \dots u_{i_N}^{fin}. \end{aligned} \tag{3.31}$$

Substitute $\mathbf{M}\vec{u}^{in}$ for $\vec{u}^{fin}$ and note that according to the Poincare Integral Invariant Theorem [36] the volume element is constant under canonical transformations, i.e. $d^6\vec{u}^{fin} = d^6\vec{u}^{in}$. Then

$$\begin{aligned} < u_{i_1} u_{i_2} \dots u_{i_N} >^{fin} &= \int d^6\vec{u}^{in} h^{in}(\vec{u}^{in}) (\mathbf{M}\vec{u}^{in})_{i_1} (\mathbf{M}\vec{u}^{in})_{i_2} \dots (\mathbf{M}\vec{u}^{in})_{i_N} \\ &= < (\mathbf{M}\vec{u}^{in})_{i_1} (\mathbf{M}\vec{u}^{in})_{i_2} \dots (\mathbf{M}\vec{u}^{in})_{i_N} > . \end{aligned} \tag{3.32}$$

In the case of second moments and linear transport this becomes

$$< u_{i_1}^{fin} u_{i_2}^{fin} >=< \sum_{j_1} (m_{i_1 j_1} u_{j_1}^{in}) \sum_{j_2} (m_{i_2 j_2} u_{j_2}^{in}) > . \tag{3.33}$$

This can be conveniently expressed as a matrix product; the matrix multiplication rule is

$$(\mathbf{AB})_{ik} = \sum_j (a_{ij} b_{jk})) \tag{3.34}$$

so that

$$(\mathbf{M_2}\mathbf{V}^{in}\mathbf{M_2^T})_{i_1 i_2} = \sum_{j_1}(m_{i_1 j_1}\sum_{j_2}(v^{in}_{j_1 j_2} m_{i_2 j_2})). \tag{3.35}$$

Here $\mathbf{V}^{in}$ is the input matrix of covariances with elements $v^{in}_{ab} =< u^{in}_a u^{in}_b >$ so that

$$(\mathbf{M_2}\mathbf{V}^{in}\mathbf{M_2^T})_{i_1 i_2} =< \sum_{j_1}(m_{i_1 j_1} u^{in}_{j_1})\sum_{j_2}(m_{i_2 j_2} u^{in}_{j_2}) > \tag{3.36}$$

which is $\mathbf{V}^{fin}$. Hence

$$\mathbf{V}^{fin} = \mathbf{M_2}\mathbf{V}^{in}\mathbf{M_2^T} \tag{3.37}$$

3.2.1 Emittance

The aim of ionisation cooling is to reduce a quantity known as bunch emittance. This is closely related to the phase space volume occupied by the bunch, and is an invariant under linear beam optics.

In general, for the method followed above, the transfer map is symplectic. The condition for a matrix to be symplectic is

$$\mathbf{M^T J M} = \mathbf{J} \tag{3.38}$$

where

$$\mathbf{J} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & -1 & 0 \end{pmatrix} \tag{3.39}$$

and it can be seen that in particular the solenoidal linear transfer map $\mathbf{M_2}$ obeys this condition for canonical momenta. Note that the determinants of matrices $\mathbf{A}$ and $\mathbf{B}$ are related by

$$|\mathbf{AB}| = |\mathbf{A}||\mathbf{B}| \tag{3.40}$$

and the determinant of the matrix and its transpose is related by

$$|\mathbf{A}| = |\mathbf{A^T}|. \tag{3.41}$$

Then, for the transfer matrix to obey the symplectic condition it must have determinant 1; using (3.32) it can be seen that the covariance matrices are related by

$$|\mathbf{V}^{fin}| = |\mathbf{V}^{in}|. \tag{3.42}$$

The normalised beam emittance in six dimensions is defined by

$$\epsilon_n^{6D} = \frac{\sqrt[6]{|\mathbf{V}^{6D}|}}{m}. \tag{3.43}$$

In this thesis, the emittance will be calculated in kinetic phase space variables. The transformation from canonical phase space to kinetic phase space under the linear approximation is given by

$$p_t^c = p_t^k \tag{3.44}$$

$$p_x^c = p_x^k - \frac{yB_0}{2} \tag{3.45}$$

$$p_y^c = p_y^k + \frac{xB_0}{2}. \tag{3.46}$$

This is a skew transformation so conserves phase space volume and hence emittance.

In the literature it is common to make various alterations to this basic definition. For example, the trace space coordinate system is often considered where $\vec{u}_{trace} = (t, dt/dz, x, dx/dz, y, dy/dz)$ [37]. In this case,

$$\epsilon^{trace} = \sqrt[6]{|\mathbf{V}^{6D}(\vec{u}_{trace})|} \tag{3.47}$$

is known as the unnormalised beam emittance. This is constant in the paraxial approximation where $(dt/dz, dx/dz, dy/dz)$ and only while the beam momentum is constant. A normalised emittance can be defined by taking $< p_z > \epsilon^{trace} \approx \epsilon_n$ that is constant even under a change in momentum. Also, sometimes energy is replaced by momentum or momentum in the $\vec{z}$ direction.

In solenoidal beam optics the longitudinal and transverse phase spaces are decoupled to first order. This means that (x, p_x, y, p_y) are independent of (t, p_t), so it is possible to define transverse and longitudinal emittances that are also conserved by the first order transfer map. Explicitly

$$\epsilon_n^{4D} = \frac{\sqrt[4]{|\mathbf{V}^{4D}(x, p_x, y, p_y)|}}{m} \tag{3.48}$$

and

$$\epsilon_n^{2D} = \frac{\sqrt{|\mathbf{V}^{2D}(t, p_t)|}}{m} \tag{3.49}$$

where these equations are valid for both canonical and kinetic momenta. By considering the motion in an appropriate rotating frame, decoupled transverse motion and 2D emittances can also be found [28]. However, the 4D case is sufficient for this thesis.

3.2.2 Geometric Interpretation of Emittance Conservation

Emittance conservation has been derived in a highly mathematical manner. It is instructive to consider the geometrical interpretation of emittance and relate it to Liouville's Theorem. In a two dimensional space with $\vec{u} = (u_1, u_2)$, the equation for an ellipse is

$$1 = u_1^2/v_{22} + 2u_1u_2/v_{12} + u_2^2/v_{11} = \vec{u}^T\mathbf{V^{-1}}\vec{u} \tag{3.50}$$

where $\mathbf{V}$ is a two by two matrix that defines the ellipse. The area of the ellipse is

$$A = \pi\sqrt{|\mathbf{V}|}. \tag{3.51}$$

When compared to the expressions for emittance given above, it is apparent that emittance is intimately related to the area occupied by an ellipse aligned to the bunch by the covariance matrix (or, in higher dimensions, the content of a hyper-ellipsoid [38]).

A linear transformation is defined as one where

$$x^{fin} = m_{11}u_1^{in} + m_{12}u_2^{in}, \tag{3.52}$$

$$y^{fin} = m_{21}u_1^{in} + m_{22}u_2^{in}, \tag{3.53}$$

that is u_1^{fin} and u_2^{fin} are a linear combination of u_1^{in} and u_2^{in}. Geometrically, linear transformations correspond to rotations, stretches, and skews. Crucially, under a linear transformation straight lines stay straight, and ellipses stay elliptical. The first order transfer map was a linear transformation.

Additionally, great pains were taken to ensure that the transfer map is symplectic. A symplectic transfer map obeys certain laws; crucially, the Poincarre Integral Invariant tells us that under a symplectic transfer map, the area (content) of elements of phase space is conserved.

These two properties are required for emittance to be conserved. Firstly, the area contained within contours of particle density must be conserved. Secondly, elliptical contours must remain elliptical. In muon cooling, emittance reduction is achieved by placing material in the beamline, a process that is non-symplectic. In MICE, emittance growth occurs from chromatic aberrations due to the very large momentum spread of the beam; here the linear approximation of beam optics breaks down and second and higher order transfer maps need to be considered where beam transport is no longer a linear transformation.

3.2.3 Canonical Angular Momentum

In general canonical angular momentum is also a conserved quantity in solenoidal beam optics. The cylindrical symmetry of the solenoid ensures that the Hamiltonian is cyclic under rotation about $\vec{z}$. Under such a rotation the Hamiltonian will remain invariant. As the generating function for a rotation is the canonical angular momentum along $\vec{z}$, this ensures that this component of canonical angular momentum is invariant [36].

The canonical angular momentum is given by

$$L_{can} = xp_y^c - yp_x^c \tag{3.54}$$

and its mean for an ensemble of particles is exactly

$$< L_{can} >=< xp_y^c > - < yp_x^c > . \tag{3.55}$$

In terms of kinetic momentum under the linear approximation this is given by

$$< L_{can} >=< xp_y^k > - < yp_x^k > + \frac{B_0}{2}(< x^2 > + < y^2 >) \tag{3.56}$$

where $< xp_y^k > - < yp_x^k >$ is the kinetic angular momentum L_{kin}.

In designing a cooling channel, care must be taken to ensure canonical angular momentum is conserved. As is demonstrated in the next chapter, muons passing through material lose kinetic angular momentum on average, which can lead to a change in canonical angular momentum. By flipping the solenoid field polarity in the cooling channel or designing the magnetic lattice so that the field is zero at the absorbers, the canonical angular momentum can be held constant. This concept is considered in more detail in later sections.

3.3 Single Particle Emittance

Single Particle Emittance (SPE), ϵ_i, is another useful quantity, defined by

$$\epsilon_i = \epsilon_n \vec{U}^T \mathbf{V^{-1}} \vec{U}, \tag{3.57}$$

also referred to as particle amplitude. This quantity is of considerable interest both because it is conserved and because it is closely related to emittance. Consider

$$\epsilon_i^{fin} = \epsilon_n \vec{U_{fin}}^T \mathbf{V_{fin}^{-1}} \vec{U_{fin}} \tag{3.58}$$

and use the transport equations described above so that

$$\epsilon_i^{fin} = \epsilon_n \vec{U_{in}}^T \mathbf{M_2^T} (\mathbf{M_2 V_{in} M_2^T})^{-1} \mathbf{M_2} \vec{U_{in}}. \tag{3.59}$$

The inverse of a general matrix product $\mathbf{C} = \mathbf{AB}$ is given by $\mathbf{C^{-1}} = \mathbf{B^{-1}A^{-1}}$ so that

$$\epsilon_i^{fin} = \epsilon_n \vec{U_{in}}^T \mathbf{M_2^T} (\mathbf{M_2^T})^{-1} (\mathbf{V_{in}})^{-1} \mathbf{M_2^{-1}} \mathbf{M_2} \vec{U_{in}} \tag{3.60}$$

or

$$\epsilon_i^{fin} = \epsilon_n \vec{U_{in}}^T (\mathbf{V_{in}})^{-1} \vec{U_{in}}, \tag{3.61}$$

which is just the definition of ϵ_i^{in}, so that the SPE is conserved. By analogy with the treatment above, longitudinal, transverse and 6D SPEs can be defined because each phase space is decoupled. This amplitude is a useful quantity as it allows a treatment of cooling without requiring a bunch.

3.3.1 Geometric Interpretation of SPE

Just as there was a somewhat intuitive geometric interpretation of emittance, so there is a similar geometric interpretation of SPE. If the definition of SPE is examined, it will be noted that this too is the area of an ellipse. In this case, it is the ellipse similar to the ellipse defined by RMS emittance, but with radius corresponding to the particle position in the bunch. So particles with SPE equal to the RMS bunch emittance sit exactly on the RMS ellipse. Further, the conservation of SPE indicates that particles will always stay on the same elliptical contour, so long as the linear approximation applies.

3.3.2 Relationship between SPE and Bunch Emittance

It can also be shown that the bunch emittance is proportional to the mean of the SPEs of the bunch. The definition of SPE can be written as

$$\epsilon_i = \epsilon_n \sum_{j}^{j=2N} \sum_{k}^{k=2N} (u_j u_k v_{jk}^{-1}) \tag{3.62}$$

where v_{jk}^{-1} is the jth, kth element of the inverse of the covariance matrix. Taking the mean over all particles in the bunch, for a 2N dimensional phase space,

$$< \epsilon_i >= \epsilon_n \sum_{j}^{j=2N} \sum_{k}^{k=2N} (< u_j u_k > v_{jk}^{-1}). \tag{3.63}$$

Now, the jth diagonal element in the matrix product $\mathbf{VV^{-1}}$ is given by

$$(\mathbf{VV^{-1}})_{jj} = \sum_{k}^{k=2N} (v_{kj} v_{jk}^{-1}) = 1. \tag{3.64}$$

Note that the ordering of indices is interchangeable as $\mathbf{V}$ is symmetric. Then (3.63) becomes

$$< \epsilon_i >= 2N\epsilon_n. \tag{3.65}$$

This is a useful relationship. It demonstrates that bunch emittance is conserved even in the limit of few particles and also demonstrates the relationship between particle SPEs and bunch emittance.

3.4 Higher Order Transfer Maps

It is instructive to consider the effects that second order corrections might introduce to the beam transport outlined above. Recall that the Hamiltonian is given by

$$\begin{aligned} K_z(\vec{u};z) = & - [(p_t^c)^2 + p^2 + 2p_0^c p_t^c - (p_x^c - qA_x)^2 - (p_y^c - qA_y)^2]^{1/2} \\ & - \frac{p_t^c + p_0^c}{v_0}. \end{aligned} \tag{3.66}$$

Expanding K_z as a Taylor series it is seen that

$$K_3(\vec{u};z) = \frac{p_0^c p_t^c K_2}{2p^2}. \tag{3.67}$$

The second order transfer map from z to $z + dz$ is given by

$$\mathbf{M_3} = \mathbf{M_2} + : K_3(\vec{u};z) := \mathbf{M_2} + \frac{p_0^c}{2p^2}(p_t^c : K_2 : + : p_t^c : K_2). \tag{3.68}$$

Then in transverse phase space, second order aberrations arise from the $p_t^c : K_2 : p_0^c/(2p^2)$ term and are purely chromatic.

Third order terms are more complicated. It is noted that they fall into three categories: terms independent of B_z; terms proportional to B_z; and terms proportional to $\partial^2 B_z/\partial z^2$. Thus in regions where the magnetic field gradient is swiftly changing or close to zero, the third order terms may behave differently to other regions.

3.5 Lattice Stability

Define the linear transfer map over a single cell of a periodic lattice, $\mathbf{M_L}$. Then the linear transfer map through n cells is $\mathbf{M_L}^n$. In order for the lattice to be stable, $\mathbf{M_L}^n$ must be finite and real in the limit of large n.

Suppose that $\mathbf{M_L}$ can be decomposed into decoupled two dimensional transfer maps $\mathbf{N_i}$ by some transformation $\mathbf{R}$ with $|\mathbf{R}| = 1$, that is

$$\mathbf{M_L} = \mathbf{RNR^{-1}} \tag{3.69}$$

where

$$\mathbf{N} = \begin{pmatrix} \mathbf{N_1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{N_2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{N_3} \end{pmatrix} \tag{3.70}$$

and each $\mathbf{N_i}$ is a two by two matrix. Then

$$\mathbf{M_L}^n = (\mathbf{RNR^{-1}})^n. \tag{3.71}$$

Expanding 3.71 and collapsing pairs of $\mathbf{R^{-1}R}$ gives

$$\mathbf{M_L}^n = \mathbf{RN}^n\mathbf{R^{-1}}. \tag{3.72}$$

Then the condition $\mathbf{M_L}^n$ is finite is equivalent to the condition that $\mathbf{N}^n$ is finite, which is true if each of $\mathbf{N_i}^n$ are finite. This condition is obeyed if the eigenvalues of $\mathbf{N_i}$ are finite and real. The eigenvalues of $\mathbf{N_i}$, λ_i are defined by

$$\mathbf{N_i U_i} = \lambda \mathbf{U_i} \tag{3.73}$$

where $\mathbf{U_i}$ is a two-vector in the phase space of $\mathbf{N_i}$. Taking the determinant gives

$$\lambda^2 - \lambda(n_{11} + n_{22}) + n_{11}n_{22} - n_{12}n_{21} = 0. \tag{3.74}$$

Defining an angle μ such that $\cos\mu = \frac{n_{11}+n_{22}}{2}$. Then eigenvalues are

$$\lambda = \cos\mu \pm i\sin\mu \tag{3.75}$$

and $\mathbf{N_i}$ can be decomposed to

$$\mathbf{N_i} = \mathbf{I}\cos\mu + \mathbf{J_i}\sin\mu. \tag{3.76}$$

Here $\mathbf{I}$ is the identity matrix and $\mathbf{J_i}$ is a matrix with parameters

$$\mathbf{J_i} = \begin{pmatrix} \alpha_i & \beta_i \\ -\gamma_i & -\alpha_i \end{pmatrix}. \tag{3.77}$$

Then for n cells,

$$\mathbf{N_i}^n = \mathbf{I}\cos(n\mu) + \mathbf{J_i}\sin(n\mu). \tag{3.78}$$

For stability, $\cos n\mu$ and $\sin n\mu$ are required to be stable for all values of n. This means they must be less than 1 for all n, that is

$$-2 < Tr(\mathbf{N_i}) < 2. \tag{3.79}$$

The trace of $\mathbf{N}$ is invariant under a transformation of the form $\mathbf{R^{-1}NR}$, so that an alternative condition is

$$-2 < Tr(\mathbf{M}_{//}) < 2 \tag{3.80}$$

$$-4 < Tr(\mathbf{M}_{\perp}) < 4, \tag{3.81}$$

where $\mathbf{M}_{//}$ and $\mathbf{M}_{\perp}$ are the decoupled longitudinal and transverse transfer matrices.

This condition describes in general whether a lattice is stable. In unstable lattices, higher order transfer maps become dominant and the first order approximation is no longer valid. In particular, beam emittance increases, the beam becomes very large and transmission is very poor. As the transfer matrix is momentum dependent the lattice can be divided into momentum regions where the lattice is non-resonant and transmission is acceptable, *pass-bands*, and where the lattice is resonant and transmission is very poor, *stop-bands*. Ionisation cooling channels typically contain very high emittance beams with large energy spreads, and care must be taken to ensure that the pass-bands are large enough to encompass the full momentum spread of the beam.

3.5.1 Decoupling Transformation

By considering transport in the rotating Larmor frame, the equations of motion for x and y are decoupled for kinetic momenta. Thus the transformation $\mathbf{R}$ comes in two parts, first to kinetic momentum coordinates and then to the decoupled motion. The kinetic phase space vector $\mathbf{U^k}$, with kinetic momenta instead of canonical ones, is related to the canonical phase space vector $\mathbf{U^c}$ by the transformation $\mathbf{R_{c\to k}}$

$$\mathbf{U^k} = \mathbf{R_{c\to k}}\mathbf{U^c} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & \frac{B_0}{2} & 0 \\ 0 & 0 & 1 & 0 \\ -\frac{B_0}{2} & 0 & 0 & 1 \end{pmatrix} \mathbf{U^c}. \tag{3.82}$$

Then the transformation to the Larmor frame $\mathbf{R_{lar}}$ is given by [28]

$$\mathbf{U^k_{lar}} = \mathbf{R_{lar}}\mathbf{U^k} = \begin{pmatrix} \cos\theta & 0 & \sin\theta & 0 \\ -\kappa\sin\theta & \cos\theta & \kappa\cos\theta & \sin\theta \\ \sin\theta & 0 & \cos\theta & 0 \\ -\kappa\cos\theta & \sin\theta & -\kappa\sin\theta & \cos\theta \end{pmatrix} \tag{3.83}$$

where $\kappa = d\theta/dz = B_0/2$. In this frame, not only is motion decoupled but it is also symmetric in x and y. Hence transverse $\mathbf{N}$ are identical. When the lattice is unstable, it is described as *resonant*.

3.5.2 Twiss Parameters

The functions β, α and γ are the Twiss parameters and are periodic functions of the lattice. For a solenoid there are two sets of independent Twiss parameters, for longitudinal and transverse phase spaces. The requirement that $|\mathbf{N}| = 1$ gives the relationship $\beta_i\gamma_i - \alpha_i^2 = 1$.

Consider a beam with a covariance matrix $\mathbf{W}$ in the frame described by the decoupled coordinate system of $\mathbf{N_l}$. It can be seen that under a transformation through a single cell, $\mathbf{W}$ will remain unchanged by $\mathbf{N_l}$ if the covariances are related by

$$\mathbf{W_i} = A \begin{pmatrix} \beta_i & -\alpha_i \\ -\alpha_i & \gamma_i \end{pmatrix} \tag{3.84}$$

where A is an arbitrary constant.

Under the assumption of a cylindrically symmetric beam the transverse covariance matrix can be written as

$$\mathbf{V}_\perp = mc\epsilon_n \begin{pmatrix} \beta_\perp/p_0 & & & \\ -\alpha_\perp & \gamma_\perp p_0 & & \\ 0 & \beta_\perp B_0 - \mathcal{L} & \beta_\perp/p_0 & \\ -(\beta_\perp B_0 - \mathcal{L}) & 0 & -\alpha_\perp & \gamma_\perp p_0 \end{pmatrix}. \tag{3.85}$$

Terms above the diagonal are not included as the covariance matrix is symmetric. The leading constant has been chosen according to the emittance definition. $\mathcal{L}$ is the constant canonical angular momentum normalised to the beam emittance, $\mathcal{L} = mc\epsilon_n/L_{can}$. $\beta_p erp$ is related to $\alpha_\perp$, $\gamma_\perp$ and $\mathcal{L}$ by

$$\gamma_\perp = \frac{1 + \alpha_\perp^2 + (\beta_\perp \kappa - \mathcal{L})^2}{\beta_\perp}. \tag{3.86}$$

The determinant $|\mathbf{V}_\perp|$ is given by

$$\sqrt{|\mathbf{V}_\perp|} = (< x^2 >< p_x^2 > - < xp_x >^2 - < xp_y >^2) \tag{3.87}$$

When $\beta_\perp$ and $\sqrt{1+\mathcal{L}^2}\beta_i$ are equal, the transverse covariance matrix is described as matched and the covariance matrix is periodic.

In general the beam β function can be calculated by

$$\beta_\perp = \frac{(< x^2 > + < y^2 >)p}{mc\epsilon_n}. \tag{3.88}$$

Under the action of $\mathbf{R_{c\to k}^{-1} M_2 R_{c\to k}}$ on $\mathbf{V}_\perp$, $\alpha_\perp$ is related to $\beta_\perp$ by

$$-2\alpha_\perp = \frac{d\beta_\perp}{dz} \tag{3.89}$$

and

$$2\beta_\perp \beta_\perp'' - (\beta_\perp')^2 + 4\beta_\perp^2 B_0^2 - 4(1+\mathcal{L}^2) = 0. \tag{3.90}$$

3.6 RF and Solenoid Combined Transfer Map

Slight modifications are required to calculate transport through solenoids and cavities as shown. A general treatment to arbitrary order is given for pillboxes in [39] and for general RF cavities in [40]. Only pillboxes are treated here and only to first order.

3.6.1 Hamiltonian Including RF

The vector potential from a cylindrically symmetric RF cavity with angular frequency ω is

$$A_r = \sum_n \frac{ik_{zn}}{k_{rn}} E_0 J_1(k_{rn}r) e^{i(\omega T - k_{zn}z)} \tag{3.91}$$

$$A_z = \sum_n \frac{E_n}{\omega} J_0(k_{rn}r) e^{i(\omega T - k_{zn}z)} \tag{3.92}$$

$$A_\phi = 0. \tag{3.93}$$

Here n is the space harmonic index, $k_{rn}^2 + k_{zn}^2 = \left(\frac{\omega}{c}\right)^2$, $k_{zn} = k_z + 2\pi n/d$ and J_m are Bessel functions of the first kind. In ionisation cooling channels cavities usually operate in the TM_{010} mode where $n = 0$ and $k_{zn} = 0$. The Bessel functions can be expanded as a power series [41]

$$J_m(k_{rn}r) = \sum_{l=0}^{\infty} \left(\frac{-1^l}{l!(m+l)!} (k_{rn}r/2)^{2l+m} \right) \tag{3.94}$$

Adding to the solenoidal vector potential and truncating as a second order polynomial gives

$$\vec{A}_2 = \frac{1}{2q} \begin{pmatrix} -B_0 y \\ B_0 x \\ A_z(r=0,T) - K_c(x^2+y^2) \end{pmatrix} \tag{3.95}$$

where $A_r = 0$ as only the TM_{010} mode is considered and $K_c = E_0 e^{i\omega T_0} k_{rn}^2/(4\omega)$.

In general, the Hamiltonian can also be truncated as a second order polynomial. As before, the Hamiltonian is expanded using the binomial expansion, now retaining A_z, giving

$$H = -qA_z - \frac{1}{2}\left(\frac{p_t^2}{p} - \frac{(p_x^c)^2 + (p_y^c)^2}{p} + \frac{2q}{p}(A_x p_x^c + A_y p_y^c) - \frac{q^2}{p}(A_x^2 + A_y^2) \right). \tag{3.96}$$

Substituting for first order terms in the expansion of $\vec{A}$,

$$\begin{aligned} H_2 &= -qA_z(r=0) + (x^2+y^2)\left(\frac{K_c}{2} + \frac{B_0^2}{8p}\right) - \frac{p_t^2}{2p\beta^2\gamma^2} \\ &+ \frac{(p_x^c)^2 + (p_y^c)^2}{2p} - \frac{B_0}{2p}(x p_y^c - y p_x^c) \end{aligned} \tag{3.97}$$

3.6.2 Transfer Map

It can be seen from (3.97) that to leading order the longitudinal and transverse phase spaces are still decoupled. Once again the transfer map is calculated by applying the equations of motion

$$\frac{d\vec{u}}{dz} = -[H, \vec{u}] \tag{3.98}$$

to the Hamiltonian above. The transfer map for canonical phase space variables in the transverse phase space (x, p_x^c, y, p_y^c) is given by

$$\mathbf{M}_{2\perp} = 1 + \begin{pmatrix} 0 & 1/p & B_0/2p & 0 \\ -K & 0 & 0 & B_0/2p \\ -B_0/2p & 0 & 0 & 1/p \\ 0 & -B_0/2p & -K & 0 \end{pmatrix} dz \tag{3.99}$$

where $K = K_c + B_0^2/4p$.

The transfer map in longitudinal phase space (t, p_t^c) is

$$\mathbf{M}_{2\perp} = 1 + \begin{pmatrix} 0 & -1/\beta^2\gamma^2 p \\ K_c & 0 \end{pmatrix} dz. \tag{3.100}$$

The inclusion of pillboxes then introduces additional terms in the transverse and longitudinal transfer map. In longitudinal phase space, the pill box produces a focusing force, which provides longitudinal phase space stability. In transverse phase space the pill box produces a weak defocussing force. This is negligible in comparison to the solenoidal focussing in ionisation cooling channels.

3.7 Solenoidal Lattice Design

The solenoidal lattice is important in cooling channel design for several reasons. It focusses the beam onto the absorbers which enhances cooling and it contains the beam to prevent scraping. As seen above the focussing is highly momentum dependent, and indeed at some momenta is resonant.

The focussing can be expressed as a function of the Fourier components of the on-axis field over a cell of the magnetic lattice [28] [29] and this is used to classify different lattices [42]. Different Fourier components lead to different resonances. Wang and Kim give the formula for the $n\pi$ resonance as [29]

$$\sqrt{\theta_0} = n \pm \frac{1}{2}\left|\frac{\theta_n}{\theta_0}\right| + \frac{5}{16}\left|\frac{\theta_n}{\theta_0}\right|^2 + \ldots \tag{3.101}$$

where $\theta_0 =< B_z^2 > qL/(2\pi p)$.

For this document, the SFoFo lattice is particularly interesting. The SFoFo lattice has a minimum in the β function at the beginning and centre of the cell. The on-axis magnetic field is approximately given by [28]

$$B_z = a_1 sin(\pi z/\lambda) + a_3 sin(3\pi z/\lambda). \tag{3.102}$$

By choosing a_1 and a_3 carefully, the resonances can be chosen to be outside of the momenta of interest for the particular beam.

The field is often chosen to have odd symmetry to prevent growth of canonical angular momentum. As shown above, a beam with no kinetic angular momentum in free space will aquire kinetic angular momentum in a solenoidal field. This will be proportional to the on-axis field so that canonical angular momentum is conserved. If the beam has kinetic angular momentum as it passes through material it loses it according to $\Delta L_{kin}/L_{kin} \approx \Delta p_z/p_z$. This leads to build up of canonical angular momentum that can cause a mismatch. By making the field odd, any build up of canonical angular momentum in the first half of the cell is removed in the second half. Also the absorbers can be placed in the zero field region where the beam has no kinetic angular momentum.

Chapter 4

Ionisation Cooling

Ionisation cooling is the technique by which it is proposed to reduce the emittance of the muon beam so that a larger number of muons will fit into the acceptance of the Neutrino Factory acceleration system and storage rings. In ionisation cooling dissipative processes are used to reduce the beam emittance as the muon beam passes through some material.

4.1 Principle of Ionisation Cooling

The significant physical processes that act on muons passing through material at energies of interest are twofold: firstly, the muon momentum is reduced as atoms are ionised: secondly, muons scatter off atoms.

The principle of ionisation cooling is to reduce transverse and longitudinal momentum of muons as they pass through material. This is replaced by longitudinal momentum in RF cavities. Multiple scattering and stochastic effects in the ionisation energy loss (energy straggling) causes heating that reduces or entirely counteracts the cooling effect. This can be related in an expression for the change in emittance due to the passage of a beam through material.

Unless otherwise specified, momenta without a superscript will be used to refer to kinetic momenta in this and subsequent chapters. Energy will be denoted by E rather than p_t^k and p_t will refer to transverse momentum.

4.1.1 Transverse Phase Space

In transverse phase space the change in emittance is given by multiple scattering effects and ionisation energy loss [43]. Second order effects due to energy straggling are neglected. The mean energy change of the beam per unit length of the absorber, dE/dz, is given by the Bethe-Bloch equation and the mean squared

increase in angular spread per unit length is approximated by a Gaussian fit to the Moliere scattering formula with variance [10],

$$\frac{d < \Theta_x^2 >}{dz} \approx \frac{13.6^2}{(p\beta_{rel})^2 L_R}. \tag{4.1}$$

Here Θ_x is the angular divergence of a particle from the beam axis in the x direction, p is the beam momentum, c is the speed of light, β_{rel} is the ratio of momentum to energy and L_R is the material's radiation length.

On passing through a thin piece of material, only the momentum is affected by the material. Recall that the determinant of the transverse covariance matrix $|\mathbf{V}_\perp|$ is given by (3.87),

$$\begin{aligned} \sqrt{|\mathbf{V}_\perp|} &= (< x^2 >< p_x^2 > - < xp_x >^2 - < xp_y >^2) & (4.2) \\ &\approx p_z^2(< x^2 >< \Theta_x^2 > - < x\Theta_x >^2 - < x\Theta_y >^2), & (4.3) \end{aligned}$$

and the beam emittance is given by (3.48)

$$\epsilon_n = \frac{\sqrt[4]{|\mathbf{V}_\perp|}}{m}. \tag{4.4}$$

The change in emittance on passing through a thickness dz of material is

$$\frac{d\epsilon_n}{dz} = \frac{1}{2m^2\epsilon_n} \frac{d\sqrt{|\mathbf{V}_\perp|}}{dz} \tag{4.5}$$

Under (4.1) terms of the form $< q_i q_j >$ and $< q_i \Theta_j >$ are constant, so only p_z and $< \Theta_i^2 >$ terms contribute to the emittance change such that

$$\frac{d\epsilon_n}{dz} \approx \frac{1}{2m^2\epsilon_n}(2\frac{dp_z}{dz}\frac{\sqrt{|\mathbf{V}_\perp|}}{p_z} + < x^2 > p_z^2 \frac{d < \Theta_x^2 >}{dz}). \tag{4.6}$$

Using $E\,dE/dz \approx p_z\,dp_z/dz$ and (4.1),

$$\frac{d\epsilon_n}{dz} \approx \frac{1}{\beta_{rel}^2 E}\left\langle \frac{dE}{dz} \right\rangle \epsilon_n + \frac{1}{2m^2\epsilon_n} < x^2 > \frac{13.6^2}{(\beta_{rel})^2 L_R}. \tag{4.7}$$

Also recall the definition of $\beta_\perp$ (3.88),

$$\beta_\perp = \frac{< x^2 > p}{m\epsilon_n}. \tag{4.8}$$

Then

$$\frac{d\epsilon_n}{dz} \approx \frac{1}{E}\left\langle \frac{dE}{dz} \right\rangle \epsilon_n + \frac{1}{2m}\frac{13.6^2}{L_R}\frac{\beta_\perp}{\beta_{rel}^3 E}. \tag{4.9}$$

Note that $< dE/dz >$ is always negative, so the first term is a cooling term and the second term is a heating term. A more focussed beam with small $\beta_\perp$ will undergo less heating. This is because the multiple scatter is smaller relative to

the angular spread of the beam. Also, as the cooling term is proportional to emittance, there will be an equilibrium emittance below which heating, rather than cooling, will occur, given by

$$\epsilon_n(equilibrium) = \frac{1}{2m}\frac{13.6^2}{L_R}\frac{\beta_\perp}{\beta_{rel} < \frac{dE}{dz} >}. \tag{4.10}$$

This is known as the equilibrium emittance.

4.1.2 Longitudinal Phase Space

In longitudinal phase space, stochastic effects in the energy loss (energy straggling) lead to an increase in the energy spread of the beam, while the curvature of the Bethe-Bloch relation leads to either a cooling or heating effect as muons with lower energy will have a different energy loss to muons with higher energy.

Consider the longitudinal phase space coordinates (t, E) relative to the reference particle. Longitudinal emittance is given by

$$\epsilon_n = \frac{1}{m}\sqrt{< E^2 >< t^2 > - < Et >^2}. \tag{4.11}$$

The Bethe-Bloch curve tends to increase the energy spread for particles with energy below minimum ionising and decrease the energy spread for particles with energy above minimimum ionising. The change in the energy spread due to this effect is given by

$$\frac{d < E^2 >}{dz} = 2\frac{d}{dE}\frac{dE}{dz} < E^2 > . \tag{4.12}$$

The energy straggling distribution follows the Vlasov distribution [10]. This is taken to be approximately Gaussian, so that variances add in quadrature; in reality it is quite heavily skewed, as shown in Figure 4.1. The total change in energy spread is then given by

$$\frac{d < E^2 >}{dz} = \left(2\frac{d}{dE}\frac{dE}{dz}\right) < E^2 > + \left(\frac{d < E^2 >}{dz}\right)_{Vlasov}. \tag{4.13}$$

The width of the Vlasov distribution is proportional to γ^2, so that energy straggling leads to more heating at higher energies.

The optimum energy for cooling is driven by a number of different requirements. Transverse cooling is stronger at lower energies, as is heating; the equilibrium emittance is lower for higher energies. Longitudinal heating due to the Bethe Bloch curve occurs at lower energies while heating due to energy straggling is stronger at higher energies. Additional constraints may arise from the focussing system. Typically, momenta around 200 MeV/c are considered for ionisation cooling channels.

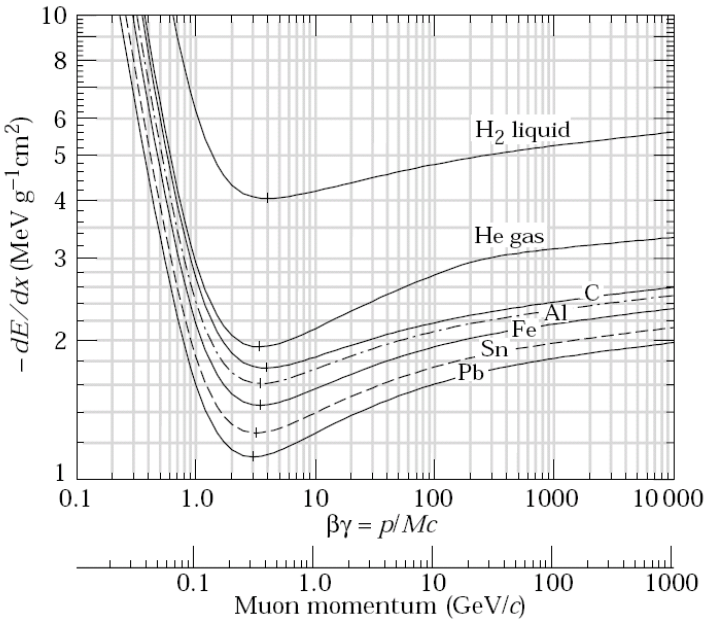

(a)

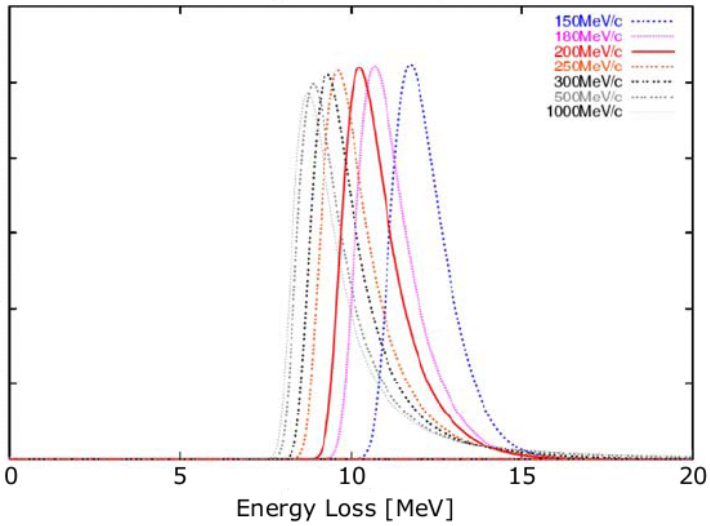

(b)

Figure 4.1: (a) The Bethe-Bloch energy loss per unit length as a function of muon energy in various materials [10] and (b) the Vavilov energy loss distribution through a 350 mm liquid hydrogen absorber for various muon energies [31].

4.1.3 Emittance Exchange

So far the transverse and longitudinal phase spaces have been assumed to be independent. This is true in a linear cooling channel, but when bending magnets are included the two become coupled. As muons travel through the bend higher

energy particles follow a higher radius path and the beam becomes wider. Six dimensional emittance is conserved due to the energy-position correlation that develops (dispersion).

By including a wedge-shaped absorber, particles with a larger energy pass through a thicker portion of the wedge and so the energy spread is reduced. This allows emittance to be moved from the longitudinal plane to the transverse plane so that the beam remains above equilibrium emittance for longer allowing enhanced cooling. This is known as emittance exchange.

4.2 Cooling Channel Schemes

Several different muon cooling channel schemes have been proposed. They can be divided into linear cooling channels, that cool in transverse phase space only, and emittance exchange schemes that cool both transversely and longitudinally.

4.2.1 Linear Cooling Channels

Two linear cooling channels based on liquid Hydrogen absorbers have been proposed and one based on solid Lithium Hydride. The European scheme [11] relies initially on 44 MHz cavities and later 88 MHz cavities providing 2 MV/m and 4 MV/m RF voltage, with transverse focussing provided by 2 T and 5 T solenoidal fields respectively.

The US liquid Hydrogen cooling scheme takes advantage of micro-bunches created by the phase rotation scheme to operate at 201.25 MHz [12]. This has the advantage that a higher RF voltage can be used which decreases the overall cost of the cooling channel. The US scheme employs the SFoFo (Super-Focussing-Focussing) lattice with a peak field of 5 T. In the full cooling channel, the β function is tapered. When the emittance is high, the lattice is designed to focus less in the absorbers and more in the RF cavities to reduce scraping. As the emittance is reduced, tighter focussing can be employed in the absorber resulting in enhanced cooling. Overall, the channel has an equilibrium emittance of about 2 mm (transverse). This scheme is especially interesting as the MICE cooling channel is based on a similar SFoFo lattice. The cooling channel performance is shown in Figure 4.2.

The US Lithium Hydride scheme is more cost optimised and employs a much simpler magnetic lattice [1]. Here, 2.8 T coils are placed between each RF cavity with each coil carrying equal and opposite currents to its neighbour. This is known as an alternating solenoid (aSol) lattice. While it provides less focussing

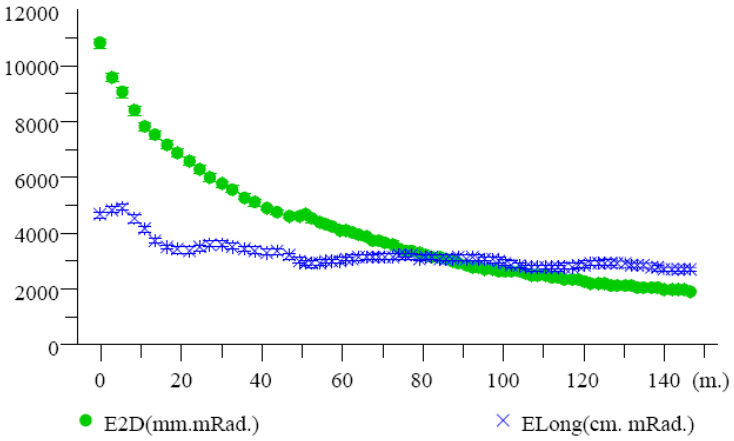

(*a*)

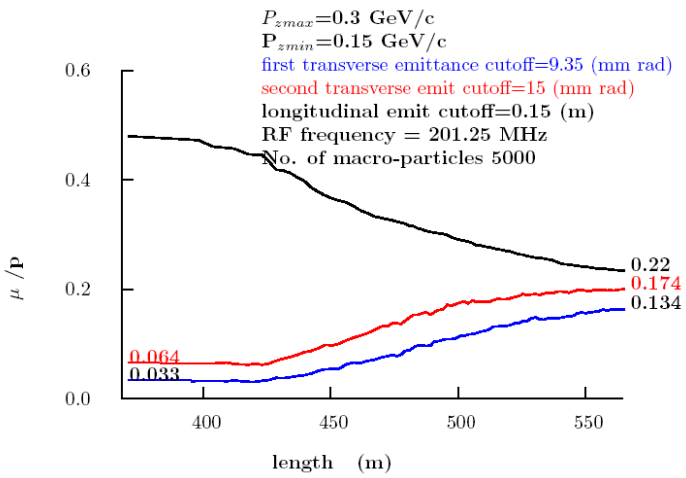

(*b*)

Figure 4.2: The cooling performance and transmission of the Feasibility Study II linear cooling channel [12]. Longitudinal emittance is roughly constant, as muons are lost from the RF bucket and contribute to scraping rather than causing emittance growth. Transverse emittance is significantly reduced and for the tight focussing in the final lattice approaches equilibrium at roughly 2 mm rad.

than the SFoFo lattice, there is no need for a coil to be placed round the outside of the RF cavities so the device is cheaper. Each RF cavity is bounded by Lithium Hydride windows coated with Beryllium that do the actual cooling. This scheme provides a more moderate amount of cooling, with an equilibrium emittance of about 5.5 π mm rad. It is significantly cheaper than the liquid Hydrogen scheme and has been adopted as the baseline design for the ISS.

4.2.2 Emittance Exchange Schemes

Several cooling schemes have been envisaged that take advantage of emittance exchange, although none of them have been incorporated into a neutrino factory design. Three types of ring have been designed, distinguished by different magnetic lattices. In addition, effort has been put into so-called "Guggenheim" coolers and helical coolers.

The ring-based coolers use solenoidal focussing lattices as in the linear case. Bending is provided by tilted solenoids [44] or FFAG magnets [45]. In all these cases solenoidal focusing introduces a kinetic angular momentum into the system that causes the dispersion to rotate between the ring's radial and axial directions. The wedges are then placed to absorb dispersion in the vertical direction i.e. so that the wedge points downwards.

Ring coolers are generally cheaper as they do not require long linacs and provide more cooling due to the emittance exchange. However, injection and extraction is thought to be extremely difficult for the low radii typical of most ring designs. Also for liquid Hydrogen wedges thermal heating due to the passage of the muon beam may be an issue. A Guggenheim cooler [46], where the ring is stretched out into a helix, may solve these issues. In addition, such a design would enable tapering of the beta function as in the case of linear cooling channels, providing additional cooling. Although the cost benefit is lost and additional magnetic shielding is required, such a system would have significantly more cooling than a linear channel.

The final emittance exchange scheme that is envisaged is a so-called helical cooling channel . The helical cooling channel is a linear channel made up of helically arranged dipoles wrapped in a solenoidal focussing field [47]. The reference path is helical in this case, with reference particles of higher momentum having a larger radius helix than reference particles of lower momentum. The higher momentum particles will thus see a longer path length. Filling the channel with absorbers or high pressure gas will cause emittance exchange to occur without the need for complicated wedges. Beam stability is provided by the combined solenoidal and dipole fields. The high pressure gas scheme is attractive as the gas acts as an insulator in the cavities preventing breakdown, even at high magnetic fields. Testing of RF cavities filled with high pressure gas is in progress and a significant improvement in achievable gradients has been observed with no dependence on magnetic field while operating at high pressure [48].

Chapter 5

The Muon Ionisation Cooling Experiment

Ionisation cooling will be demonstrated for the first time by the Muon Ionisation Cooling Experiment (MICE) [49]. MICE will measure the change in 6D emittance over a single 5.5 m cooling cell. The MICE design is based on the US liquid Hydrogen cooling scheme outlined above and is summarised more fully in the MICE Technical Reference Document (TRD) [23]. Muons will be produced from a target in the ISIS proton synchrotron and transported along a pion-muon beamline into MICE. Muons pass through the cooling channel one by one. The muons' phase space coordinates will be measured by time of flight scintillators (TOF 1 and 2) and scintillating fibre tracking detectors (SciFis or trackers) upstream and downstream of the cooling channel. Muons will be distinguished from other particles in the beam using a combination of the spectrometers and the so-called Particle Identification (PID) detectors, three time of flight walls (TOF 0, 1 and 2), a Cerenkov detector and an electron-muon calorimeter. The full MICE design is shown in Figure 5.1.

5.1 Beamline

The MICE muon beamline, shown in Figure 5.2 provides a low intensity muon beam for MICE. Pions will be transported from a target dipping into the fringe of the ISIS proton beam, through a pion decay channel, into a muon transport line and then into MICE. For efficient use of muons it is desirable to have a reasonably good match between the transport beamline and the cooling channel, with selection performed in analysis. Also, the beamline must suppress non-muon events from entering the cooling channel.

The target consists of a $10 \times 10 \times 1$ mm^3 titanium cuboid. The cuboid will

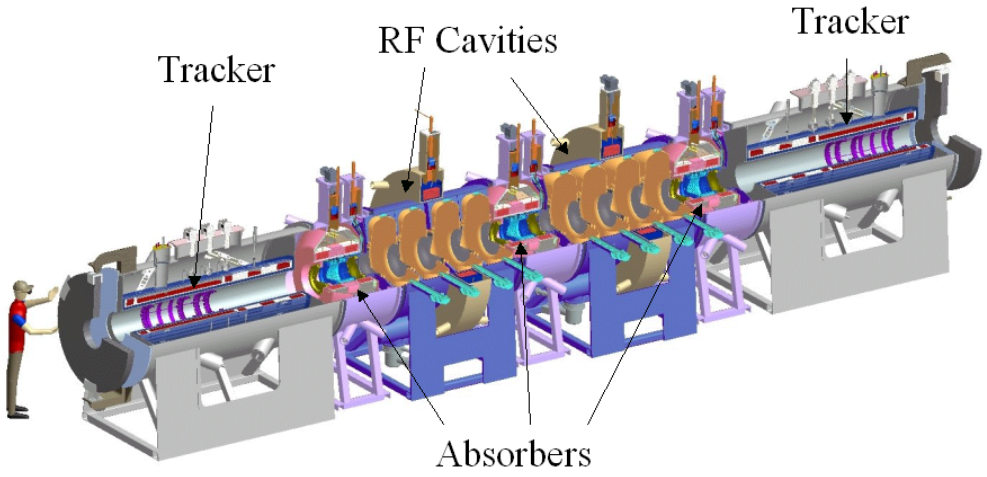

Figure 5.1: Cutaway view of MICE.[23].

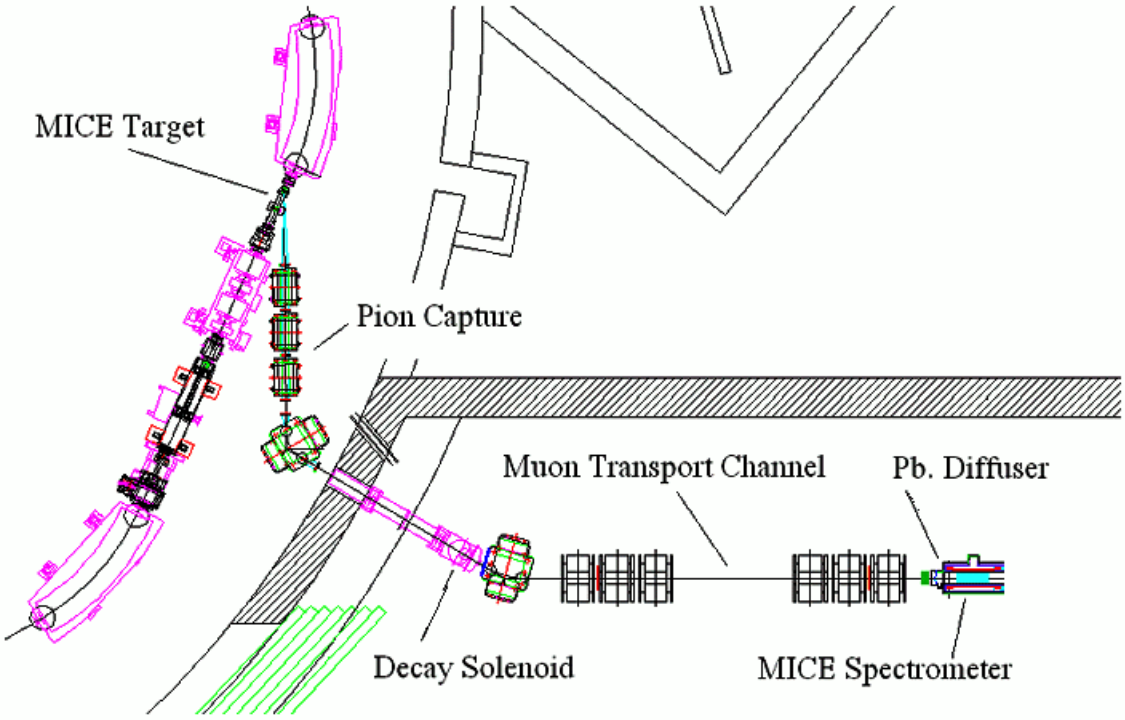

Figure 5.2: Hall layout for the MICE beamline. Particles are transported from the ISIS beamline in the top left to MICE in the bottom right [23].

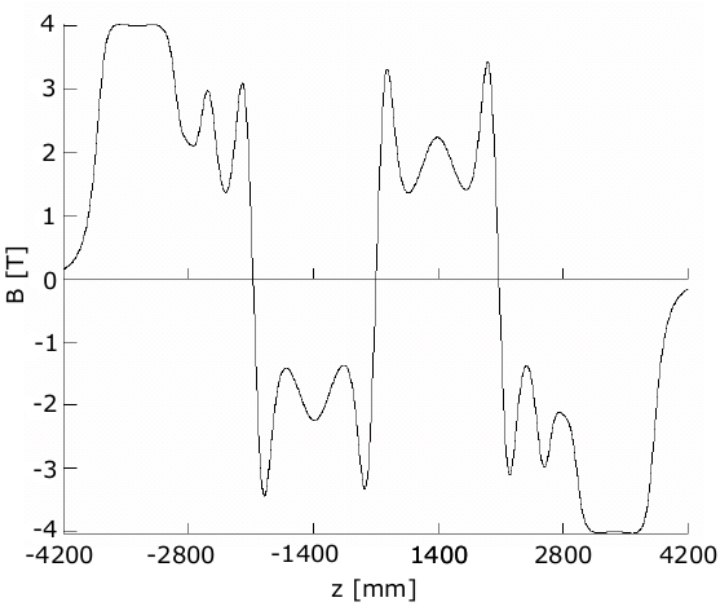

Figure 5.3: On-axis field for step 6 in field flip mode. The SFoFo region is between -2750 mm and +2750 mm [23].

be propelled into and withdrawn from the ISIS beamline by an electromagnetic linear drive, taking protons only at the end of the acceleration cycle when the proton energy is in the range 700-800 MeV.

Downstream of the target, pions are transported through a quadrupole triplet and momentum selection dipole. Most pions decay in a 5 m long, 60 mm bore, 4 T solenoid. Beyond the solenoid remaining pions and protons are absorbed in a thin plastic disc. Momentum selection of muons is achieved using a second dipole before two quadrupole triplets transport them through the PID detectors and into the first spectrometer solenoid. The muon beam optics are controlled by the quadrupole triplets and additionally by a thin lead diffuser of variable thickness that increases the muon beam emittance.

5.2 Cooling Channel

The MICE cooling channel is made up of three distinct components. Current carrying coils provide transverse focussing, absorbers provide cooling and RF cavities replace momentum lost in the absorbers.

	z_{centre} [mm]	Length [mm]	Inner Radius [mm]	Thickness [mm]	Current Density [A/mm^2]
End2	-5951.5	110	258	68.2	142.492
Centre	-5201	1294	258	22	146.9
End1	-4451.5	110	258	61.6	136.8
Match2	-4051.5	197	258	26.4	161.3
Match1	-3611	198	258	52.8	104.6115
Focus3	-2955	210	263	84	113.95
Focus2	-2545	210	263	84	-113.95
Coupling	-1375	250	725	116	-96.21
Focus1	-205	210	263	84	-113.95
Focus1	205	210	263	84	113.95
Coupling	1375	250	725	116	96.21
Focus2	2545	210	263	84	113.95
Focus3	2955	210	263	84	-113.95
Match1	3611	198	258	52.8	-104.6115
Match2	4051.5	197	258	26.4	-161.3
End1	4451.5	110	258	61.6	-136.8
Centre	5201	1294	258	22	-146.9
End2	5951.5	110	258	68.2	-142.492

Table 5.1: Geometry for the MICE coils and currents operating in flip mode [50].

5.2.1 Magnetic Lattice

The MICE magnetic lattice is made up of 18 coils. The MICE coils are described in Table 5.1 for the field flip configuration and the on-axis B_z is shown in Figure 5.3. Each spectrometer contains five coils. Three solenoid coils provide a constant field for the MICE spectrometer and two matching coils match the beam from the spectrometers into the cooling channel. Each of the three absorbers sits between two focus coils. These focus the muon beam in the absorbers to provide optimal cooling. At the centre of each RFCC module a single coupling coil creates the second harmonic term in the field that make the MICE channel SFoFo. The MICE coils are all superconducting with current densities between 16 and 140 A mm^{-2}.

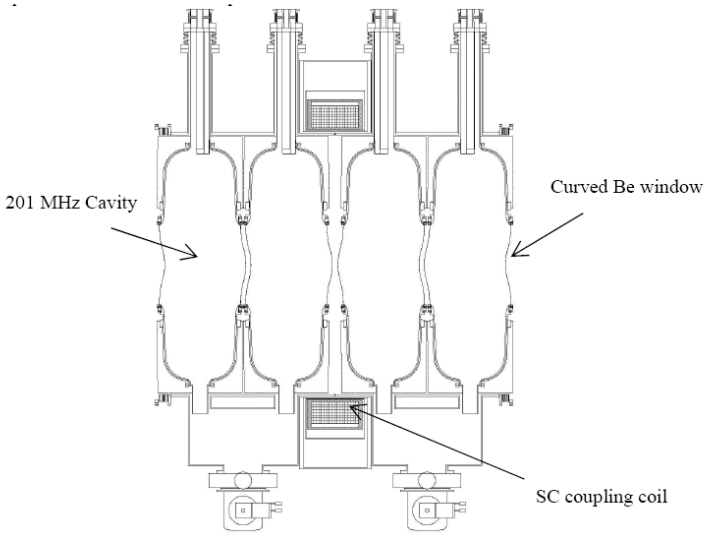

Figure 5.4: The MICE four-cavity RF assembly [23].

5.2.2 RF Cavities

The MICE linacs consist of four normal-conducting copper cavities approximately 430 mm in length and 610 mm in radius with a cross sectional profile shown in Figure 5.4. There are several challenges in the construction and operation of the MICE cavities.

The RF cavities operate at 201.25 MHz with a peak field of 8 MV/m. The US scheme on which MICE is based phased the cavities at 50°. Due to financial constraints MICE will have insufficient power to replace all of the energy lost in the absorbers while running in this mode. Instead, the full muon energy will be replaced for muons that pass through the rf cavities on-crest. As there is no time distribution in the incoming beam, muons at all phases will be sampled.

The cavities must have a large aperture in order to avoid significant scraping from the high emittance beam. Hence 0.38 mm thick Beryllium windows are used to electromagnetically seal the cavities to avoid a high shunt impedance. These are nearly invisible to the muon beam. The windows have been designed to prevent electromagnetic heating from buckling the windows.

Peak surface field is only 8 MV/m compared with the Kilpatrick limit of around 15 MV/m. At full RF power the cavities are designed to operate near to this limit and care must be taken to avoid breakdown. To this end, the cavity walls are carefully cleaned and constructed non-parallel. Despite these preparations, the Kilpatrick limit may be reduced in strong fields preventing the RF cavities from operating at full power. A MICE-like RF cavity is under test at Fermilab in order to understand the behaviour of the breakdown limit in the

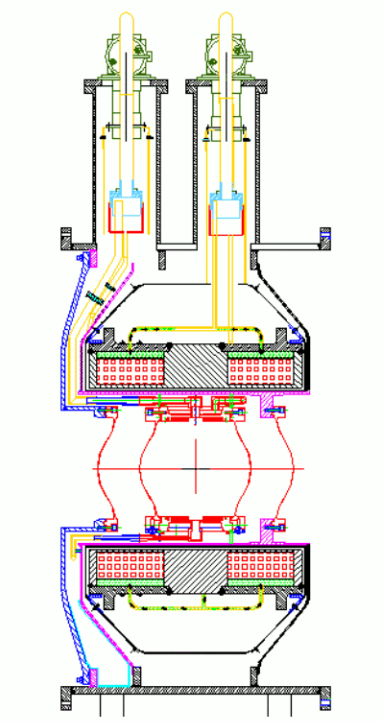

Figure 5.5: A liquid Hydrogen absorber. Focus coils and absorber windows are shown in red. Other colours depict cooling apparatus and supports [23].

presence of intense solenoidal fields. Tests at 805 MHz show reduction by a factor 2 [22].

Even without breakdown, electrons stripped from the cavities by the strong electric fields (dark currents) will be accelerated through MICE into the absorber modules. Most of these electrons should stop in the absorbers, but secondary x-rays may be transported through the absorbers into detector systems. It is expected that this will not degrade SciFi performance, but the effect on the TOFs is an area of active research.

5.2.3 Absorbers

The baseline MICE absorber consists of a 21 litre, 350 mm long liquid Hydrogen vessel sealed with a pair of curved, cylindrically symmetric aluminium windows at each end and cooled using cryocoolers, as shown in Figure 5.5. The absorber is removable, so that different materials may be used. For example, solid absorbers constructed from materials such as Lithium Hydride may replace the liquid Hydrogen absorbers.

The absorber windows are curved so that they can withstand a higher pressure than cylindrical windows, enabling a much thinner construction. A second safety

window is placed in front of the first to prevent Hydrogen leakage in the event of a failure of the absorber windows. A 2400 litre volume between the two windows, normally in vacuum, provides a safety buffer for any leaking Hydrogen.

5.3 Detectors

The detectors measure muon phase space coordinates and eliminate non-muons from the beam. Measurement of phase space coordinates is provided by a TOF and SciFi at either end of the MICE channel. Upstream particle identification is provided by the upstream TOF and SciFi in addition to a threshold Cerenkov detector and another TOF. Downstream particle identification is provided by the downstream TOF and SciFi in addition to an electron-muon calorimeter.

5.3.1 MICE Trackers

A SciFi tracking detector sits on each end of the cooling channel to measure the momentum and position of each particle as it passes through MICE [54]. Each tracker is made up of five stations of scintillating fibres seperated by several centimetres where the position of the muons are measured. By examining the helical track of the muon as it passes through the solenoidal field, the muon momentum and position can be reconstructed at a plane at the inside edge of each tracker.

Each station contains three planes of fibres with threefold rotational symmetry, to give a stereo view of the position of each muon hit. The third plane is redundant but provides additional data to refine the space point resolution. The 350 μm scintillating fibres are bundled into groups of seven, which are then connected to light guides which are read out by cryogenically cooled Visible Light Photon Counters (VLPCs). The VLPCs provide photomultiplication and are subsequently read out to Data Aquisition (DAQ) electronics.

Tracks are reconstructed using a Kalman filter that seeks to minimise the mean square difference between the measured position of hits and the position of hits on a helical path characteristic of a muon in a constant field. The helix is modified to account for energy loss on traversing fibre planes. Position and momentum resolution is effected by the space point resolution of the stations and multiple scattering in the fibre material.

A detailed simulation has been performed of the MICE tracker using measured values for parameters such as light yield and a detailed knowledge of tracker geometry. In table 5.2 the resolution of the tracker in kinetic phase space variables

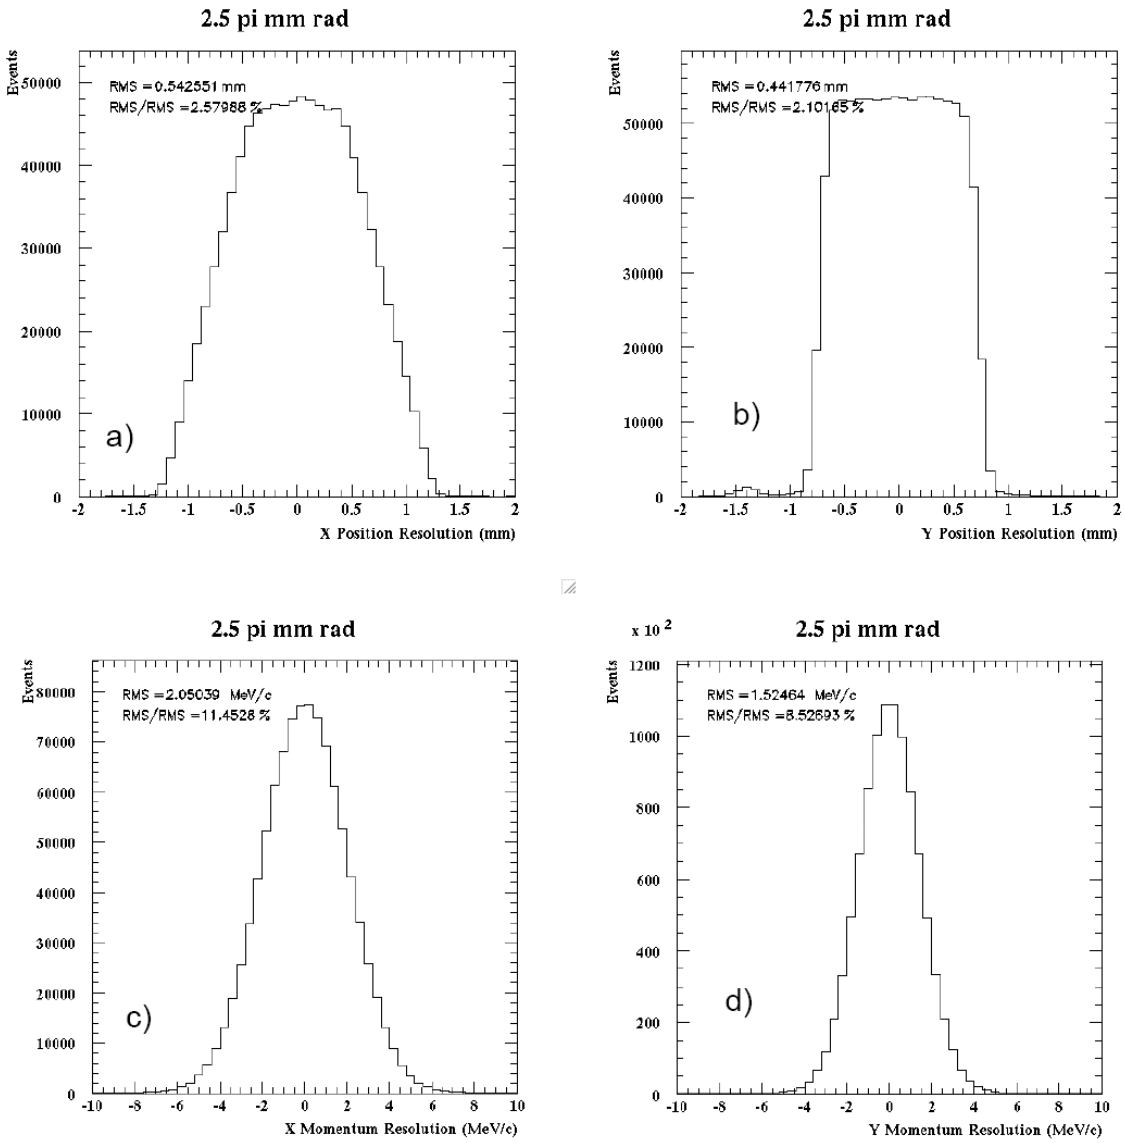

Figure 5.6: Transverse phase space residuals in the MICE tracker [54].

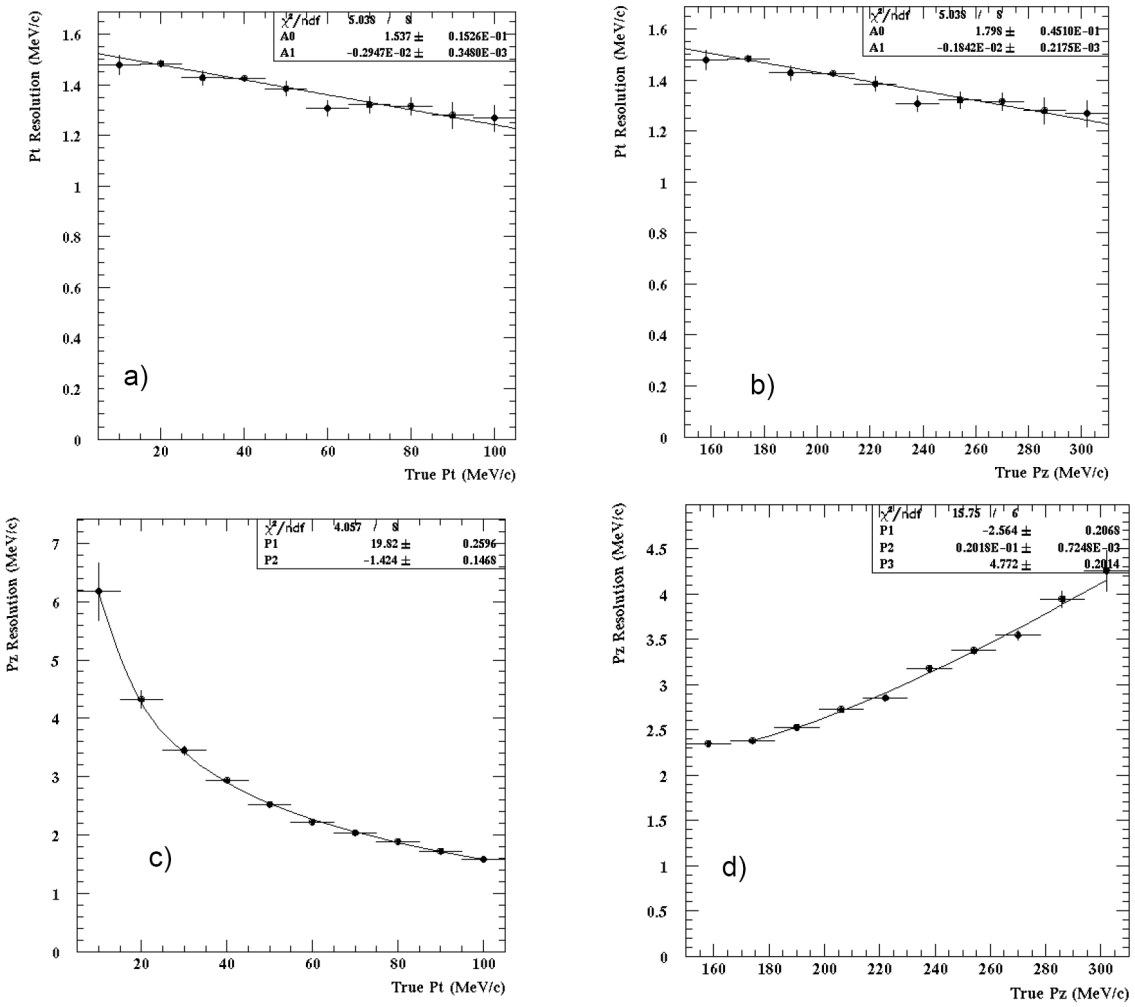

Figure 5.7: Correlations between truth and standard deviation errors of longitudinal and transverse momentum in the MICE tracker[54].

u_i	$\sigma(du_i)$
x [mm]	0.54
y [mm]	0.44
p_x [MeV/c]	2.05
p_y [MeV/c]	1.52
E [MeV]	3.49

Table 5.2: The standard deviation of the residual in kinetic phase space variable u_i for an input beam at equilibrium emittance passing through the upstream tracker [54].

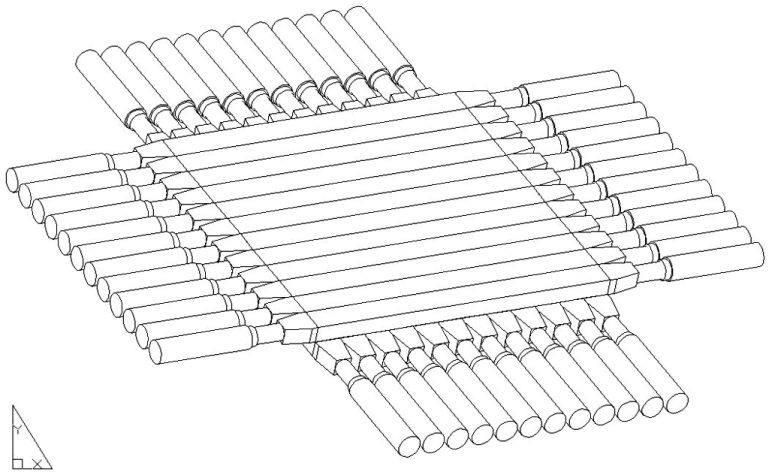

Figure 5.8: Diagram of one of the TOF stations. Plastic scintillator bars lie across the beamline with PMTs at either end. Vertical bars lie behind the horizontal ones [23].

is listed for a beam with 2.5 mm transverse emittance, $\beta_\perp$ of 333 mm, α of 0, mean momentum of 200 MeV/c and angular momentum characteristic of a 4 T field. The residual distributions in the transverse phase space variables are shown in Figure 5.6.

There is a strong correlation between longitudinal and transverse phase space coordinates. Muons with a small divergence will follow a straight track. In this case the p_z resolution is more sensitive to space point resolution and multiple scattering, resulting in a worse momentum resolution. The relation between longitudinal and transverse momentum resolution is shown in Figure 5.7.

5.3.2 TOFs

The time relative to the RF phase at which muons pass through MICE will be measured by a pair of TOF stations, TOF 1 and TOF 2. A diagram of one of the TOF stations is shown in Figure 5.8 [23]. The TOFs are each constructed from two planes of four plastic scintillator bars placed across the beamline. Muons cause the plastic to scintillate on travelling through the bars. The scintillation light passes into photomultiplier tubes and out to a DAQ system.

Each TOF plane provides 90 ps resolution. By taking an average measurement between the two, it is hoped to achieve a timing resolution of 70 ps at each station. Sources of error come from the crossing time of the muon and the response of the PMTs and DAQ. Care must be taken to ensure that environmental variations do not introduce significant variations in the transit time along cables. In addition, PMTs are sensitive to the fringe field of the spectometer solenoids necessitating magnetic shielding around the PMTs.

5.3.3 Upstream PID

Upstream of MICE it is necessary to reject beam impurities using a selection of PID detectors. Simulations of the beamline indicate that a relatively pure muon beam will be delivered into the upstream spectrometer. The primary background comes from stray pions. These are likely to decay to muons in the cooling channel which will appear as excess heating. Stray electrons which travel into MICE are likely to either scrape out of the channel or be rejected by downstream PID, appearing as a negligible excess of scraping. Upstream PID will be provided by TOF 0, TOF 1 and a pair of threshold Cerenkov counters.

The two TOFs are used to calculate the axial velocity of particles passing through them. By comparing the velocity with the momentum measured in the spectrometer, it should be possible to reject most pions. The mass reconstruction achieved by a TOF-tracker comparison is plotted against the momentum measured in the tracker in Figure 5.9 and a clear separation is visible. In addition, good electron separation will be provided by the TOFs.

The threshold Cerenkov detectors use the phenomenon of Cerenkov radiation [10] to discern between pions and muons in the high momentum region of MICE where TOF PID is more sensitive to detector resolutions as the particle time of flight is closer to the time measurement resolution. Particles travelling through a dielectric medium polarise atoms in the medium which emit radiation. If the particles are travelling faster than the speed of light in the medium, the radiation interferes constructively and can be measured. In MICE, two planes of aerogel

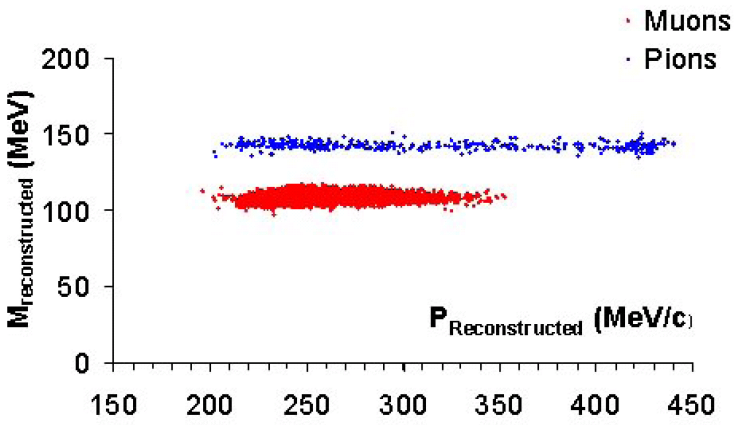

Figure 5.9: Reconstructed mass as a function of reconstructed momentum for pions and muons using the upstream MICE SciFi and TOFs [52].

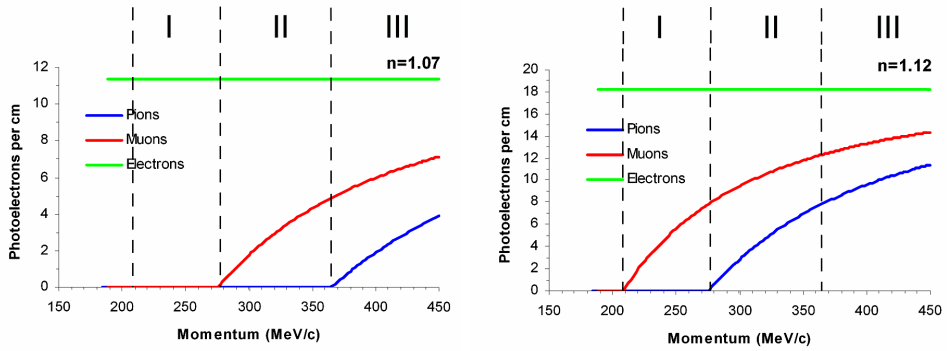

((b))

Figure 5.10: Expected photon yield for the two Cerenkov detectors: (a) n=1.07 (b) n=1.12 [51].

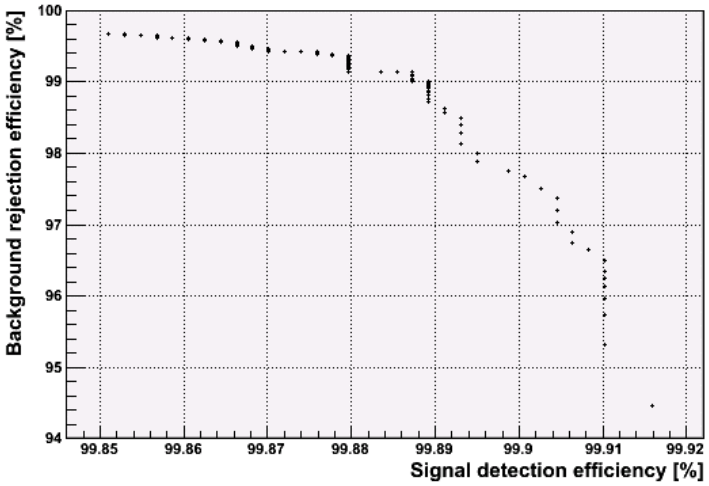

Figure 5.11: Best values for efficiency and purity of the beam after PID by the downstream PID system [53].

are used with different refractive indices of 1.07 and 1.12. The expected photon yield from the aerogel is shown in Figure 5.10 as a function of particle momentum for muons and pions. Particles that trigger only the second Cerenkov detector with momenta below 280 MeV/c are muons. Particles that trigger both detectors and have a momentum below 360 MeV/c are also muons.

5.3.4 Downstream PID

In MICE, roughly 1% of muons decay over the course of the cooling channel and these decays must be identified. In addition care must be taken not to reject muons as this may bias the cooling measurement. Downstream PID is achieved primarily by examining the time of flight between TOF I and TOF II. In addition, an electron-muon calorimeter (EMCAL) has been devised to provide separation of electron and muons downstream.

Decay electrons effectively travel at the speed of light, c, whereas muon velocities lie in the range $0.8c$ to $0.9c$. This enables good separation of muons from electrons in MICE by comparing the velocity calculated in TOF 1 and TOF 2 with the momentum measured in the downstream spectrometers, unless muons decay near to TOF 2 or electrons suffer strong deflection in the magnetic fields. In the case that this system fails, additional separation is provided by the EMCAL. Here, the energy loss characteristics of particles in lead interspersed with plastic scintillator is used to distinguish between muons and electrons. The simulated downstream PID detector performance is summarised in Figure 5.11.

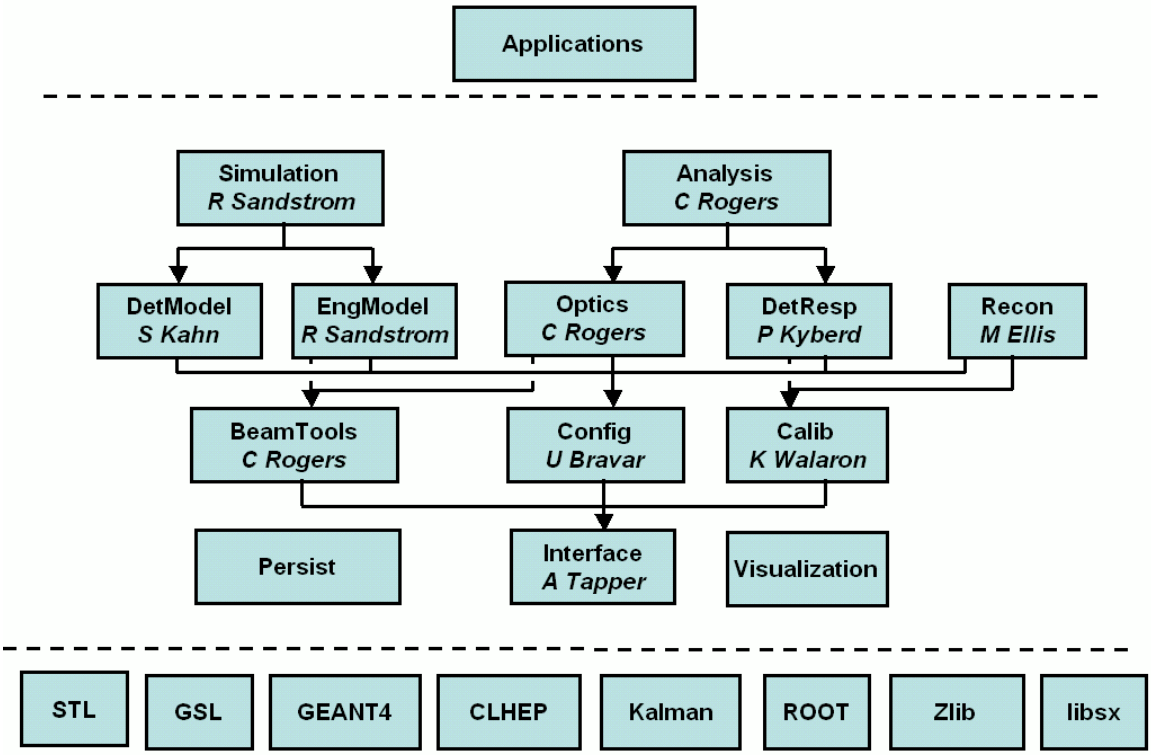

Figure 5.12: Domains in G4MICE. Domains can use libraries below them in the diagram if they, or their children, are connected. Package managers, responsible for ensuring that code is appropriately tested and documented in each domain, are shown in italics.

5.4 MICE Software Tools

Simulation, reconstruction and analysis of the MICE channel and detectors is performed using the G4MICE software package. Simulation in G4MICE is based on GEANT4 [55], a tracking and physics process modelling package. Reconstruction is based around the Kalman linear fitting package [57]. The beamline will be simulated using G4Beamline [58], also based on GEANT4, and Decay Turtle, a third order beam optics package [59]. In addition, the ICOOL software package [60] [61] is used occasionally for simulating MICE and also in the simulation of other Neutrino Factory components.

5.4.1 G4MICE Framework

G4MICE uses an object oriented methodology based on C++ that enables different areas of the software to be developed in parallel and independently of other development and in principle can support a very large number of developers simultaneously. The package is divided up into several different domains shown

in Figure 5.12. Each domain is compiled as a separate library enabling a faster compile and test cycle.

User Interface

The Interface, Calib, Config, RealData, Persist and Visualisation packages handle input control variables, input geometry information and tracking data between different modules. Run control is handled by a datacard text file that controls global data and MiceModule text files that handle geometry and object specific data. Experimental data will be handled by the RealData package. Graphical visualisation of the experiment is performed by the Visualisation package.

Simulation

The BeamTools, EngModel, DetModel and Simulation packages control simulation of particles through MICE. BeamTools models solenoids, quadrupoles and RF fields. EngModel handles modelling of physical volumes. DetModel handles modelling of detector volumes and controls the data that is output for simulation of the detector response. Simulation controls the interface with GEANT4 for run control, physical processes and tracking.

Detector Response

The DetResp package handles simulation of the detectors and electronics based on tracking information from the Simulation module. For example, information on the energy deposited in the TOFs from the tracking simulation is used to model the photon yield and calculate the electronics response.

Reconstruction

The Recon package uses information either from DetResp or from experimental data to reconstruct particle tracks and other data for analysis. For example, the electronics response modelled in DetResp can be used to calculate space point hits in the SciFi and these hits combined to calculate track momentum and position at some plane perpendicular to the z axis.

Beam Analysis

The Optics and Analysis packages provide a library for calculating beam parameters from the electromagnetic fields and output particle tracks respectively. The Optics package calculates first order transfer maps in quadrupole, solenoid and RF fields as well as modelling energy loss and multiple scattering effects in

materials. Either covariance matrices or single particles can be transported. In addition, a quick, simplistic single particle transport algorithm has been implemented by integration of the equations of motion.

Applications

The Applications package provides a convenient framework to write user-defined applications using the G4MICE libraries. Users can select which libraries are included in the application, leaving the compiler blind to other libraries improving the time taken to compile and execute applications.

5.4.2 Other Software Packages

G4Beamline provides a simple simulation for the tracking of particles through an accelerator system. It is not used for the MICE channel as it does not provide the level of sophistication and detector simulation required for the MICE simulation. On the other hand, it provides a user friendly interface that allows swift redesign convenient for the beamline simulation.

ICOOL operates on a similar principle to G4Beamline, also offering a relatively user friendly interface for particle tracking through an arbritrary accelerator system. ICOOL is based on a modified version of GEANT3. The software also comes with a built-in beam analysis tool, ECALC9, which calculates several beam parameters of interest. ICOOL has the advantage that it has been used for software simulation for several years and is well tested.

5.5 Operation of MICE

MICE will be constructed in 6 steps to enable calibration of components. The MICE steps are shown in Figure 5.13. In the first step, the PID detectors will be calibrated and the beamline will be commissioned. Subsequently the first spectrometer will be placed in the beamline enabling the match quality between the beamline and the MICE channel to be optimised. In step 3 the second spectrometer will be placed in the beamline enabling calibration of the two detectors by comparing measurements of particles travelling between them. The first cooling measurement will be made in step 4, where the first Absorber and Focus Coil (AFC) module will be placed between the two spectrometers. This is expected to be the first observation of ionisation cooling. In step 5 a half cell of the cooling channel will be constructed. A RF and Coupling Coil (RFCC) module will be placed between two AFC modules enabling a study of longitudinal and transverse

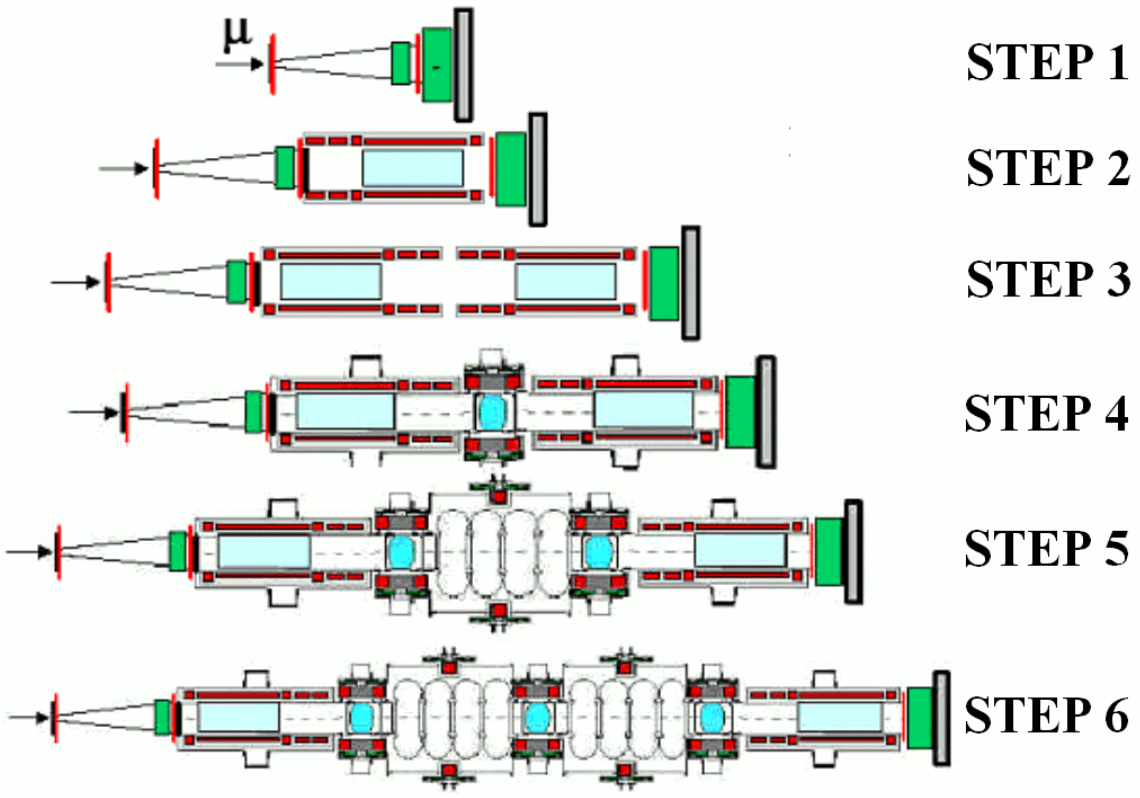

Figure 5.13: The 6 steps of MICE. Initial steps are for calibration of cooling channel elements. Cooling will be demonstrated in later steps [23].

beam dynamics. Step 6 will be made up of three AFC modules and two RFCC modules.

Step 6 will operate at central momenta between 140 and 240 MeV/c over an emittance range between 10 mm and 2.5 mm, which is roughly the equilibrium emittance for MICE running with a β function of 420 mm in the absorbers. The β function in the absorbers will be varied between 70 mm and 420 mm by changing the coil currents and muon momenta with suitable input beam parameters to control the amount of cooling achieved.

In this thesis, Step 6 will be considered in detail with currents set in flip mode, central momenta of 200 MeV/c and a β function in the absorbers of 420 mm. This will be used as an example of the simulation of MICE to examine simulation performance, cooling performance and measurement errors.

Chapter 6

G4MICE Field Modelling Library

Three libraries have been written within the G4MICE framework for simulating and studying beam dynamics in MICE. In this chapter the BeamTools library is discussed. The BeamTools library generates RF cavity, quadrupole and solenoid fields for use by the GEANT4 stepping algorithms and the G4MICE Optics and Analysis libraries.

The BeamTools package is made up of several classes that provide utilities such as field placements and phasing and several classes that calculate the actual fields. The solenoid and magnetic field map components were originally written by Daniel Elvira. The quadrupole component was originally written by Tom Roberts. However, the code has been largely rewritten by the author of this thesis to make the code maintainable and reliable.

6.1 Object Oriented Approach

A Unified Modelling Language (UML) [56] diagram displaying the classes in the BeamTools library and important external classes is shown in Figure 6.1. Classes are abstract representations of objects in G4MICE. So, for example, a BTSolenoid class is an abstract representation of a current carrying coil, which can be *instantiated* as a specific object with parameters such as current density, length and thickness.

The BeamTools package has been written with an object oriented approach in mind to make the library robust and easily extended. The class structure enables new types of field to be added with minimal knowledge of the library's structure and minimal interference in the workings of existing objects.

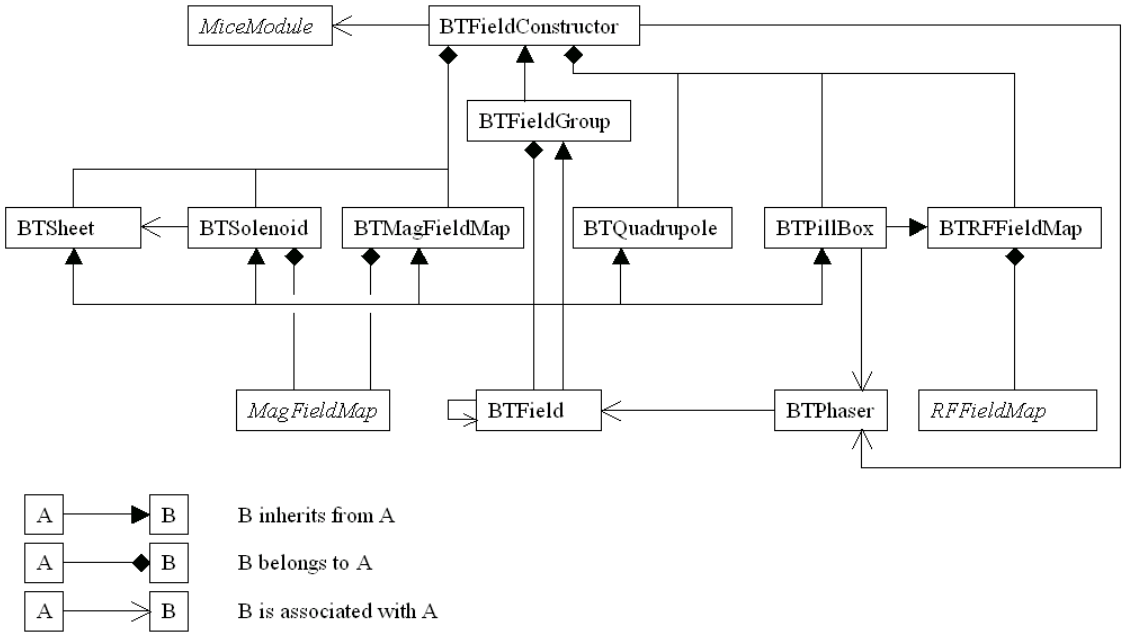

Figure 6.1: A UML diagram of the BeamTools package. Classes external to BeamTools are shown in italics.

6.1.1 BTField

The BTField class provides an Abstract Data Type (ADT) from which all other field types in the library inherit. The BTField is said to be the parent class, while other field types are said to be the child classes. Children inherit many of the parent's properties. This enables other objects to interface with a generic field object without knowing what type of field is being used. For example, any class that requires the field value from a BTField can call BTField::GetFieldValue, without knowing whether the object is a BTSolenoid, BTPillBox or some other type of field. The relevant function in the child object will be called automatically unless the function is not defined, in which case the parent function is called instead.

6.1.2 BTFieldGroup

The BTFieldGroup provides a good example of the convenience of this abstraction. This class is used to place fields and is itself a BTField, so can be placed within a BTFieldGroup. Without the ADT it would be necessary to write code for each field type to place fields, recover field values and perform other operations. Under this framework it is only necessary to write code once and any additional field types that are required can be added with ease.

6.1.3 BTFieldConstructor

The BTFieldConstructor is a BTFieldGroup that provides an interface between the user commands, GEANT4 and BeamTools. The BTFieldConstructor converts the user-supplied placement information in the MiceModule into a BTFieldGroup and then passes this information to GEANT4.

6.2 Solenoids

Transverse focussing in MICE is provided by solenoids so accurate simulation is important. Several different algorithms have been proposed for the calculation of the field produced by solenoids [63]. In G4MICE, solenoid coils are represented by a number of infinitely thin current carrying cylindrical surfaces distributed evenly through the coil volume. The simulation model will eventually be compared with measurements of the solenoid fields.

To improve processing speed the field is by default calculated at a number of points on a rectangular grid in (r, z), and then calculated at a general position by interpolation between grid points. There is no θ dependence due to the cylindrical symmetry of the problem. In addition, the grid can be generated by an external program and read into G4MICE or the field can be calculated directly from the sheets with a significant increase in processing time.

6.2.1 Sheet Model

A schematic of the sheet model is shown in Figure 6.2. The field from a single current sheet of length $2L$ and radius a at some point (r, z) is given by [62] [63]

$$B_z(r,z) = b_z(r,z+L) - b_z(r,z-L), \tag{6.1}$$

$$B_r(r,z) = b_r(r,z-L) - b_r(r,z+L), \tag{6.2}$$

where

$$b_z(r,z) = \frac{\mu_0 I'}{\pi}\frac{za}{\zeta(a+r)}\Big[K(k) + \frac{(a-r)}{2a}(\Pi(k,c) - K(k))\Big], \tag{6.3}$$

$$b_r(r,z) = \frac{\mu_0 I'}{\pi}\frac{\zeta}{4r}\Big[2(K(k) - E(k)) - k^2 K(k)]. \tag{6.4}$$

Here

$$k = \sqrt{\frac{4ar}{(a+r)^2+z^2}}, \tag{6.5}$$

$$\zeta = \sqrt{(a+r)^2+z^2}, \tag{6.6}$$

$$c = -\frac{4ar}{(a+r)^2}, \tag{6.7}$$

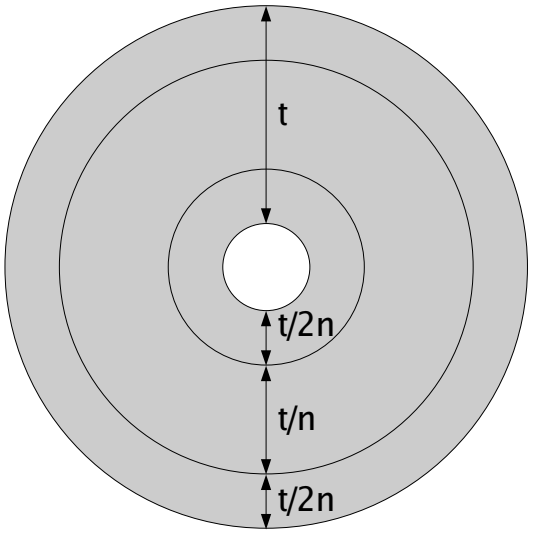

Figure 6.2: Schematic of a solenoid of thickness t and current density J constructed from n current sheets. Sheets are separated from the inner and outer radius by $t/2n$ and from each other by t/n. Each sheet is simulated with a current density of J/n.

and K, E and Π are complete elliptic integrals.

The field is identical to the field from a solenoid with finite thickness and a constant current density in the limit that the number of sheets tends to infinity.

6.2.2 Model Accuracy

Errors in this model arise from three sources: firstly, close to the sheet the field given above becomes singular; secondly if few sheets are used the field will be quite different to the field generated with a large number of sheets; and thirdly the interpolation will introduce an error on the field calculation. The error arising from these sources is examined for a single MICE focus coil, with length 210 mm, inner radius 260 mm, thickness 84 mm and carrying a current density of 113.95 A/mm^2.

The analytical value of the field strength close to a current sheet is examined in Figure 6.3. It will be noted that the field is convergent up to a fraction of a micron from the current sheet, when the calculation becomes divergent. This necessitates care when choosing grid coordinates.

In Figure 6.4 is shown the maximum discrepancy between the field calculated with a few sheets and the field calculated with 100 sheets within the Focus Coil bore, here taken to be 200 mm. This gives a good estimate of the maximum error introduced by using fewer sheets. As the calculation of fields is usually only made once, it is sufficient to set the number of sheets such that this maximum error is well below the error introduced by the interpolation algorithm.

In addition the error in the calculation of the on-axis field is shown. As shown above, the on-axis field is important in beam optics calculations. In this case, it is sufficient to set the number of sheets such that the maximum error is well below the change in field over a step in any numerical integration of the transfer map.

6.2.3 Grid Interpolation Algorithm

The solenoid field is interpolated from a pair of two dimensional grids $B_z(r_i, z_j)$ and $B_r(r_i, z_j)$; $B_\phi = 0$ due to the symmetry of the problem. A third order polynomial fit is applied on the z-axis to determine the field at (r_i, z) and (r_j, z). A linear interpolation is applied across r to determine the field at (r, z).

The accuracy of this fit is found by comparing the analytical value of the field strength with the interpolated value. The grid spacing was held fixed at 0.25 mm in z and varied in r; and held fixed at 0.25 mm in r and varied in z. The result is shown in figure 6.5. Here the field was measured randomly at 10000 points

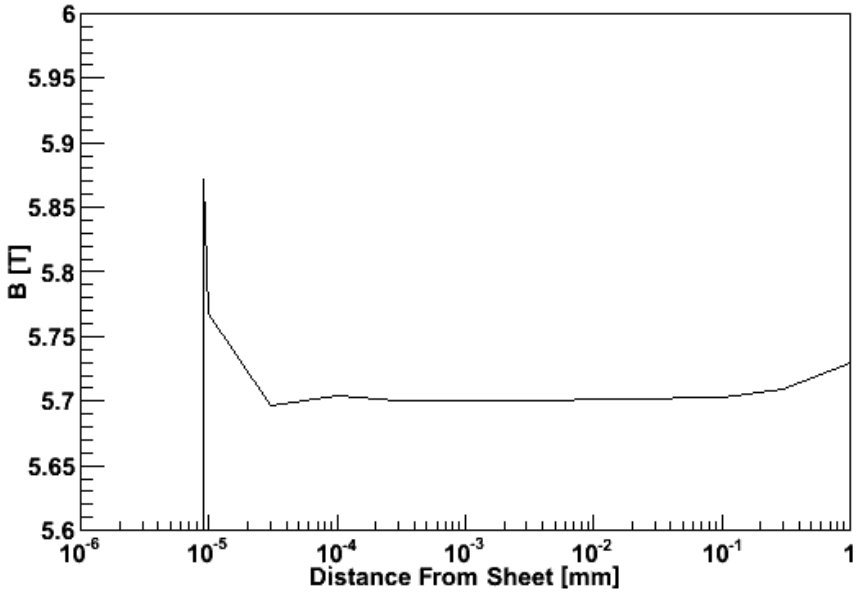

(*a*)

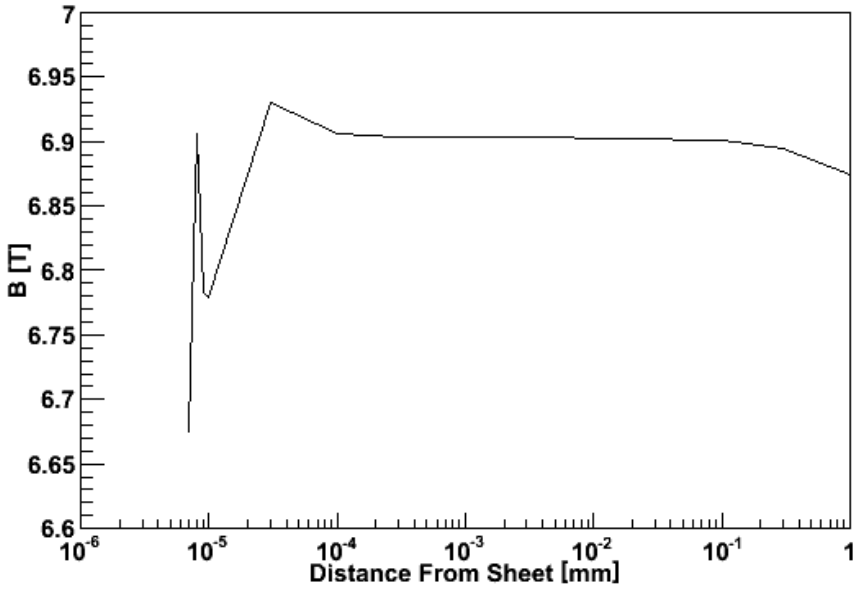

(*b*)

Figure 6.3: The axial field as a function of distance from the sheet: (a) outside of the sheet; (b) inside of the sheet.

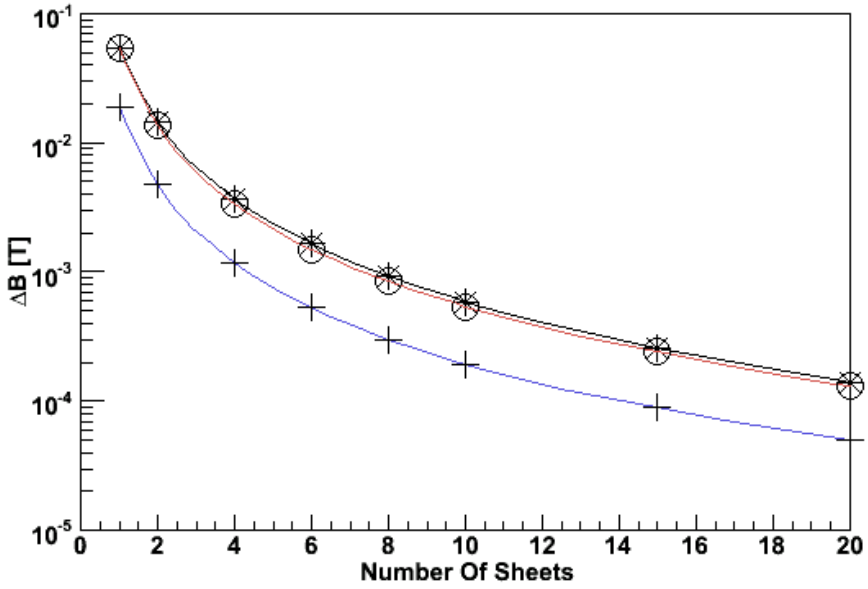

Figure 6.4: The maximum error in the field calculation as a function of the number of sheets used: (o) $B_r(r,z)$; (*) $B_z(r,z)$; (+) $B_z(0,z)$.

within a region extending 500 mm from the coil centre in z and 200 mm from the coil centre in r for each of the grid spacings; the error is taken as the mean of the difference between the analytical value and the true value. The field map is chosen to have a length of 4800 mm and a radius of 550 mm.

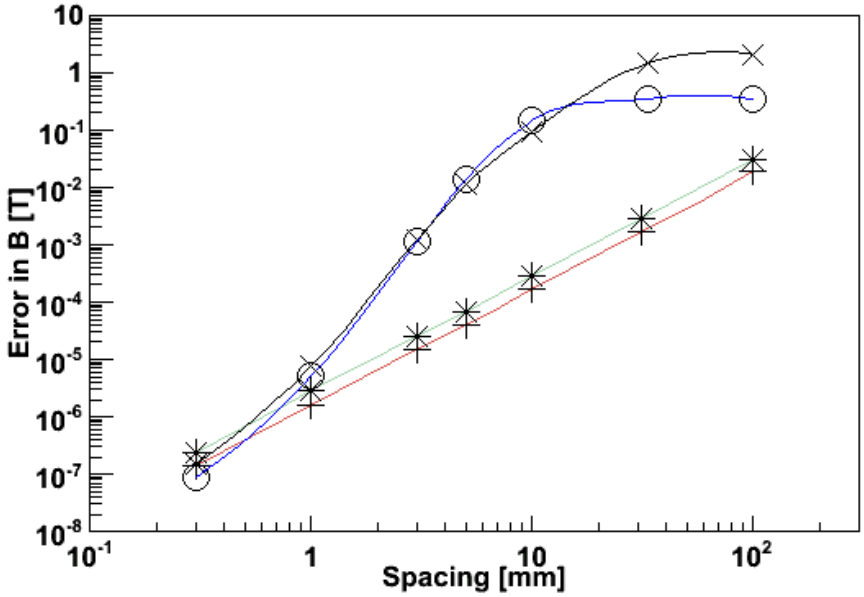

Figure 6.5: The mean errors on field components as a function of grid spacing: (o) B_r against r; (x) B_z against r (*) B_r against z (+) B_z against z

6.3 RF Model

The MICE RF cavities operate in the TM010 mode. Two methods for simulating the fields are available: firstly, the field can be calculated using the analytical model for a pill box standing wave: secondly, the field can be interpolated from a cylindrically symmetric field map.

6.3.1 Pill Box Fields

The field is calculated using the analytical solution for a standing wave in a cylindrical pill box cavity [34],

$$E_z = E_{max} J_0(2.405 \frac{r}{r_{cav}}) sin(\omega t - \phi_{ref}), \quad (6.8)$$

$$B_\theta = \frac{E_{max}}{c} J_1(2.405 \frac{r}{r_{cav}}) cos(\omega t - \phi_{ref}), \quad (6.9)$$

with all other field components 0. Here E_{max} is the peak electric field, J_0 and J_1 are Bessel functions, ω is the angular frequency of the RF cavity and ϕ_{ref} is the phase of the reference trajectory.

It is also possible to set the cavities to give an accelerating electrostatic field across the pill box. This is non-physical as such high field gradients would cause sparking across the cavity, but provides a convenient model for the occasions where precise phasing and beam dynamics are not important.

6.3.2 RF Field Maps

The RF cavities in MICE are built with non-parallel sides to reduce break down. Thin Beryllium windows enable a larger aperture, so less scraping, without perturbing the RF field. These effects cause aberrations in the electric field away from the pill box model. As an alternative to the analytical model, field maps generated in SuperFish [64] can be used in MICE. As in the case of solenoids, these field maps use a linear interpolation in r and a third order polynomial fit in z to calculate the fields away from the grid points.

The geometry and fields of the pill box and MICE cavities are shown in Figure 6.6, as generated in SuperFish. There are four such cavities in each of the MICE RFCC modules.

The accuracy of the fit was found by generating electromagnetic fields on a rectangular grid with grid spacing varied in r and z. The error on the interpolation was found by comparing this to a fit against a field generated on a grid with spacings of 0.02 mm in z and 0.15 mm in r. The grid spacing was varied in r keeping the z spacing constant and in z keeping the r spacing constant. Note

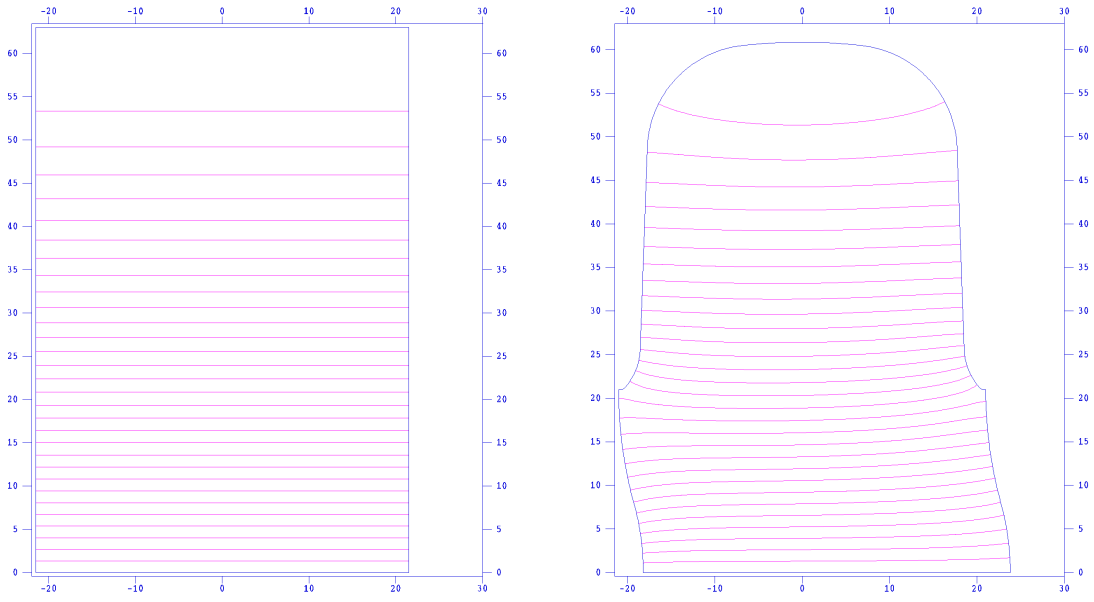

Figure 6.6: A rectangular pill box cavity and a realistic MICE RF cavity generated by SuperFish.

that the field was generated in SuperFish using a triangular mesh with spacings of approximately 2 mm, and at finer granularity SuperFish applies some interpolation of its own, so at fine granularities the accuracy of the G4MICE fit to the SuperFish fit is shown rather than the fields per se. The fitting accuracy is shown in Figure 6.7.

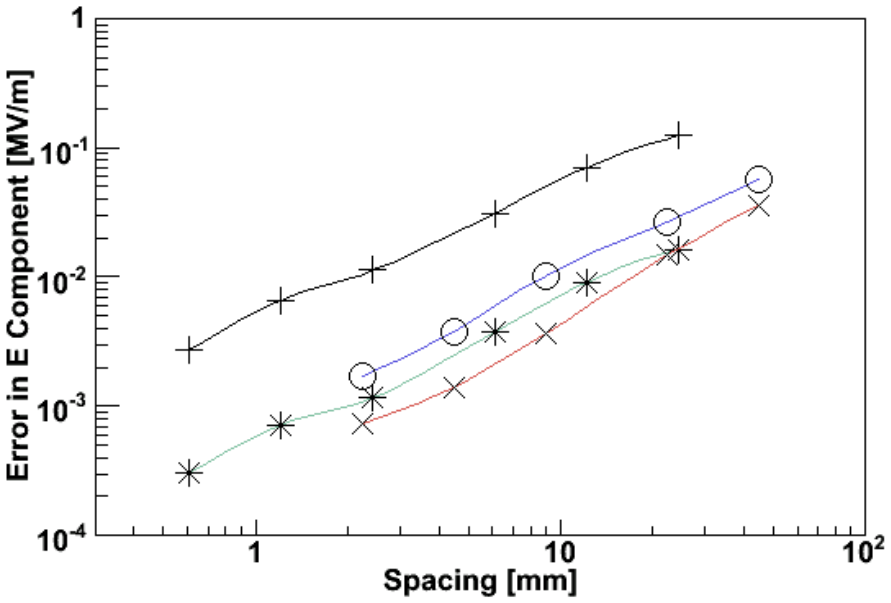

Figure 6.7: The mean errors on electric field components as a function of grid spacing: (o) E_r against r;(x) E_z against r; (*) E_r against z; (+) E_z against z.

6.3.3 Phasing

RF cavities can have their synchronous phase set either by the user, or by use of a reference particle. In the latter case, this is achieved by the following algorithm:

1. The cavity phase is set to some arbitrary value.
2. The time and energy of the reference particle is recorded at the upstream edge of the cavity.
3. The reference particle is passed through the RF cavity once with an electrostatic field.
4. The cavity phase is set to the time at which the reference particle passes through the cavity centre.
5. If the difference between the old phase and the new phase lies within some user-defined phasing tolerance, the cavity phase is left at this value.
6. If it lies outside the tolerance, the reference particle is fired again from the start of the cavity, with the phase set at the value from 3.
7. Step 3 to 5 are repeated until the cavity phase lies within the user-defined tolerance.

The phase of the cavity is observed to reach machine precision very quickly, typically within 3 or 4 iterations. This is perhaps to be expected. The convergence of the phasing is governed by the transit time of the muon across the cavity and the velocity change of the muon. If the muon crosses the cavity very quickly relative to the RF frequency, then the cavity field will be approximately constant. If the velocity change of the muon is small, then the effect of the change in the cavity field as the muon crosses the cavity will have a negligible effect on the time at which the muon reaches the cavity centre. In the case of MICE, the cavity passes through about a fifth of an RF cycle during the transit of a muon, but the muon velocity changes by only 1% in the cavity. Even if the muon is completely out of phase initially, the first estimate of the RF phase of the cavity will only be erroneous by 1%, and so subsequent estimates are likely to be highly accurate.

6.4 Quadrupoles

The aim of G4MICE is to simulate the MICE detectors and the components in between. Although the cooling channel itself is made up of solenoids, the muon

beamline contains quadrupoles. In order to simulate the combined particle identification performance of TOF0, TOF1 and the Cerenkov detector, it is necessary to track particles through the quadrupole fields. Two models are presented in G4MICE; under the hard edged model, fringe fields are taken to be negligible while under the Enge model, fringe fields are modelled using a fit of the form [30]

$$F(z) = \frac{1}{1 + exp\Big[\Sigma a_i (z/D)^i\Big]} \tag{6.10}$$

where D is the aperture of the quadrupoles and a_n are fit parameters. The field gradient of the quadrupoles is shown in Figure 6.8 for the default parameters in G4MICE. These parameters were originally measured from the Positron Electron Project (PEP) quadrupoles [30]. The measurement of these fringe field parameters has been proposed for the MICE beamline quadrupoles.

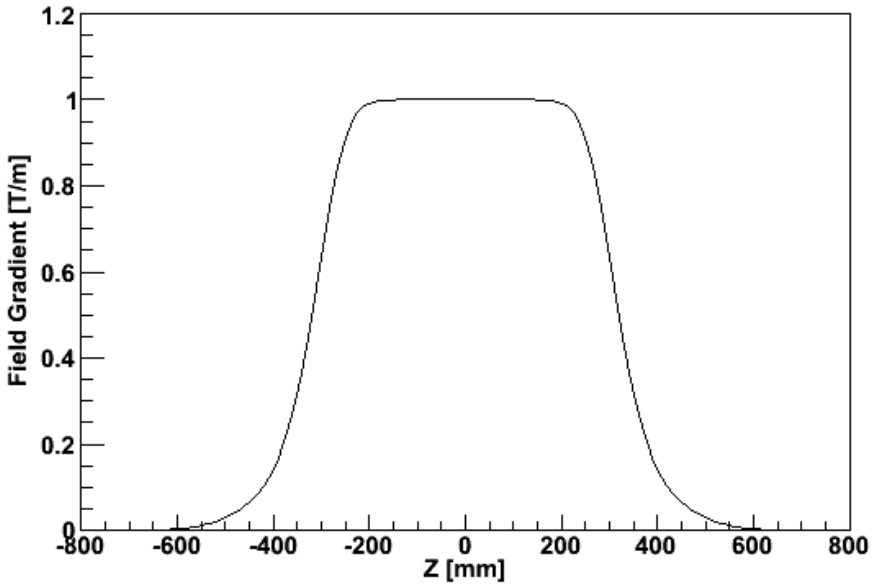

Figure 6.8: Field gradient for the default fringe field parameters, normalised to 1.

6.5 Field Tracking Performance

The tracking performance of the various field components is now examined by comparing tracking performed with a very fine map and tracking performed with maps of varying granularity. For these purposes there will be no material present in the cooling channel.

6.5.1 Solenoid Grid Size

In solenoids, the dominant error is usually chosen to be the grid size. Requiring a finer field map slows the processing on each track while requiring a larger number of sheets only slows the processing during the generation of the field map. As the field map can be reused, this is not expected to be a significant factor in tracking performance.

To study the effect of grid size on tracking accuracy, a bunch of 1000 particles was tracked through the MICE Step 6 fields from $z = -5500$ mm to $z = 5500$ mm, where $z = 0mm$ corresponds to the centre of the central absorber. This point was chosen as it is well inside the constant field region of the solenoid. The beam was chosen so that β is constant through the tracker solenoid field and matched to the SFoFo lattice.

The number of points in the field grid of each solenoid was varied for r and z. The radial position and transverse momenta of each track was compared with the radial position and transverse momenta of tracks with 250 and 500 points in r and z respectively. In Figure 6.9 the mean of the difference in r and p_t, momentum transverse to the beam axis, is shown.

For tracking in MICE, the error is required to be significantly smaller than the experimental resolution. The tracker resolution is 0.5 mm in x, 0.4 mm in y, 2.0 MeV/c in p_x and 1.5 MeV/c in p_y. By choosing to use 100 grid points in z and 150 grid points in r for each solenoid, the error introduced by the solenoid grid is constrained to be significantly less that 10^{-2} mm and 10^{-2} MeV. This grid spacing is equivalent to a 32 mm spacing in z and a 5.5 mm spacing in r for the focus coil outlined above, indicating that errors of order 10^{-3} in $\vec{B}$ give an appropriate tracking accuracy and 10 sheets or greater should be used in the generation of field maps.

6.5.2 RF Grid Size

In RF cavities, the only source of error arises from the choice of grid size. The effect of the grid size on tracking was studied in a similar manner to that of solenoid grid size above and the results are shown for transverse and longitudinal phase space variables in Figures 6.10 and 6.11.

As the RF field affects both longitudinal and transverse phase space, the bunch was chosen to have a standard deviation of 10 MeV and 0.3 ns. These values are smaller than those seen in a Neutrino Factory front end, but muons at larger amplitudes that sit outside the RF bucket are not of interest in the study of cooling. The cavities were simulated running on-crest as this is the MICE

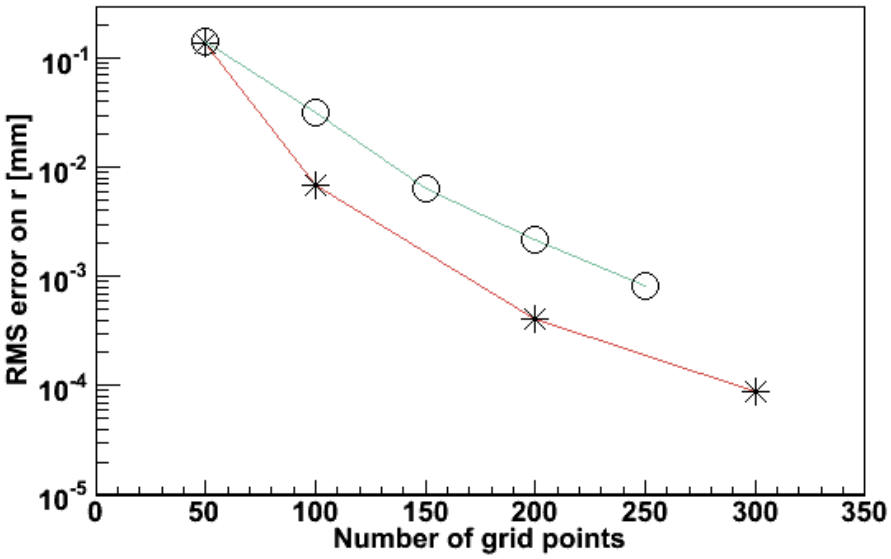

(*a*)

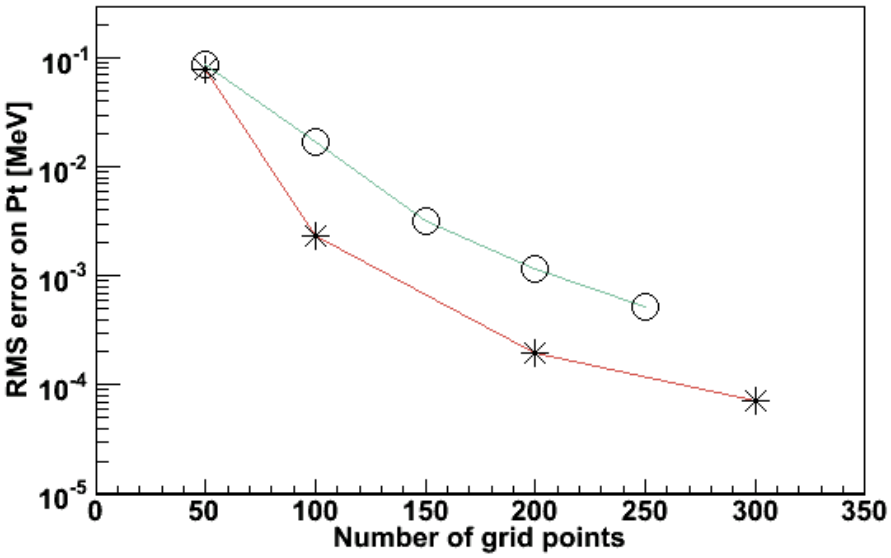

(*b*)

Figure 6.9: Mean difference in (a) r and (b) p_t as a function of number of grid points on each solenoid: (o) number of grid points in r; (x) number of grid points in z.

baseline mode and the mode where cavities will affect the tracking of the bunch to the greatest extent. In this mode it is not possible to match the longitudinal phase space; instead the bunch was chosen to initially have no correlation between energy and time.

Again the RF grid size was varied. In this case the error in energy, time, transverse momentum and transverse position were examined by comparison with tracking performed with an RF field map with 2000 points in z and 400 points in r. The field map had a radius of 608 mm and a length of 448 mm.

The energy resolution of the tracker is 3.5 MeV and the time resolution of the Time Of Flight counters is 70 ps. Relative to the experimental resolution, the dominant error is in r where 100 radial grid points and 500 longitudinal grid points give errors an order of magnitude less than the tracker resolution.

6.6 GEANT4 Stepping

The computational model used by GEANT4 is described in the GEANT4 user manual [65]. An experiment is represented in GEANT4 as a series of volumes made up of a particular material. Tracking is performed by integrating the equations of motion numerically for each muon. GEANT4 terminates the step if a process takes effect, such as multiple scattering or particle decay, or if a track passes through a volume boundary. If a track intersects a volume boundary, the position of the track is estimated at the boundary of the volume and stepping continues within the new volume.

GEANT4 offers three principal parameters to set the tracking accuracy: DeltaOneStep; DeltaIntersection; and the Miss Distance. These set the error on a step within a volume, the error on stepping into a new volume and the error on the calculation of whether a track clips another volume.

Integration Error

The parameter DeltaOneStep, denoted δ_1, is used by GEANT4 to determine the error on a single step within a volume. The default tracking algorithm uses fourth order Runge Kutta numerical integration of the equations of motion. The error on a step is estimated using a Richardson extrapolation [66]. The estimated error on the position and momentum vectors, **dq** and **dp**, are required to obey $\mathbf{dq}^2 < \epsilon_q^2$ and $\mathbf{dp}^2 < \epsilon_p^2$ where (ϵ_q, ϵ_p) are the allowed error on a step, related to δ_1 by $(\epsilon_q, \epsilon_p) = (\delta_1, \delta_1/h_{step})$ where h_{step} is the initial estimate for the step length used. Additional parameters EpsilonMax and EpsilonMin limit the maximum and minimum allowed values of ϵ_{err}. GEANT4 will shorten its step length until

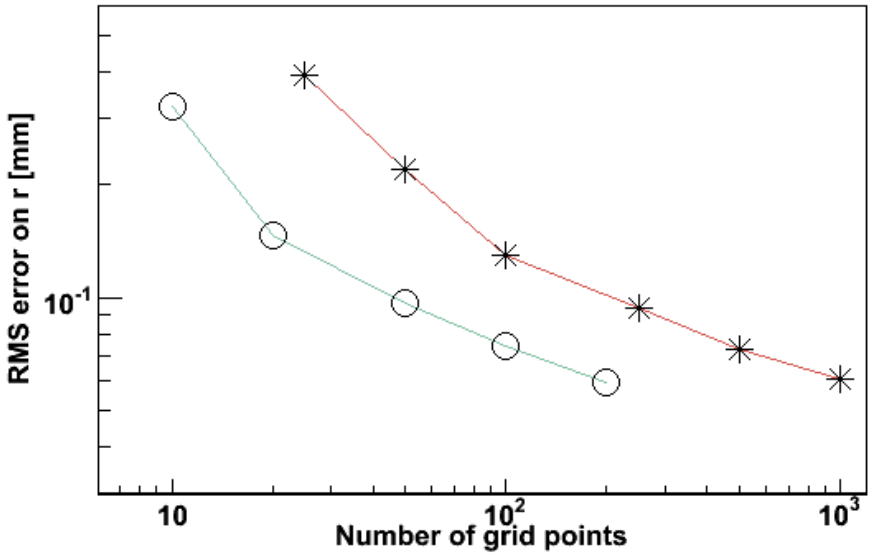

(*a*)

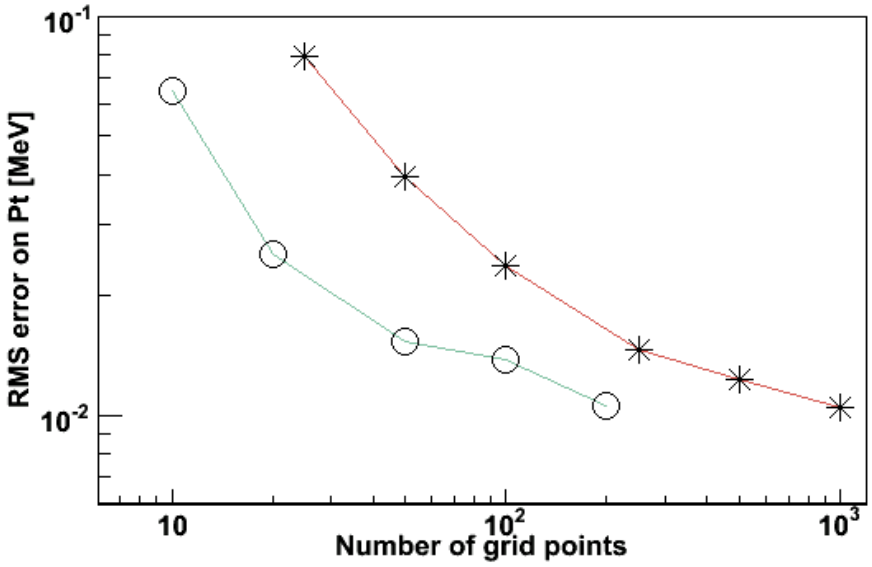

(*b*)

Figure 6.10: Tracking error in (a) r (b) p_t as a function of number of grid points in each RF cavity: (o) number of grid points in r; (x) number of grid points in z.

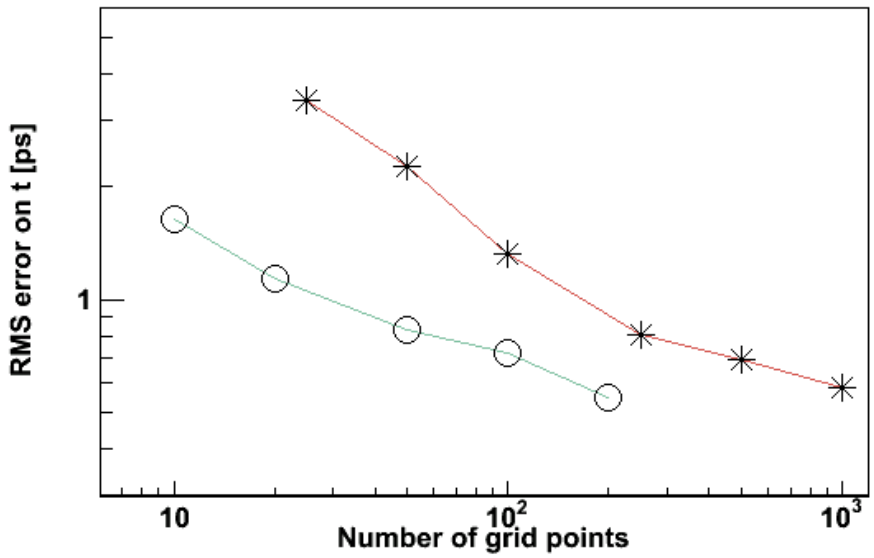

(*a*)

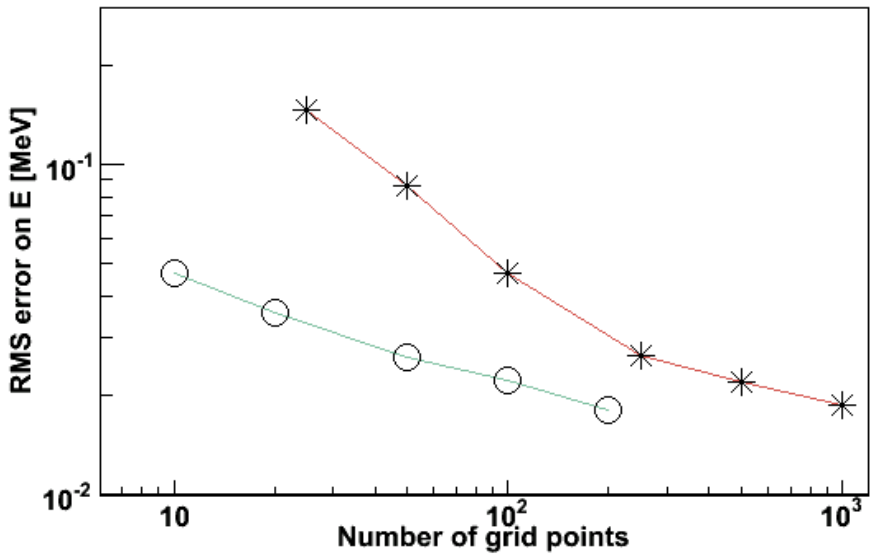

(*b*)

Figure 6.11: Tracking error in longitudinal phase space variables (a) t (b) E as a function of number of grid points in each RF cavity: (o) number of grid points in r; (x) number of grid points in z.

errors obey these requirements. If GEANT4 cannot step with the accuracy required by ϵ_{err} after 100 iterations, it performs a step with the step length at the last iteration.

Other numerical integration algorithms have been implemented that offer varying levels of performance. In general, higher order integration offers better performance in slowly changing fields. In all cases errors are estimated using a Richardson extrapolation [67] and limited by the same constraints outlined above. In this thesis, numerical integration will be performed using the fourth order Runge Kutta algorithm only.

DeltaIntersection

When a track is found to have crossed a boundary, the position of the track end point at the material boundary is estimated with an accuracy determined by the parameter deltaIntersection. A smaller parameter value makes the estimation of the position on the material boundary more accurate.

Miss Distance

When GEANT4 calculates whether a track has crossed a material boundary, it approximates the trajectory as a straight line between the start and end points and examines whether this straight line has crossed any boundary. The maximum distance between the true path and linear interpolation is known as the miss distance. A small miss distance improves the estimation of volume boundaries where tracks clip the corner of a volume. It also decreases the step length.

In MICE, tracks are not expected to clip very often. Volume boundaries are approximately perpendicular to the beamline except at the scraping aperture. The dominant effect of the Miss Distance parameter is then expected to be to limit the step size.

6.6.1 Stepping Accuracy

The behaviour of GEANT4 tracking is shown in Figure 6.12 for the position and time of muons tracked through the MICE Step VI fields. The full MICE cooling channel representation was used, but with all materials set to vacuum. In this way an accurate representation of the volume boundaries is used without material processes interfering with the tracking. It should be noted that in this case multiple scattering off of the MICE vacuum, which is modelled with small but finite density, limits the maximum step length to a few centimetres.

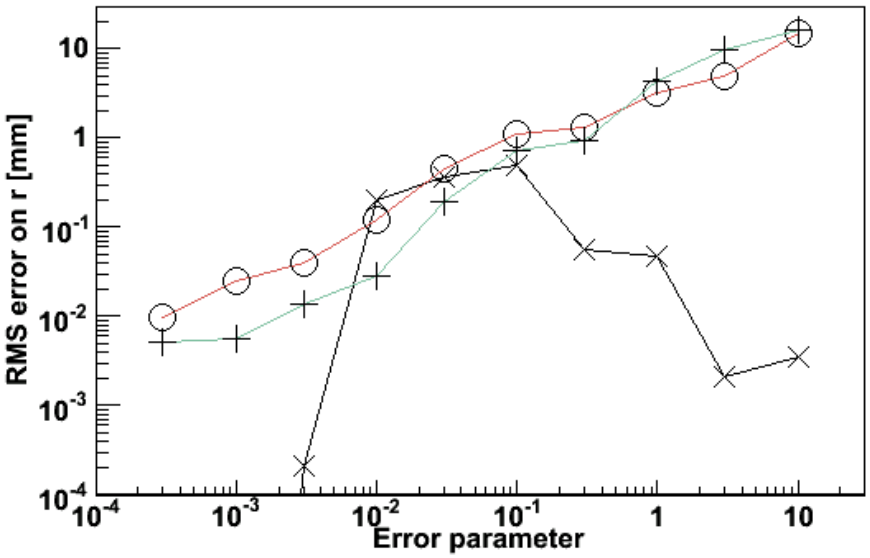

(*a*)

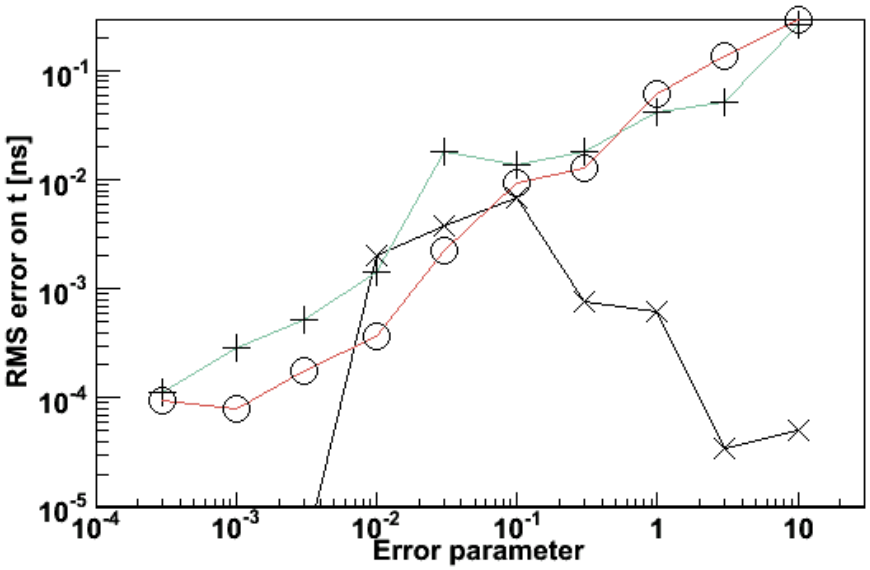

(*b*)

Figure 6.12: Accuracy of (a) r (b) t against GEANT4 parameters after tracking through the MICE fields. All other parameters were held constant while the following parameters were varied: (x) deltaOneStep; (o) missDistance; (+) deltaIntersection.

A matched beam was tracked through the MICE fields and deltaOneStep was varied in steps of a factor $\sqrt{10}$ with epsilonMin set to 10^{-10} and epsilonMax set to 10^{10}. DeltaIntersection and missDistance were set to 10. The error introduced during tracking was calculated by taking the difference between the position of a muon tracked from -5500 mm to 5500 mm with a particular value of deltaOneStep and the position tracked for the next smallest value of deltaOneStep. So, for example, r was calculated for deltaOneStep set to 10 and deltaOneStep set to $\sqrt{10}$ and the RMS of the difference was used to estimate the error on the stepping. This procedure was followed for each step size in order to estimate the stepping error. 1000 muons were tracked through the MICE apparatus. 100 muons were used to estimate the RMS value of dr and the tracking error was calculated for the other 900 muons with a cut at 10 RMS on dr to remove outliers.

The behaviour is shown in Figure 6.12 and the stepping does not converge in a straight forward manner. This behaviour was found to be independent of deltaIntersection, but highly dependent on missDistance. A small value of missDistance tends to limit the error on tracking for larger values of deltaOneStep; while for small values of deltaOneStep and missDistance, the same behaviour was observed as shown in Figure 6.12.

Similar plots are shown for the parameters missDistance and deltaIntersection. MissDistance and deltaOneStep were set to large values and then one or the other was varied. As before the error was calculated by comparing muons tracked with different precision. In these cases the tracking error converges, although the missDistance tracking error is not monotonic. The dominant error is once more in r, where values of deltaIntersection less than 1e-4 and missDistance less than 3e-4 give tracking errors of order 10 microns.

In order to understand the dependency of the stepping accuracy of δ_1 on missDistance, the procedure outlined above was repeated for several different values of missDistance. Results are shown in Figure 6.13. In this case, each point represents only 100 muons. DeltaIntersection was set to a 10^-4 for this simulation. DeltaOneStep is seen to be quite dependent on the missDistance, but there is no evidence of convergence for different values of deltaOneStep.

6.7 Summary

The field model of the RF fields, solenoids and quadrupoles in G4MICE have been detailed. The tracking performance has been studied compared to the grid size of the solenoid field map, RF field map and various tracking parameters. The tracking parameters that give tracking performance much less than the detector

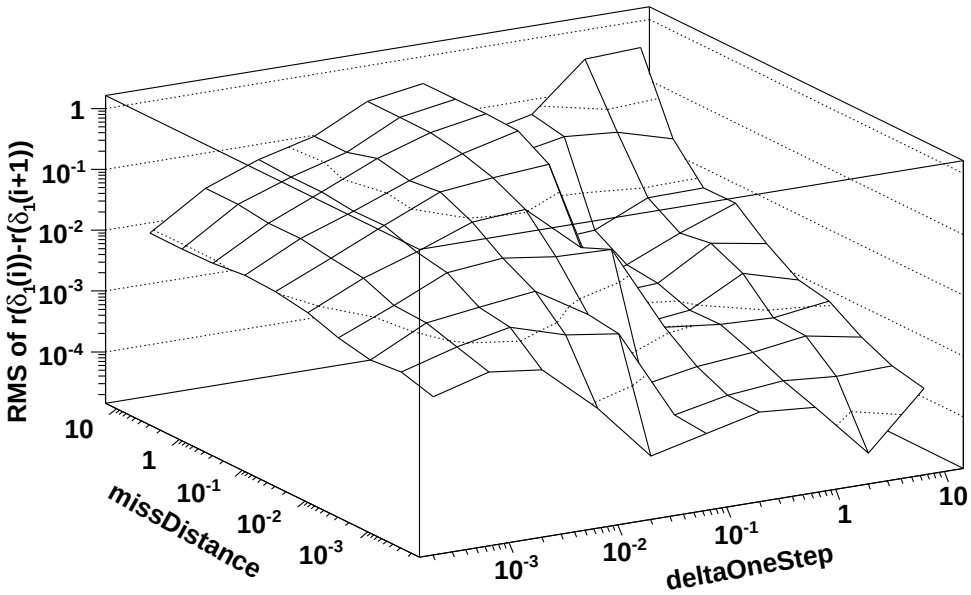

Figure 6.13: Accuracy of r as a function of deltaOneStep for different values of missDistance.

Parameter	Value
Solenoid number of sheets	10
Solenoid number of Z points	100
Solenoid number of R points	150
RF number of Z points	2000
RF number of R points	400
DeltaOneStep	3*
DeltaIntersection	1e-4
MissDistance	3e-4

Table 6.1: Tracking parameter settings for tracking errors less than 10 microns. * Convergence not found for DeltaOneStep.

resolution are listed in table 6.1. In all cases, the dominant error is in r, where results are shown for tracking errors to accuracies of approximately 10 microns.

Chapter 7

MICE Beam Dynamics

MICE employs an SFoFo lattice to provide transverse focusing in the cooling channel. A linear transport code has been developed that takes advantage of the G4MICE field modelling library to transport points in phase space and covariance matrices. The code is used here to find beams matched to the lattice and to study linear resonances. The results are compared with tracking through G4MICE using integration through the field maps.

7.1 Optics Package

The components of the G4MICE optics package are displayed in Figure 7.1. Points in phase space are represented by the class PhaseSpaceVector, which uses the CLHEP three and four vector classes to store the momentum, position, four-vector potential and electromagnetic fields. Probability distributions are represented by the class CovarianceMatrix, which stores the second moment matrix and mean of the distribution. Functions are provided to calculate the canonical and kinetic covariance matrices, Twiss parameters, angular momenta and emittances.

The TransferMap class stores the transport matrix and has the methods necessary to transport particles. Covariance matrices are transported using the relationship discussed in Chapter 3,

$$\mathbf{V}^{fin} = \mathbf{M_2}\mathbf{V}^{in}\mathbf{M_2^T}. \tag{7.1}$$

The transfer map is calculated by the TransferMapCalculator class by integration of the transport matrix through the solenoid field. The integration is performed using the GNU Scientific Library (GSL) implementation of fourth order Runge Kutta with adaptive step size [68].

The Material class calculates the change in the covariance matrix and the

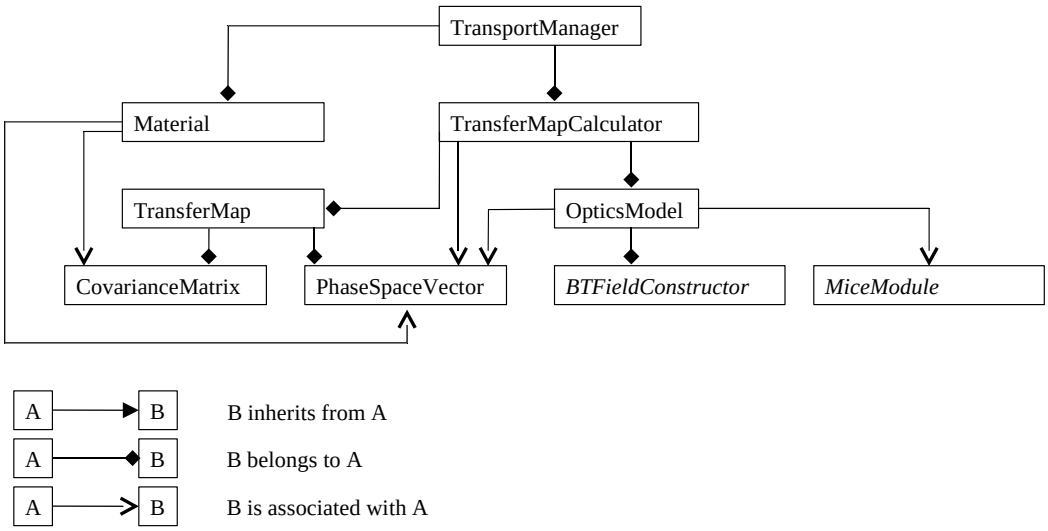

Figure 7.1: A UML diagram displaying the components of the Optics package in G4MICE. Classes external to Optics are shown in italics.

change in the phase space vector of a muon as it passes through some material. The Bethe Bloch curve is used to calculate energy loss and the Moliere formula is used to calculate the change in covariance for muons. Implicit in this operation is that the probability distribution described by the covariance matrix and the moliere scattering distribution are approximately Gaussian so that addition of errors in quadrature is appropriate.

The TransportManager class ties the material model and the transfer map model together. The TransportManager calculates transfer maps up to and between each material; on transport, covariance matrices are alternately transported by transfer maps and materials. It is assumed that materials fill the beamline and that scraping is negligible.

7.1.1 Integration of Transfer Map

In transverse phase space the focussing strength of the fields is calculated from the on-axis components of the solenoidal field and the RF field, and the quadrupole field gradient calculated numerically from the field map. This is used to calculate the change in the transfer map according to equations (3.99) and (3.100). The equations are integrated through z as 16 coupled differential equations in transverse phase space and 4 coupled differential equations in longitudinal phase space. Longitudinal and transverse phase spaces are assumed to be decoupled for faster integration.

The GSL provides a number of parameters to set the precision of the integration. Two parameters control the numerical precision in the MICE Optics module: the relative error constrains the error across a step relative to the value at the step end; and the absolute error constrains the absolute error across a step. In Figure 7.2 the determinant of the transfer matrix is displayed as a function of these precision parameters after integration of the transfer matrix was taken through the MICE step 6 fields from $z = -5500$ mm to $z = 5500$ mm; for transport to be symplectic, the transfer matrix should have a determinant of 1 and this is used as a measure of the accuracy of the integration. It can be seen that the determinant is convergent on 1.

7.2 Transverse Beam Dynamics

The transverse β function and magnetic field for various MICE steps has been studied previously [69] using the ICOOL tracking code, under the assumption of a cylindrically symmetric beam. Here the process is repeated by way of verification

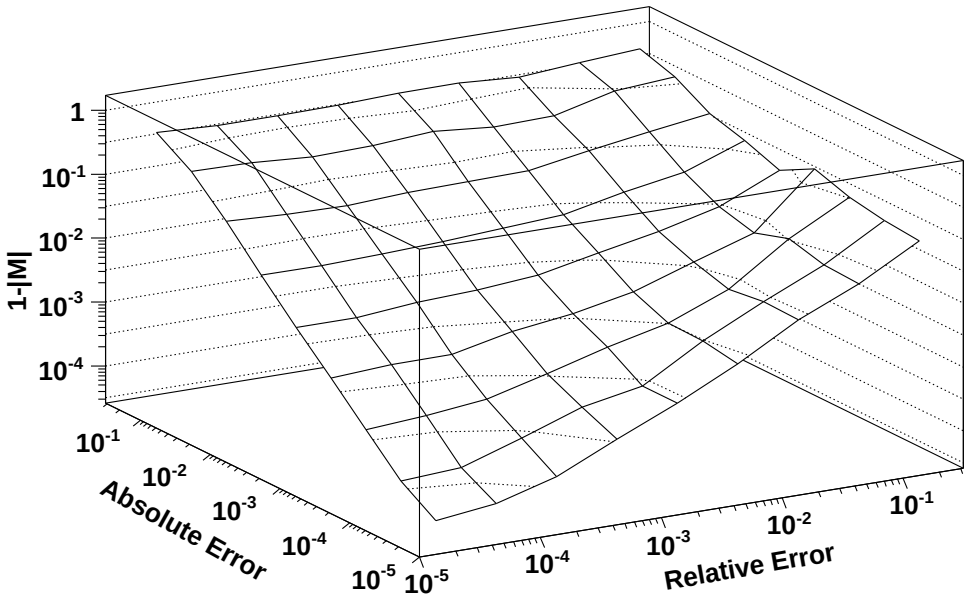

Figure 7.2: The deviation of the determinant of the transfer map $\mathbf{M}$ from unity as a function of relative and absolute error of the numerical integration.

of the G4MICE field modelling tools and transfer map code. In addition, by comparing a monochromatic beam with a beam with a modest energy spread second and higher order effects are studied.

7.2.1 Tracking in Magnetic Field

In Figure 7.3 the β function is shown for a number of different beam configurations and compared with the β function calculated by the linear optics code described above.

The linear transfer map was calculated, which was then used to transport a covariance matrix. The covariance matrix was chosen to have β constant through the solenoids and canonical angular momentum 0, in line with the MICE baseline; this is designed to be periodic across each of the 2.75 m cells of the lattice.

Subsequently, several beams were tracked through the MICE fields. Three beams were chosen; a monochromatic beam with small emittance such that the linear approximation holds; a monochromatic beam with a canonical 6 mm emittance such that only third and higher order terms contribute to the lattice transport; and a polychromatic beam with an energy spread slightly lower than that typical of a Neutrino Factory beam and a 6 mm emittance beam such that second and higher terms contribute to the transport.

The low emittance bunch shows good agreement with the linear calculation of the β function evolution through the lattice, as expected. However, significant deviation from the linear calculation is observed for the bunches with large emittance. This is to be expected, as $\sigma(E)/p, \sigma(p_x)/p$ are in the range $5 - 15\%$ so even third order effects are significant.

This also leads to a significant emittance growth, as observed in Figure 7.4, where the emittance change of the same three beams is shown. The monochromatic beam with low transverse emittance shows a very small emittance growth, as the linear approximation holds. The monochromatic beam with large transverse emittance shows more emittance growth, and the polychromatic beam with large transverse emittance shows considerably more emittance growth.

Second and higher order terms give non-linear beam optics, determined to n^{th} order by terms of order $O(u_{i_1} u_{i_2} \dots u_{i_n})$ in the transfer map. As shown previously, in transverse phase space second order terms are purely chromatic so that

$$H_3 = \sum (p_t v_i) \tag{7.2}$$

where $\vec{v}$ is the vector of transverse phase space coordinates. If the beam is monochromatic, however, emittance growth is constrained to third and higher order terms. Third order terms come in two flavours; those that arise due to

the z-dependent Hamiltonian and those that arise from cubic terms in the vector potential. Cubic terms in the vector potential are proportional to $\partial^2 B/\partial z^2$, which is peaked in the absorbers. Thus it is no surprise to see large changes in the emittance near to the absorbers for a beam dominated by these third order terms.

A full understanding of these non-linear terms and the emittance growth is beyond the scope of this thesis, but the emittance growth is of especial importance to MICE.

7.3 Longitudinal Beam Dynamics

The Feasibility Study II muon cooling channel operates RF at 50° in order to provide a large energy gain with the relatively large bucket required to transport the high emittance beam characteristic of a Neutrino Factory. In MICE it is envisaged that the experiment will sample muons relative to the RF phase set at 50° or 90° depending on the RF power available.

As MICE uses the ISIS ring as a proton source for its muons, the input beam is expected to have an approximately flat distribution in time on the scale of the MICE RF, with macro-pulses every few ms. Any energy-time distribution will be introduced during analysis.

In this section the solution for a periodic energy-time relation is found for the case where the RF is 50°. In addition, the effects of the energy gain and RF cavities on transverse phase space are examined in both the linear and non-linear regimes.

7.3.1 RF bucket

In Figures 7.5 and 7.6 the effects of various phenomena on the shape of the RF bucket are examined. Contours in the Hamiltonian are calculated using [34]

$$\frac{\omega}{cm(\beta\gamma)^3}p_t^2 + E_0\Big(\cos(\omega t) - \cos(\phi_0) + (\omega t - \phi_0)\sin\phi_0\Big). \tag{7.3}$$

Muons were passed through a repeating SFoFo lattice with a length of 2750 mm and their position in phase space relative to the reference particle was examined. The SFoFo was made up of 350 mm cylindrical liquid Hydrogen blocks followed by sets of four RF cavities. This is similar to the central cooling section of the MICE lattice, with the spectrometers removed and repeated many times. Muons were fired along the accelerator axis for energies above and below the reference energy.

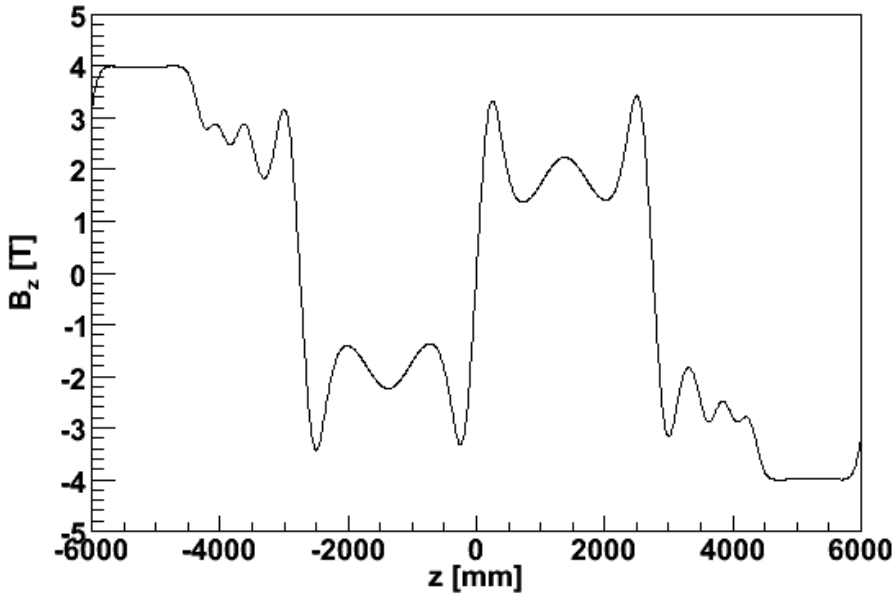

(*a*)

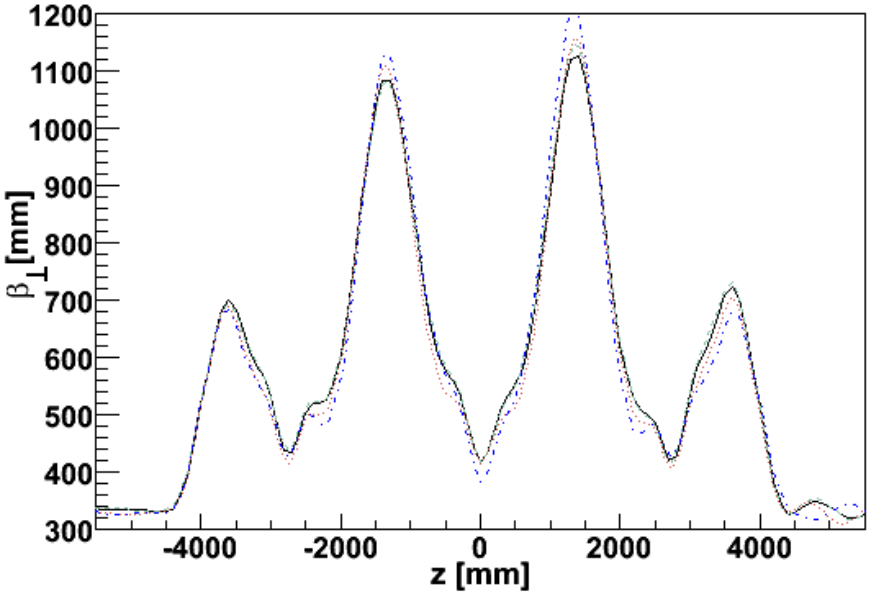

(*b*)

Figure 7.3: (a) The on-axis magnetic field through the MICE cooling channel for the final MICE configuration, step VI. (b) The β function tracked through G4MICE for various beams and as calculated by the G4MICE Optics package.

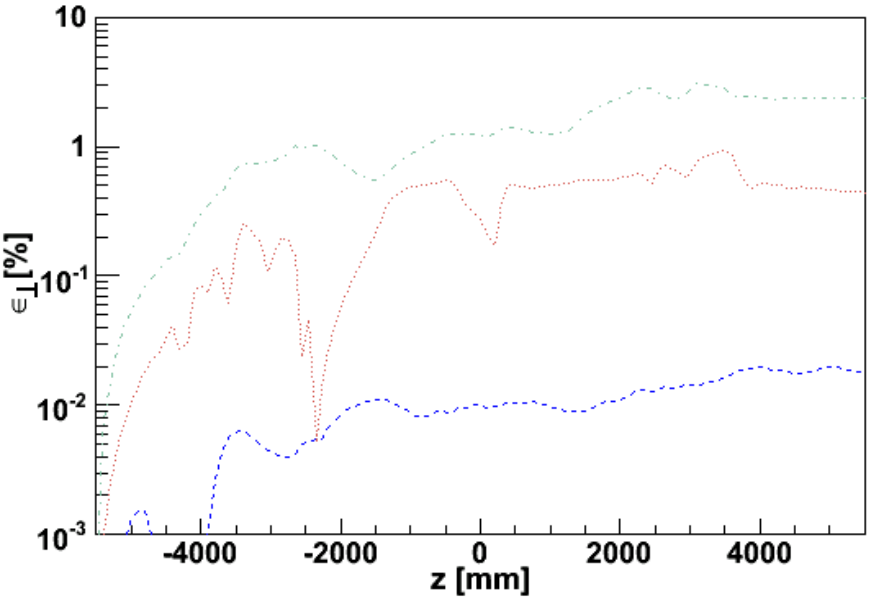

Figure 7.4: Fractional change in emittance over MICE step VI for various different beams.

Results for linear transport are shown in Figure 7.5 for the RF phase set to 50°. Points represent the position in longitudinal phase space of muons at the end of each 2750 mm lattice. Two models for energy loss in material are used: firstly the muons are assumed to be minimum ionising so that all muons receive the same energy loss; and secondly the muons see energy loss given by the Bethe Bloch curve without stochastic processes. In the former case, muons describe an elliptical path as expected. In the latter case, a slight emittance increase is observed due to the gradient of the Bethe Bloch formula; slow muons receive a slightly greater energy loss while fast muons receive a slightly smaller energy loss.

In Figure 7.6 muons are examined travelling around the RF bucket as tracked in G4MICE. Muons are simulated transported by RF field maps of the MICE cavities and by ideal pillboxes. Energy loss is governed by the Bethe Bloch curve but stochastic processes are ignored. In both cases, slight emittance growth is seen due to the gradient of the Bethe Bloch formula. Slightly larger emittance growth is seen when RF cavities are used but the distortion in the RF bucket due to the field map is quite small and much smaller than the distortion due to the Bethe Bloch curve.

A similar exercise is repeated in Figure 7.7, but with the RF set to 90°. In both cases here an RF field map is used; but in the latter case, energy straggling and multiple scattering are activated. In this case, these stochastic processes have a relatively small effect on the muon.

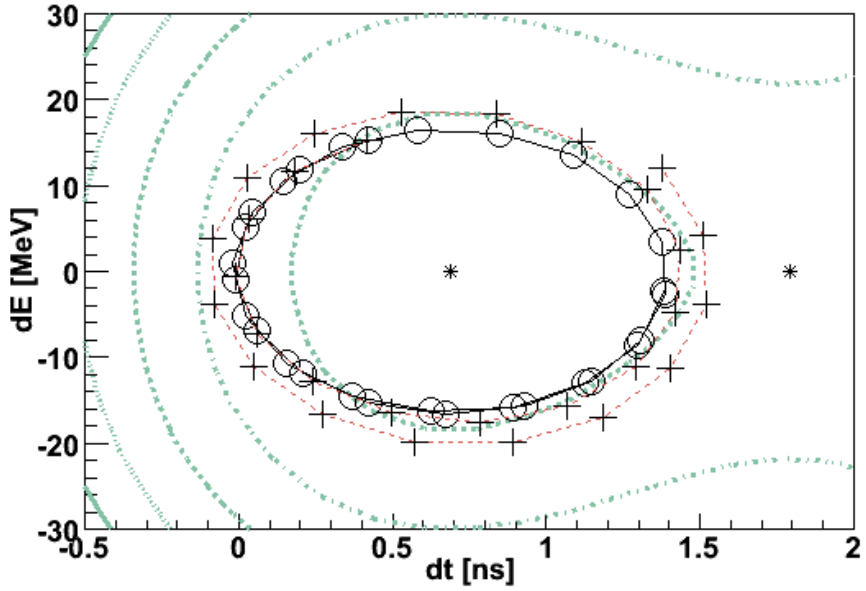

Figure 7.5: Muons moving around the RF bucket under linear beam optics: (full) energy loss in liquid Hydrogen assumed minimum ionising (dashed) energy loss in liquid Hydrogen under the Bethe Bloch model. Contours in the Hamiltonian are shown in turquoise.

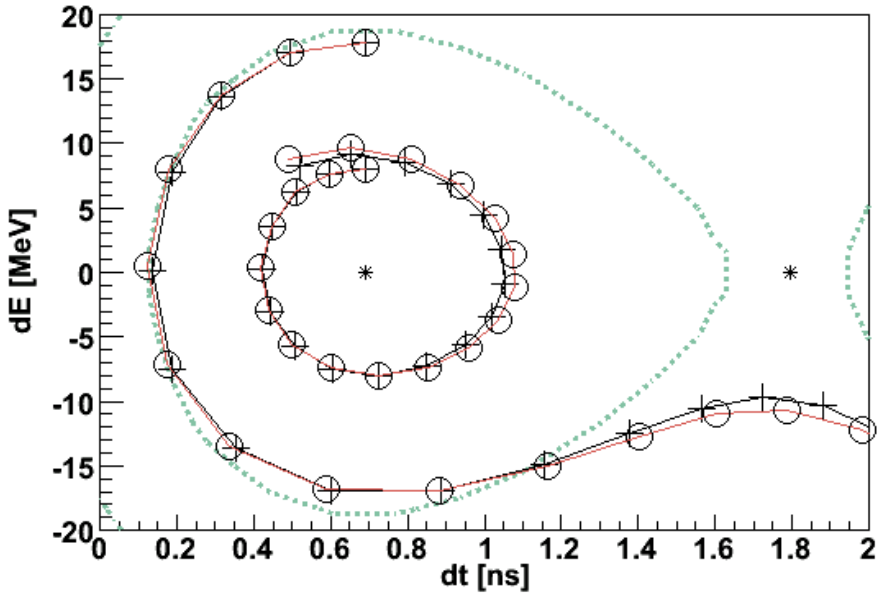

Figure 7.6: Muons tracked around the RF bucket with RF phased at 50° (o) using RF field maps (+) using the pillbox model. Contours in the Hamiltonian are shown in turquoise.

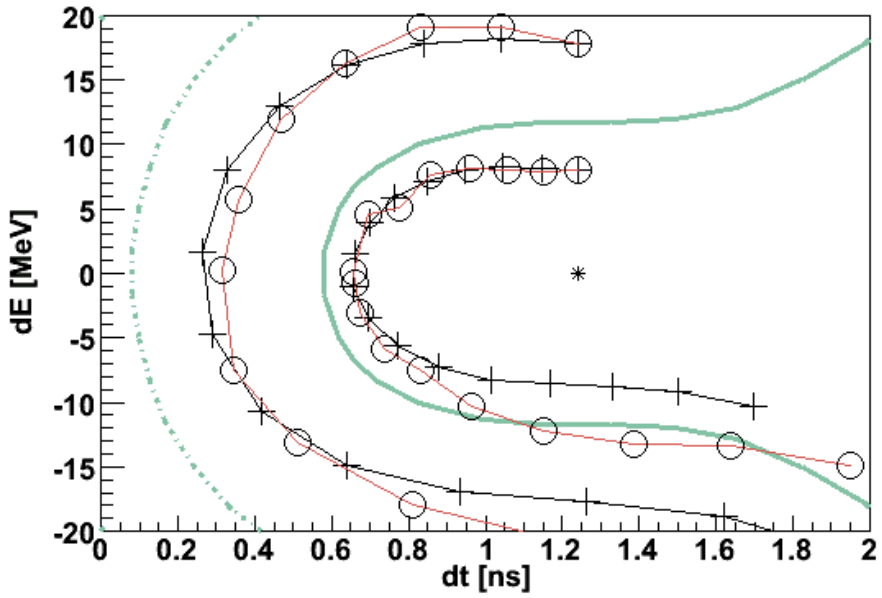

Figure 7.7: Muons in longitudinal phase space with RF phased at 90°: (+) with no stochastic processes; (o) with multiple scattering and energy straggling. Contours in the Hamiltonian are shown in turquoise.

7.3.2 Longitudinal matching at 50°

It is not possible to find a periodic solution for the longitudinal β function when the RF operates at 90° as there is no RF bucket. Instead, a beam will be used that is initially matched for the RF bucket operating at 50°.

The periodic solution to the β function in the SFoFo region of MICE is plotted in Figure 7.8. The periodic solution was found by calculating the transfer map through a 2750 mm SFoFo lattice and finding the value for which β and α are periodic using a minimising routine on the input covariance matrix. The transfer map was then calculated for the longitudinal drift space before the 2750 mm lattice and the inverse of this transfer map was used to transport the covariance matrix back through the absorbers. For this simulation, the absorber was taken to be split into two halves, each 175 mm long, so that the correct energy loss would be seen at the beginning and end of the SFoFo lattice. Energy straggling was taken to have a negligible effect on the bunch covariance and a constant RF field was used to calculate the longitudinal focussing. For a bunch with a moderate energy spread over a short lattice, these are reasonable approximations. The mean energy loss was calculated using the Bethe Bloch curve.

The periodic solution was then compared with Monte Carlo tracking in G4MICE and the results are also shown in Figure 7.8. The same model was used for MICE.

However, both pill boxes and RF field maps were used in the tracking and the results compared. For these cases a transverse emittance of 6 μm emittance was used. Each case shows slight deviations but in general good agreement is found with the linear approximation.

Finally, a bunch was considered with a larger transverse emittance of 6 mm emittance but the same longitudinal covariance matrix. In this case, a considerable coupling between the longitudinal and transverse phase spaces is observed, particularly when the beam is tightly focussed. In Feasibility Study II [12] it is noted that bunches develop a so-called amplitude-momentum correlation between transverse particle amplitude and longitudinal momentum. Muons with high transverse amplitude tend to have a larger transverse momentum, particularly in regions where the beam is tightly focussed. For a given energy muons with higher transverse momentum have a smaller longitudinal momentum, so that $dt/dz = E/p_z$ is larger and these muons move forwards through the bunch. This phenomenon has been discussed with particular reference to FFAGs [70]. It has been observed empirically that a second order term of third moments can be combined with the normal emittance calculation to give a quantity that is more closely conserved. However, such an analysis of the non-linear optics is beyond the scope of this thesis.

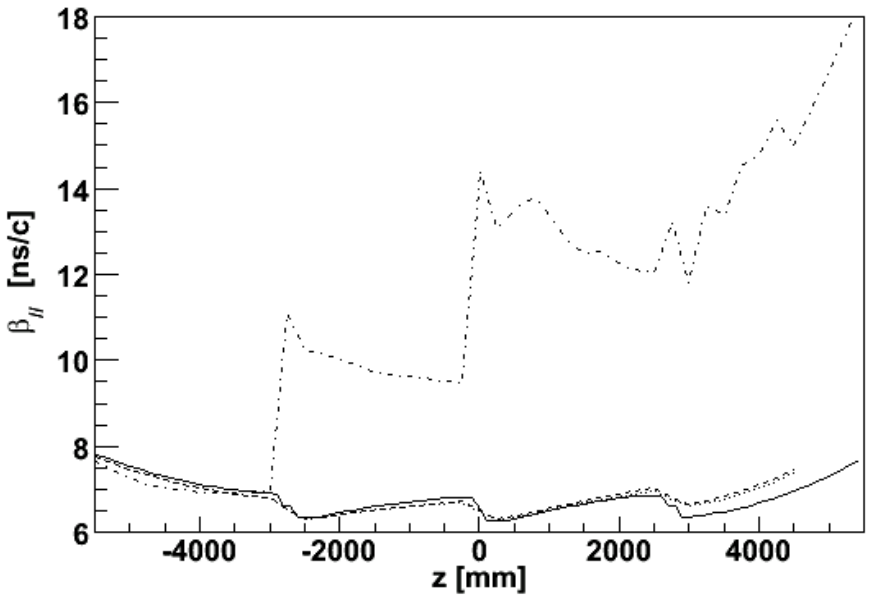

Figure 7.8: Longitudinal β function through MICE Stage 6: (full line) calculation from linear optics model; (dashed) Monte Carlo tracking with pill box cavities; (dotted) Monte Carlo tracking using RF cavities; (dot-dashed) Monte Carlo tracking using RF cavities with 6 mm transverse emittance.

7.3.3 Transverse Phase Space

In Figure 7.9 the transverse β function is plotted for muons under the linear model and using G4MICE. All processes were included and RF cavities were present. In this case the linear β function has a slightly stronger mismatch than previously observed. Under tracking through G4MICE non-linear effects can also be observed. It is interesting to note that there is no observable difference between tracking through the pill box field and the RF field map on transverse phase space.

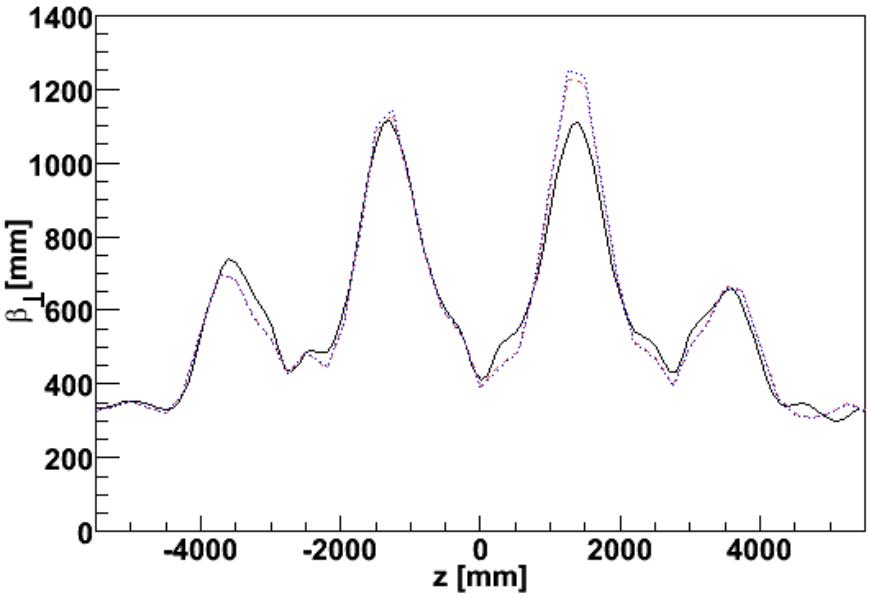

Figure 7.9: Transverse β function through MICE Stage 6: (full line) calculation from linear optics model including material and RF cavities; (dashed) Monte Carlo tracking with pill box cavities; (dotted) Monte Carlo tracking using RF cavities.

7.4 Resonances

In cooling channels, strong linear resonances are observed due to the large spread in longitudinal momentum of the beam. This causes reduced transmission and enhanced emittance growth in certain momentum ranges for a given magnetic lattice. A detailed study has been performed of the resonance structure of the Feasibility Study II lattice [42]. In this section, the linear resonance structure of the MICE lattice between the detector ends is compared with that of the Feasibility Study II lattice. Even in a short lattice like MICE, a beam with tails

in the resonance regions sees enough emittance growth to drown out transverse cooling.

7.4.1 Linear Transfer Map

In regions where the trace of the linear $2N$-dimensional transfer map is larger than $2N$ the transfer map is in resonance. In Figure 7.10 the trace of the transfer map for the periodic 2750 mm SFoFo lattice is compared against the trace of the transfer map in MICE Stage VI for transport from -4790 mm to +4790 mm and from -5201 mm to +5201 mm. In all cases the effect of material and RF is neglected. In MICE emittance change will be measured between the innermost tracker planes at +/-4790 mm; while +/-5201 mm is the centre of the tracker solenoid and so makes a natural point about which to define a periodic lattice in order to compare the linear theory with Monte Carlo.

The SFoFo lattice has a weak resonances at 175 MeV/c. Strong resonances are observed at 145 MeV/c and 240 MeV/c which bound the 200 MeV/c pass-band. The pass-band for the MICE lattice is narrower. If the lattice is taken between +/- 5201 mm weak resonances are observed at 160 MeV/c and 180 MeV/c with stronger resonances also at 145 MeV/c and 240 MeV/c. If the lattice is taken between +/- 4766 mm strong resonances are observed at 250 MeV/c and 140 MeV/c with weaker resonances at 220 MeV/c, 180 MeV/c and 160 MeV/c. These weak resonances may give some emittance growth in the cooling channel, although as the lattice is not periodic in MICE it is not clear how much effect they will have. However, in these regions it is not possible to define a β function.

7.4.2 Monte Carlo

The effect of the resonances on a repeating lattice can be seen in Figure 7.11. A repeating lattice was simulated in ICOOL made up from six 10402 mm MICE step VI lattices and twenty four 2750 mm MICE SFoFo lattices. A 6 mm beam, gaussian in transverse phase space and chosen to be periodic at 200 MeV/c was passed through the lattice, with a rectangular distribution in momentum. As above, no RF or materials were included in the simulation.

The transmission reflects the stop bands described above. In the case of the SFoFo lattice, the transmission drops quickly on either side of the 200 MeV/c pass-band; a weak resonance may be present at 190 MeV/c. In the case of the MICE 10402 mm lattice, the 200 MeV/c pass-band is again bounded by strong resonances either side of the 200 MeV/c pass band. However, after 6 lattice periods the transmission in the pass band is generally poor, perhaps indicating

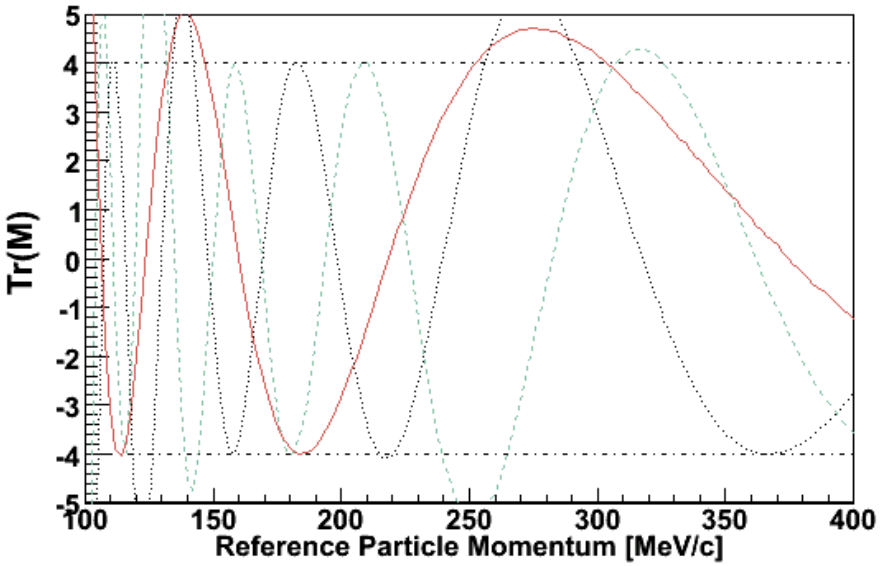

Figure 7.10: Trace of the linear transfer map: (full line) calculation for the 2750 mm SFoFo lattice; (dashed) calculation for MICE step VI across +/-5201; (dotted) caculation for MICE step VI across +/-4766. Horizontal lines are shown for $|Tr(\mathbf{M})| = 4$, beyond which the lattice is resonant.

the effect of the weaker resonances.

7.5 Summary

Various aspects of the beam dynamics in the MICE cooling channel have been investigated. The transverse beam dynamics have been simulated in G4MICE and show good agreement with a linear model. Muons have been simulated traversing the RF bucket and good agreement has been shown with analytical models. The longitudinal dynamics of the cooling channel have been investigated and linear resonance phenomena examined.

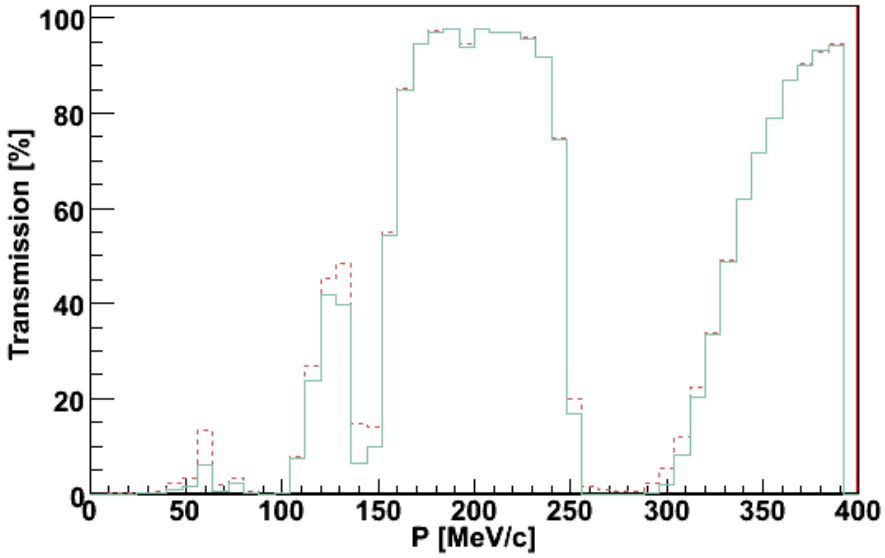

(*a*)

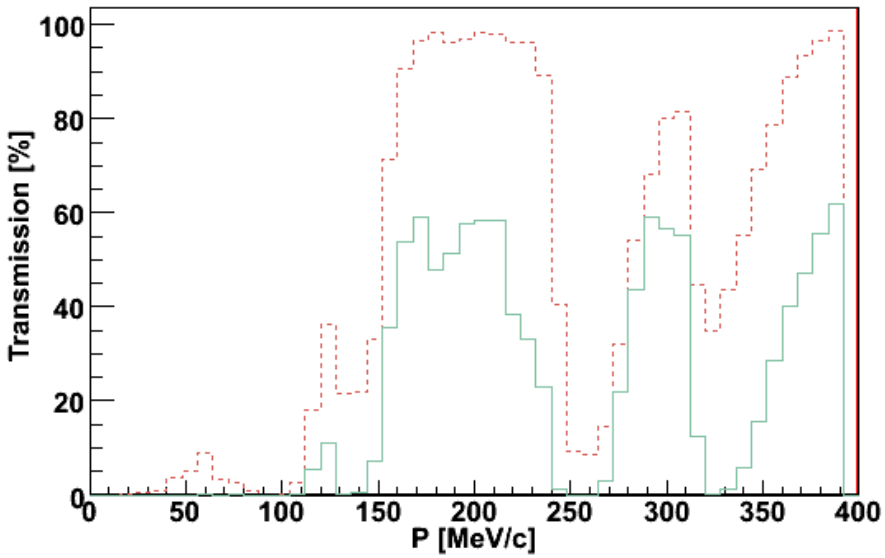

(*b*)

Figure 7.11: Transmission for a repeating lattice: (left) calculation for the 2750 mm SFoFo lattice; (right) calculation for the 5201 mm MICE step VI lattice.

Chapter 8

Cooling in G4MICE

MICE and MICE-like cooling channels have been simulated in a number of tracking codes, including ICOOL [12], COSY-INFINITY [31] and G4Beamline [58]. In this chapter simulations of the cooling channel using G4MICE are described. Materials effects on the beam dynamics are studied for MICE Stage 6; the effects of absorbers, absorber windows is reviewed and the overall cooling performance is examined.

8.1 G4MICE Cooling Channel Representation in G4MICE

The G4MICE interface allows an arbitrary degree of detail in the physical representation of MICE through a text interface known as MiceModules. The MiceModules provide an interface with G4MICE geometry objects such as cylinders and cuboids. In addition, specialised objects can be represented such as polycones, quadrupole apertures and the magnetic fields discussed elsewhere.

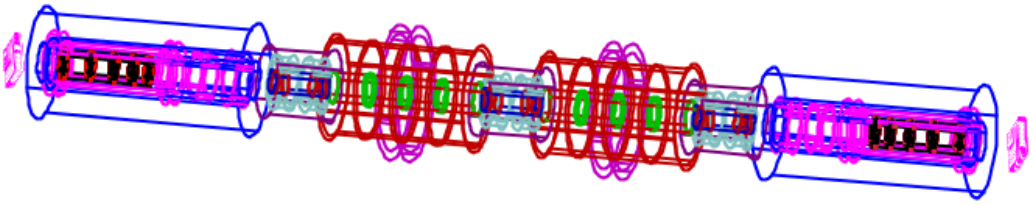

Figure 8.1: Visualisation of the MICE cooling channel from G4MICE.

8.2 GEANT4 Material Model

In the MICE baseline case, cooling is provided by liquid Hydrogen absorbers. Typically, liquid Hydrogen is difficult to simulate as simple Moliere scattering models do not simulate low Z materials well. The material model for the principle materials in MICE, Beryllium, Aluminium and liquid Hydrogen is studied below.

8.2.1 Multiple Scattering Model

Multiple Coulomb scattering limits the cooling performance that can be achieved in ionisation cooling channels. Scattering from Beryllium RF windows, Aluminium absorber windows and scattering in the absorbers themselves all contribute.

In Figure 8.2 the GEANT4 model of multiple scattering is compared with the Fano model [71] as implemented in ICOOL and a Gaussian fit to the Moliere formula (4.1) [10]. 10^5 muons with 200 MeV/c longitudinal momentum and no transverse momentum were simulated travelling through a 350 mm Hydrogen cylinder, a 180 μm Aluminium cylinder and a 380 μm Beryllium cylinder in ICOOL and G4MICE. These are the nominal thicknesses of the Hydrogen absorbers, Aluminium absorber windows and Beryllium RF cavity windows in MICE.

In all cases the Moliere model gives a smaller beam divergence than the Fano and GEANT4 models while the Fano model gives a slightly smaller spread than the GEANT4 model. The Fano model has been shown to agree well with experiment [72].

8.2.2 Energy Loss in Liquid Hydrogen

As outlined previously, cooling is provided by the energy loss in liquid Hydrogen, while energy straggling leads to heating in longitudinal phase space. In Figure 8.3 the liquid Hydrogen energy loss model in GEANT4 is compared with the energy loss model in ICOOL and approximations to the measured energy loss [10].

In Figure 8.3a, samples of 10^4 muons with various different longitudinal momenta were simulated travelling through a 350 mm thick cylinder of liquid Hydrogen and the energy loss over that distance was calculated. Only muons within 20 standard deviations of the mean energy loss were included in the calculation to remove statistical fluctuations from outliers in the long tail of the energy straggling. The Bethe Bloch curve is also shown; here the energy loss is calculated using the mean of the energy and momentum of the muons before and after the

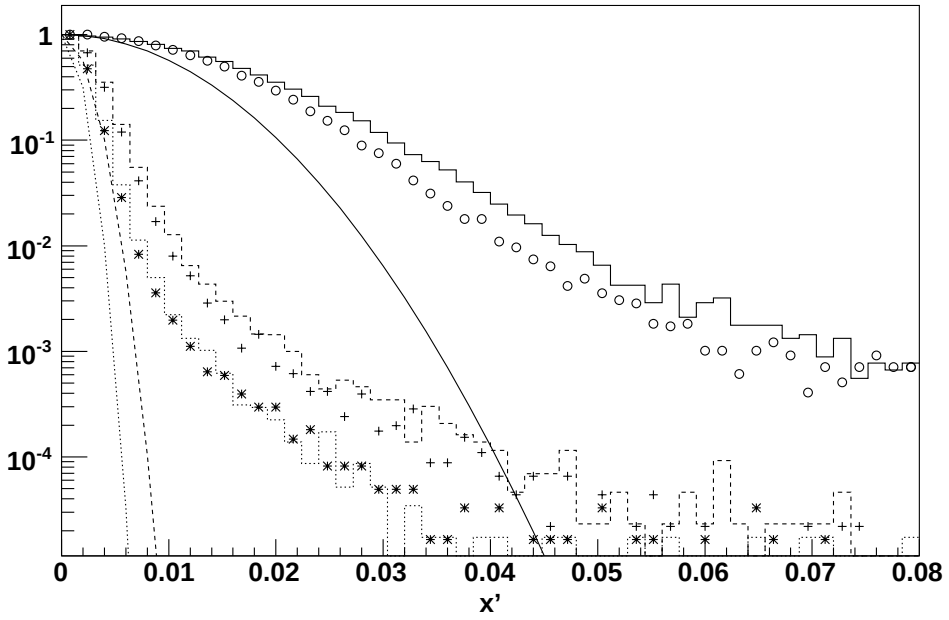

Figure 8.2: Multiple Scattering in G4MICE compared against other models: (o) Fano model in Hydrogen; (+) Fano model in Al; (*) Fano model in Be; (line histogram) GEANT4 model in liquid Hydrogen; (dashed histogram) GEANT4 model in Al; (dotted histogram) GEANT4 model in Be; (line) Gaussian approximation to the Moliere model in liquid Hydrogen; (dashed curve) Gaussian approximation to the Moliere model in Al; (dotted curve) Gaussian approximation to the Moliere model in Be.

Hydrogen block. In general this shows good general agreement with the energy loss calculated from the Bethe Bloch formula, although a deviation of order 1-2% is observed from the Bethe Bloch formula, which is of the order of the model accuracy [10].

The distribution in energy loss due to energy straggling for the 10^5 muon sample is shown In Figure 8.3b and compared with the ICOOL model and a fit to a Landau distribution with most probable value -10.369 MeV and width 0.4057 MeV. The ICOOL model shows a slightly higher energy loss while the Landau distribution gives a longer tail than both of the other distributions.

8.3 Absorber Model

In MICE, three absorbers provide transverse cooling. The simulation of the absorbers requires some care, as the absorber windows are curved and of varying thickness in order to contain the Hydrogen with the minimum amount of Aluminium. In this section, the window model is detailed and the effect on beam heating is compared to that of a cylindrical window. The overall transverse cooling performance of the absorber modules is then examined.

8.3.1 Window Shape

The absorber windows have been made as thin as possible while maintaining sufficient strength to contain the liquid Hydrogen. Each absorber has four windows; one at each end for containment, and an additional pair of safety windows. In order to maintain strength while minimising material in the beamline, the windows are curved and considerably thinner in the center than at the edge.

In G4MICE, the inner and outer radius of the window at regular z-intervals is read from a text file. GEANT4 calculates the surface of the window based on an interpolation between these points. In Figure 8.4 the implementation of the absorber window shape and thickness is shown. Muons were passed through the absorbers and the position of their tracks as they entered and left the windows was plotted. The two windows can be clearly seen; errors in the boundary position are below 20 % relative to the window thickness and most pronounced at high radii, which is populated by only a few muons. This error is especially pronounced at regions where dz/dr is small owing to the nature of the interpolation.

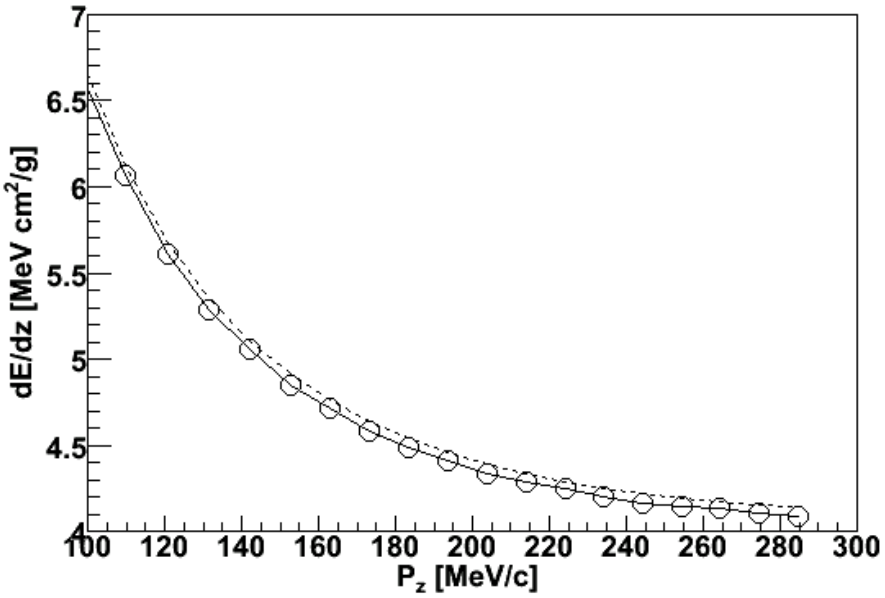

(*a*)

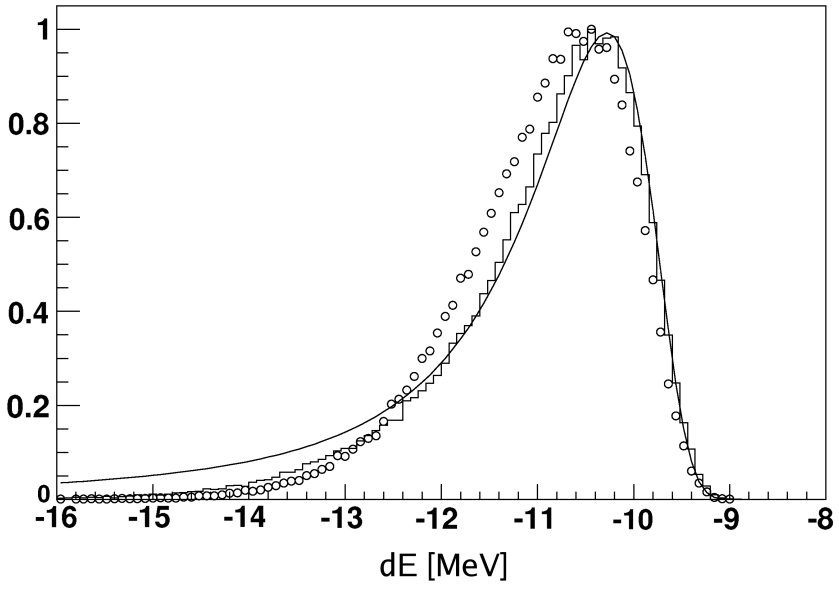

(*b*)

Figure 8.3: Energy loss in Liquid Hydrogen in G4MICE compared against other models: (a) mean dE/dz; (b) energy straggling.

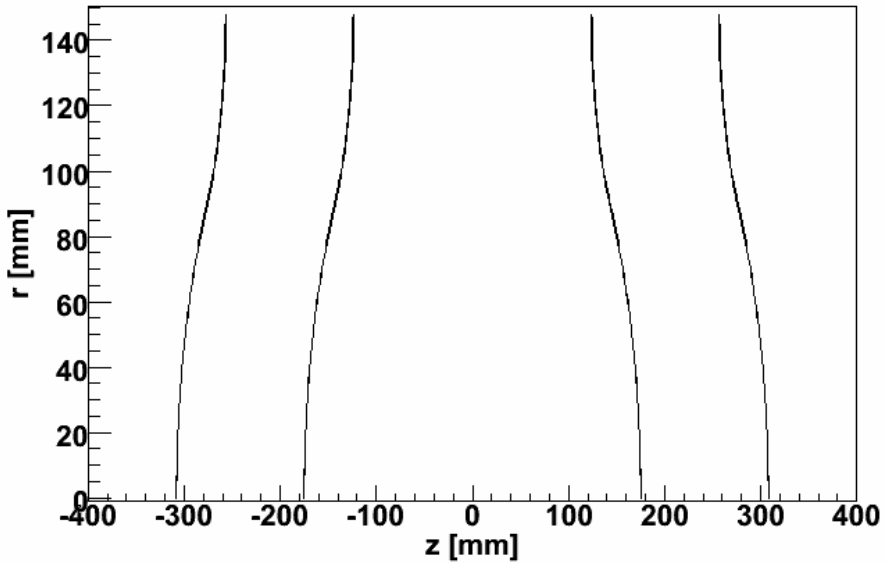

(*a*)

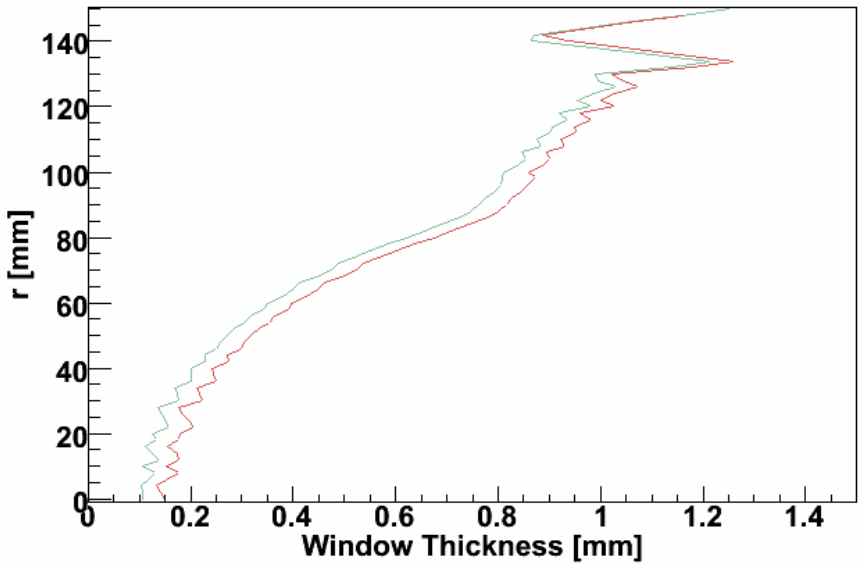

(*b*)

Figure 8.4: Absorber window shape and window thickness.

8.3.2 Emittance Growth in Windows

In Figure 8.5 the effect of the variation in window thickness is examined for an absorber window. 10^5 muons were passed through the window. The mean change in transverse single particle emittance was plotted for muons in different radius buckets to examine the effect of passing through the window at different radii. Recall from (3.65) that the bunch emittance is given by $\epsilon_\perp =< \epsilon_i > /4$ where ϵ_i is the single particle emittance, so that the change in bunch emittance is given by

$$\delta\epsilon_\perp = \frac{1}{n}\sum_{r_j} n_{r_j} < \epsilon_i(r_j) > \tag{8.1}$$

where n_{r_j} is the number of muons in each radius bucket and $< \epsilon_i(r_j) >$ is the mean change in single particle emittance in a radius bucket.

On examining Figure 8.5 it is apparent that muons at larger radius see a greater emittance growth, even for a cylindrical absorber, but the higher radius regions are not populated significantly. Surprisingly, the variation in window thickness does not have a significant effect on the beam heating. For a less focussed beam or one with a higher emittance the higher radius regions may be more highly populated. Errors are calculated assuming a relationship of the form $< \delta\epsilon_i > \approx \sigma(\delta\epsilon_i)/\sqrt{n}$ where n is the number of muons in the radius bin. The radius was taken at the downstream edge of the window while the change in emittance was taken between two planes 350 mm upstream and downstream of the window. A beam with emittance of 6 mm, β of 420 mm and α of 0 was injected at the upstream plane.

8.3.3 Cooling through Absorber Module

The transverse cooling performance of a single absorber is examined in Figure 8.6. Cylindrical and curved absorbers are examined.

For the cylindrical absorber, vacuum and absorber windows are taken to have thicknesses of 180 and 190 μm respectively and the liquid Hydrogen is taken to have a thickness of 350 mm. Muons that fall outside the absorber radius of 150 mm are discarded. For the curved absorber, the window profiles outlined above are used, with all space between the absorber windows filled with liquid Hydrogen. The thickness of the liquid Hydrogen in this case is 350 mm on-axis.

Monochromatic bunches of 10^4 muons were passed through the absorber. Beams with emittance of 6 mm and β of 420 mm at various different momenta and beams with momentum of 200 MeV/c and β of 420 mm at various different emittances were tracked through a cylindrical or curved absorber and the dependence of emittance change on momentum was examined.

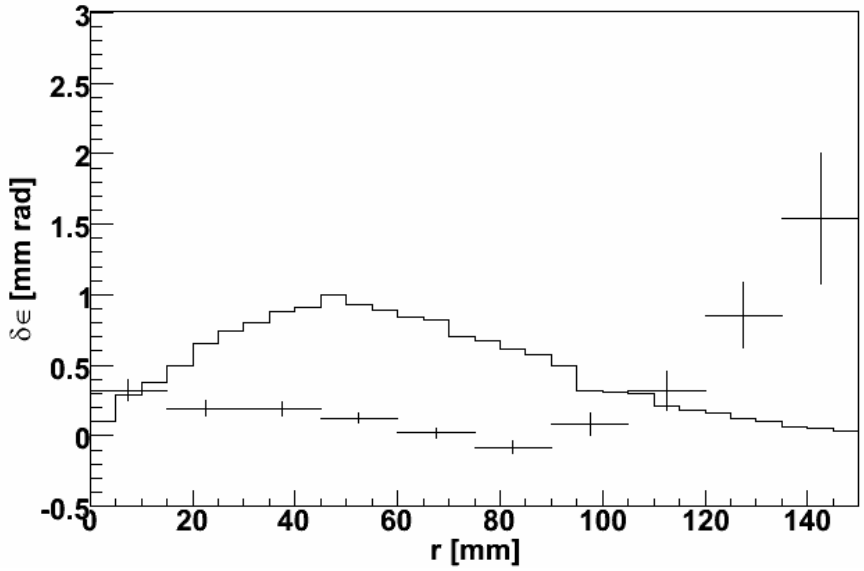

(*a*)

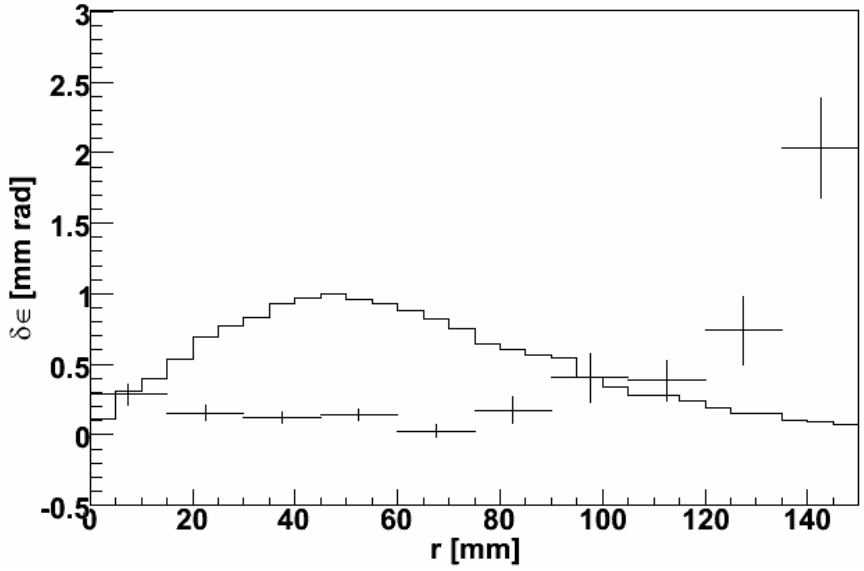

(*b*)

Figure 8.5: (Histogram) number of muons and (points) change in single particle emittance on passage through an aluminium window as a function of radius for: (a) a cylindrical window; (b) a curved window.

The beams see a weaker cooling effect at low emittances and even heating. Recall that the cooling term in the expression for emittance change is proportional to the beam emittance. The equilibrium emittance is roughly the same in both cases. The use of realistic absorber windows degrades the cooling performance slightly for higher emittances. This is to be expected as high emittance beams have more muons at large radii where muons pass through less liquid Hydrogen and more Aluminium.

The transverse cooling is stronger at low momenta as the mean fractional energy loss is greater. The use of realistic absorber windows degrades the cooling performance for lower momenta, but at higher momenta the effects of multiple scattering are considerably less pronounced. Recall from equation 4.9 that the heating term is proportional to $(E\beta_{rel}^3)^{-1}$, while the cooling term is proportional to $< dE/dz > /E$.

8.4 Simulated Emittance Change

The cooling performance of the full MICE cooling channel is now examined. There are a number of ways to assess the cooling performance of the MICE channel. The baseline measurement is to examine the change in beam emittance. MICE aims to measure the change in emittance of various beams to a precision of 0.1 %. This will be compared with simulation results and any discrepancy examined. In simulating cooling, there are two sources of errors: systematic errors due to errors in the physics model; and statistical errors due to the stochastic physical processes in materials. Errors in the physics model have been addressed in this and previous chapters.

In studying statistical errors, close attention must be paid to the quantity to be studied. There is a significant difference between attempting to understand the error in the change in emittance of a *distribution* and the error in the change in emittance of a *sample* of a distribution. If the change in emittance of a probability distribution is to be measured, then the error in sampling the distribution must be understood. For example, the covariance matrix of a randomly selected sample of particles will be different to the covariance matrix of the distribution from which they were sampled. If the change in emittance of a sample of a distribution is to be measured, then the dominant statistical error arises due to the stochastic nature of processes as the beam passes through material. This introduces an error both in the beam emittance and the beam dynamics through the rest of the cooling channel.

The statistical errors are particularly important as these contribute signifi-

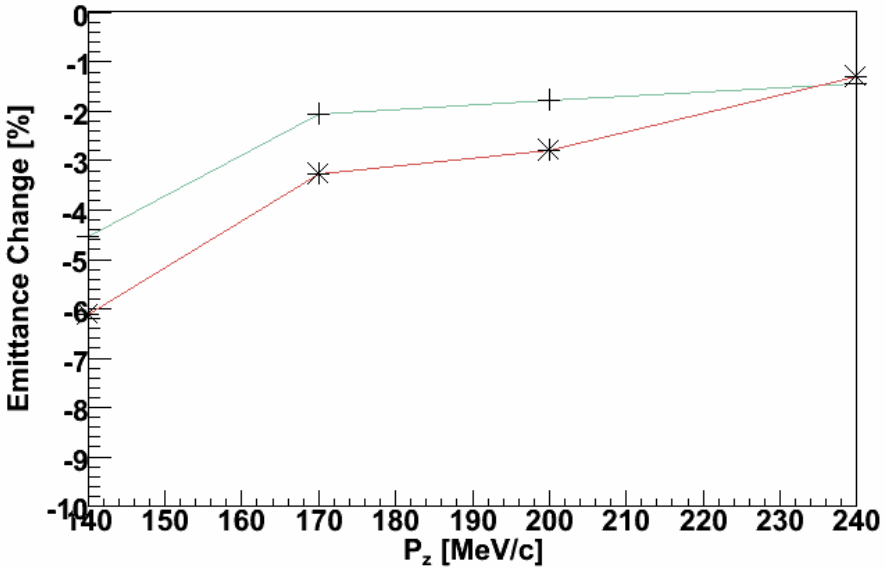

(a)

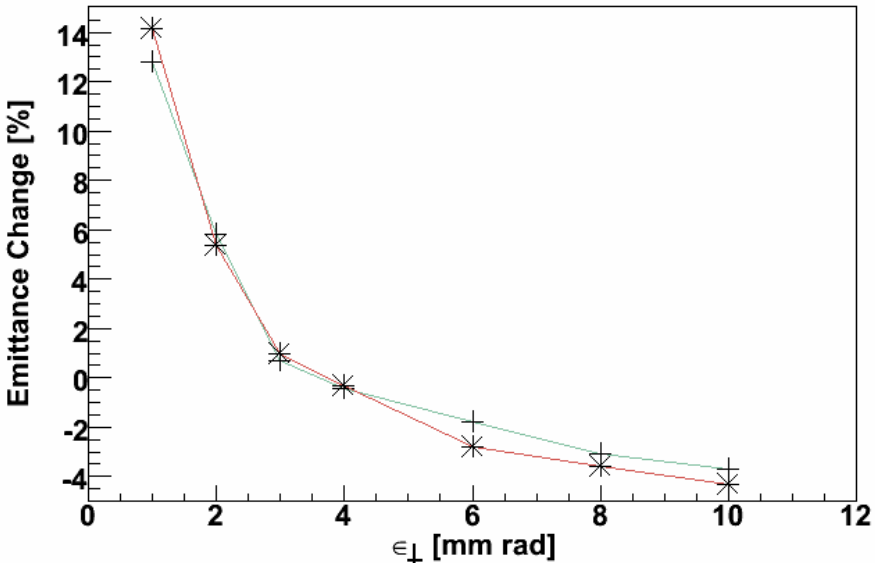

(a)

Figure 8.6: Change in emittance on passage through absorber as a function of (a) input momentum and (b) input emittance for (*) cylindrical absorbers and (+) curved absorbers

cantly to statistical errors on the emittance measurement. This is discussed in more detail in the next chapter.

8.4.1 Cooling performance

In Figure 8.7 the emittance of a sample of 10^5 muons passed through the MICE Stage VI cooling channel from -4690 mm to 4690 mm is shown as a function of position in the cooling channel. Polycone absorber windows were used as described above. In addition polycone RF windows with a constant thickness of 0.38 mm were used. The RF was phased for operation on crest, with an RF field map used, while the beam was matched for operation at 50°. It is not possible to match for operation at 90° as there is no longitudinal focusing at this phase. A beam with an initially gaussian distribution in the phase space coordinates was used, with transverse emittance of 6 mm and longitudinal emittance of 0.3 ns. The emittance of 0.3 ns corresponds to a rather low emittance for a typical Neutrino Factory beam. However, the RF bucket at 50° has an acceptance of about 0.5 ns, and muons above this amplitude are typically lost from a realistic Neutrino Factory cooling channel. In this section a smaller longitudinal amplitude cut of 0.15 ns will be used. This is necessary to prevent non-linearities from smothering the cooling.

Beam emittance is shown as a function of z along the cooling channel in Figure 8.7, calculated for muons with longitudinal amplitudes less than 0.15 ns. Much worse performance is observed for any higher amplitude muons due to beam heating in the cooling channel. Over 70% of muons are excluded due to this amplitude cut; in addition, a few % are scraped or decay and these are also excluded from the plot. In total 20218 muons survive both cuts. Even then, significant transverse emittance growth is observed near to the absorbers. The transverse emittance shows a marked decrease, but longitudinal emittance increases significantly; overall, 6D emittance increases due to the significant non-linearities introduced by the RF operating on-crest.

8.4.2 Statistical fluctuations

As discussed above, the statistical fluctuation on the transverse cooling performance can be considered for a sample of events from a distribution or for a distribution of events. These fluctuations on the MICE cooling measurement have been discussed previously[73] [74]. The two different cases are examined in 8.8: ten different samples of muons and ten sets of the same sample of muons were simulated passing through the cooling channel. The muons were sampled

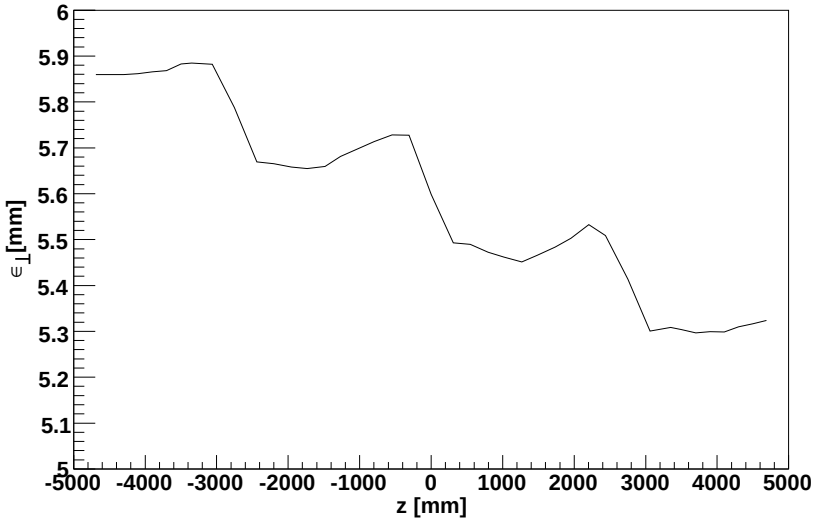

(a)

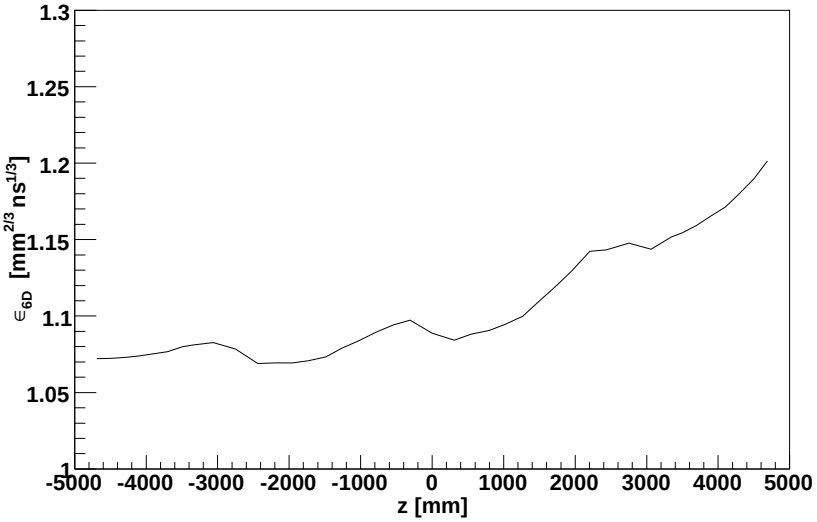

(b)

Figure 8.7: Change in emittance on passage through MICE stage VI: (a) transverse emittance; (b) 6D emittance

from the distribution outlined above and the same cut on longitudinal amplitude was applied. As expected, a marked difference was observed in the spread in transverse emittance change of the two datasets. In both cases the mean cooling performance is roughly the same, but the standard deviation on the transverse cooling performance is three times higher where different samples of muons are used. Somewhat suprisingly, the standard deviation on longitudinal cooling is roughly the same in both cases. This indicates that stochastic processes in the absorbers introduce sufficient statistical fluctuations to significantly change the longitudinal heating that is observed. The results are listed in Table 8.1.

The case where different samples of muons are passed through the cooling channel with no stochastic effects is also examined in Figure 8.8. Fluctuations in the longitudinal heating are of similar magnitude to the other two cases, indicating that selection of the beam sample introduces a similar fluctuation in heating as the stochastic processes do. Fluctuations in transverse cooling are much smaller. This is to be expected as the largest uncertainty in transverse emittance change comes from the heating term, where multiple scattering gives a mean increase in transverse momentum.

	$\sigma(\delta\epsilon_\perp)$	$<\delta\epsilon_\perp>$	$\sigma(\delta\epsilon_{//})$	$<\delta\epsilon_{//}>$	$\sigma(\delta\epsilon_{6D})$	$<\delta\epsilon_{6D}>$
Different Samples	0.148	-0.544	0.0026	0.0251	0.033	0.123
Same Sample	0.049	-0.552	0.0035	0.0251	0.020	0.056
No stochastics	0.018	-1.009	0.0030	0.0258	0.021	0.129

Table 8.1: Emittance change and RMS fluctuation when ten different input beams are used, when the same input beam is used over ten different runs and when ten different input beams are used but stochastic processes are switched off.

8.5 Other Possible Measurements

The utility of considerations such as change in beam amplitude has been shown in the previous discussion. The study of the number of muons within some acceptance is more representrative of the increase in the number of muons transmitted into the Neutrino Factory accelerator complex. Also, the aperture of the cooling channel is an important quantity as it determines the amount of focussing that can be achieved in the absorbers before scraping becomes dominant.

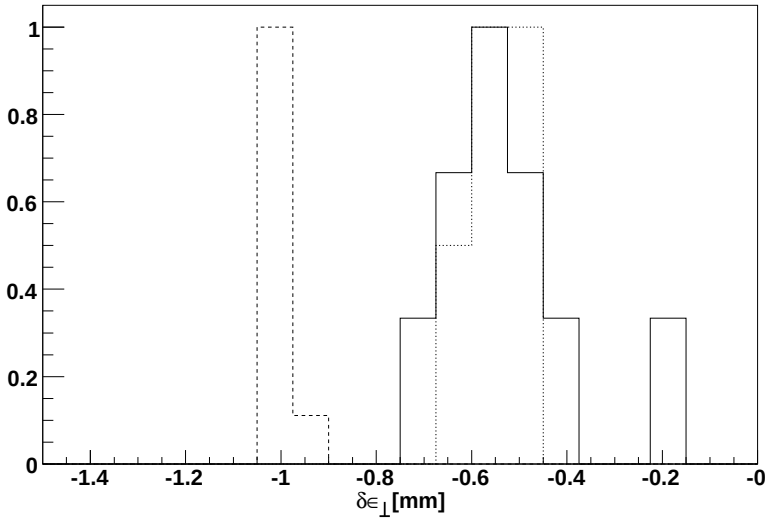

(*a*)

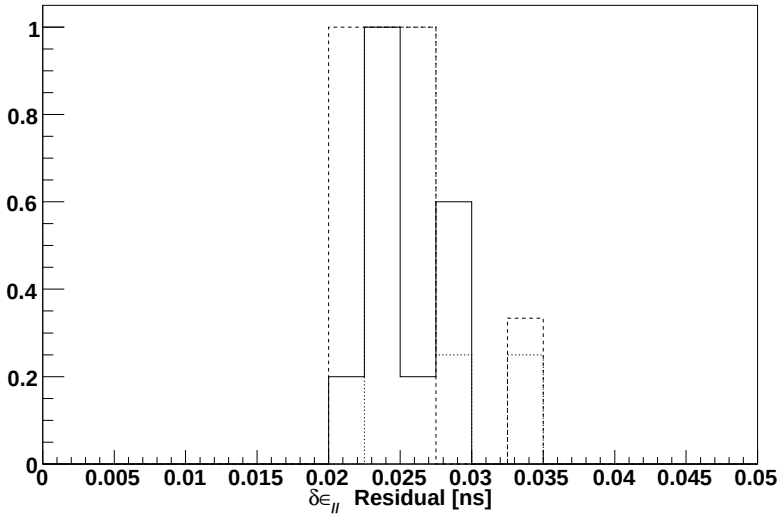

(*b*)

Figure 8.8: Emittance change for ten 2000 event samples over MICE stage VI: (a) transverse emittance; (b) longitudinal emittance

8.5.1 Scraping Aperture

The design of a long cooling channel calls for tight focussing on the absorbers in order to enhance the cooling effect. This in turn leads to more scraping. Typically, as the emittance of the bunch is reduced the amount of focussing on the absorbers is enhanced by changes to the magnetic lattice.

In Figure 8.9 the maximum radius of the cooling channel is shown using a simplified representation of the cooling channel simulated in ICOOL. In this sample, a matched bunch of 10^4 muons were transported through the cooling channel with a small energy spread. The maximum radius of the bunch is seen to be large in the RF cavities where transverse focussing is weaker.

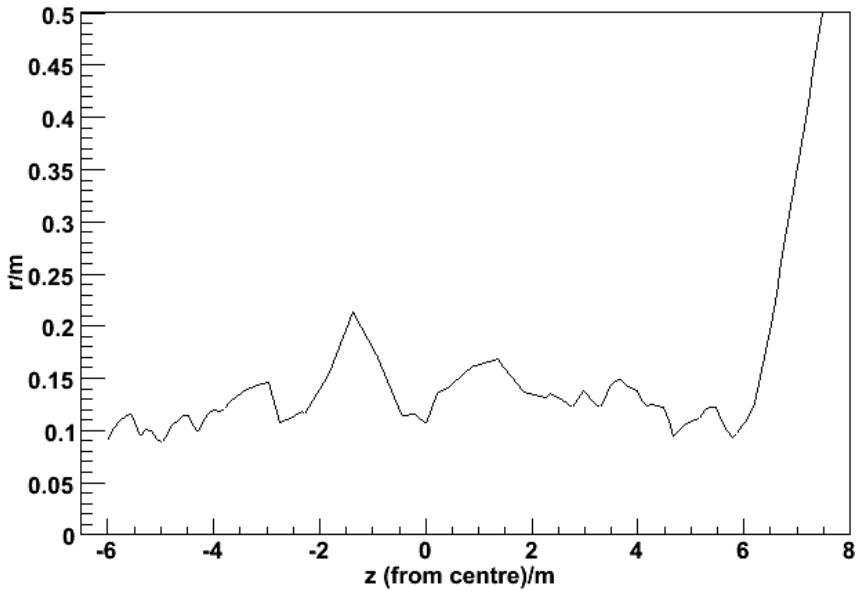

Figure 8.9: Maximum radius of muons as a function of z.

8.5.2 Transmission into an Acceptance

Emittance change is a useful indication of cooling performance, but for a Neutrino Factory the change in the distribution of particle amplitudes is also important. In 8.10 the number of muons contained within a particular acceptance is displayed along the MICE beamline for the 100000 event simulation described above. Three amplitude cuts are shown, typical of the acceptance of different Neutrino Factory accelerators. In all cases, the number of muons inside the cut increases as the bunch passes along the cooling channel. Only muons that are initially within the 0.15 ns longitudinal acceptance cut are shown.

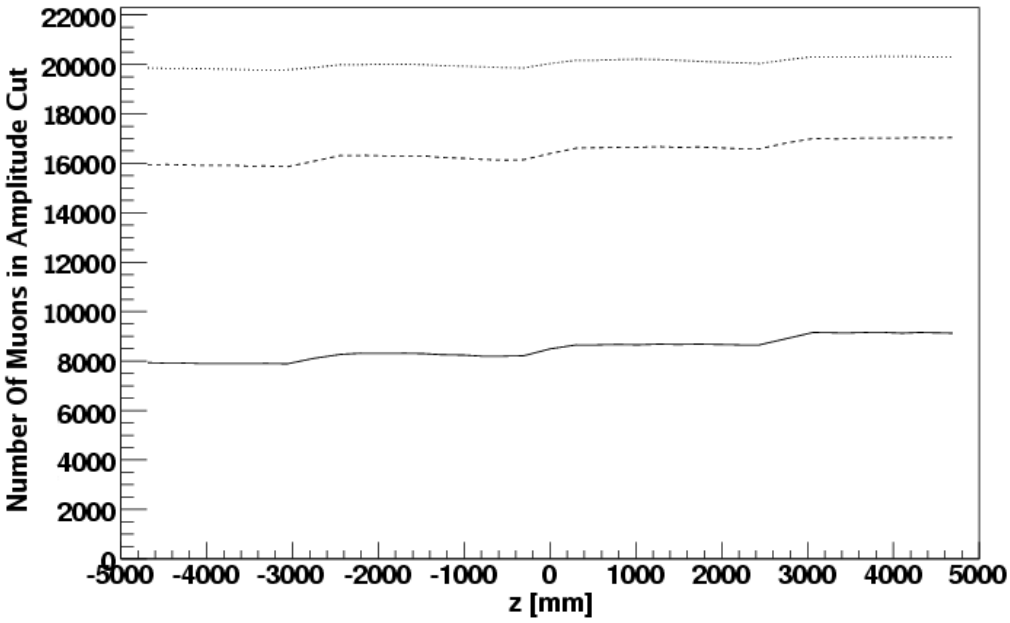

Figure 8.10: Number of muons within transverse acceptance $A_\perp^2$: (full) $A_\perp^2 < 15$ mm (dashed) $A_\perp^2 < 30$ mm (dotted) $A_\perp^2 < 45$ mm

8.5.3 Fractional change in acceptance

The fractional change in the number of muons between each end of the cooling channel is shown as a function of initial amplitude in Figure 8.11. This function has the interesting feature that it is expected to be an invariant quantity of the cooling channel for a matched beam, for amplitudes well below the scraping aperture and for linear beam optics. Also it is measurable by MICE. In Figure 8.11 10 different samples of 10^4 muons were simulated passing through the MICE cooling channel. Muons with longitudinal amplitude initially greater than 0.15 ns were excluded from results. Error bars show the standard deviation on the result.

There are three regions of interest. Below 6 mm, the number of muons increases, as the beam is cooled. Between amplitudes of 6 and 60 mm the number of muons in the beam decreases. In this region, muons are either migrating towards the centre of the beam due to cooling, migrating outwards due to heating from non-linearities or being scraped out of the beam. Above 60 mm, the number of muons increases as muons that are heated move to higher amplitudes. For higher emittance beams, high amplitude muons are expected to be lost due to scraping.

8.6 Summary

The model for material processes in GEANT4 has been evaluated including multiple scattering and energy loss. The physical representation of the cooling channel has been detailed, and the effect of realistic curved absorber windows on the cooling performance compared with cylindrical windows. The cooling performance of MICE step VI has been evaluated including statistical errors on the cooling performance, evaluated in terms of bunch emittance and particle amplitude.

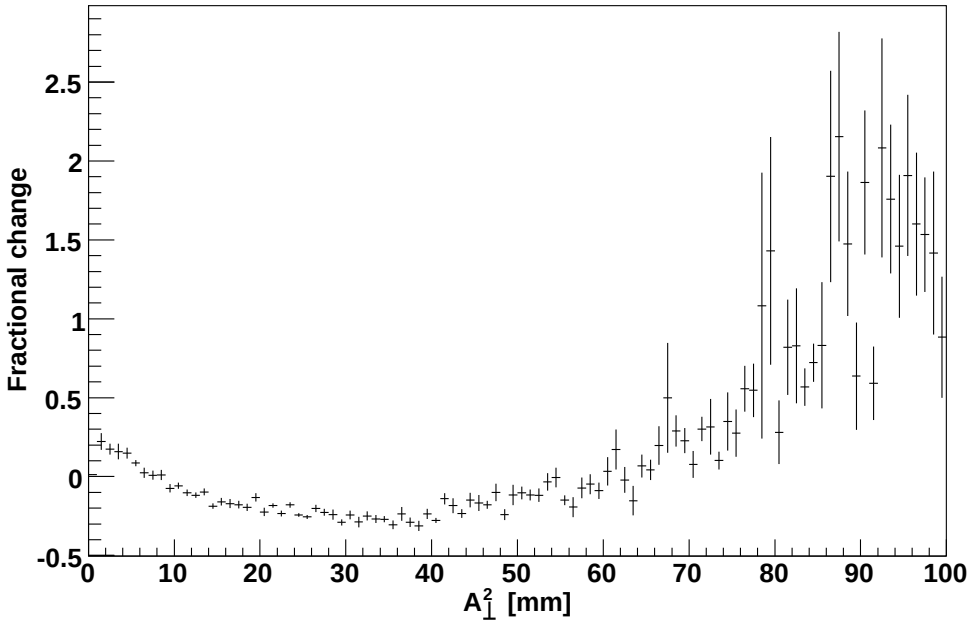

Figure 8.11: Fractional change in number of muons in transverse amplitude bin as a function of initial amplitude.

Chapter 9

Simulated Measurement of Cooling

In this chapter the measurement of the MICE cooling performance is simulated and the emittance resolution of the MICE detectors is examined.

In MICE transverse phase space coordinates and energy are reconstructed at the inside edge of the upstream and downstream tracker, while the time at which muons pass through the TOFs relative to the RF phase is reconstructed at the TOFs. In this chapter, the extrapolation of the timing measurement from the TOF to the tracker will be examined assuming some finite detector resolution in the phase space variables, and the effect that this finite resolution has on emittance measurement will be studied. It is assumed that an appropriate time distribution from the beam has been sampled.

9.1 Detector Models

The MICE trackers are used to reconstruct the momentum and position of particles as they enter and leave MICE. Reconstructed particles are used to calculate emittance. A full simulation and reconstruction of the tracker has been performed. However, the software that was used to perform the reconstruction was never published and is no longer available. A new reconstruction code is in preparation. In this chapter, fits to the tracker reconstruction are used instead. The fits are described below, and the parameters used are listed in Table 9.1

9.1.1 Tracker Model

The MICE tracker is described in detail in [54] and [23]. Measurement of the position of muons in phase space is achieved by fitting a helix to the measured

Parameter	Units	Value
$x_{1/2}$	[mm]	1.0
$y_{1/2}$	[mm]	0.7
s_{px}	[MeV/c]	2.3
s_{py}	[MeV/c]	1.8
p_0	[MeV/c]	150
c_{pt}		-0.003
c_{pz}		-0.004
s_{pz}	[MeV/c]	1.5
k_{pz}	$[(\mathrm{MeV/c})^{-1}]$	0.0006
$\sigma(t)$	[ps]	70

Table 9.1: Values of parameters used for the tracker model.

position of muons at each tracker plane and reconstructing the phase space coordinates at a well defined z-position known as the Tracker Reference Plane (TRP), here taken to be on the inside edge of each tracker, at 4.69 metres upstream and downstream of the cooling channel centre. The position measurement resolution is limited by the finite thickness of the scintillating fibres. The momentum resolution is limited primarily by the spatial resolution of the tracker. Multiple scattering at each tracker plane also has an effect.

In Figure 9.1 the simulated distribution of residuals in the position measurement is shown for x and y, modelled as a rectangular distribution with half-width of $x_{1/2}$ and $y_{1/2}$ respectively. It is assumed that the resolution is independent of other phase space variables. The square shape arises due to the ganging of the fibres. When the tracker was simulated, the tracker was aligned such that the fibres of one of the views lay parallel the x axis, with the other two views aligned at 30° each from the x axis, leading to an asymmetry between the two axes which is reflected in the reconstructed spatial resolution of the tracker.

The transverse momentum of a track is proportional to the radius of the fitted helix. In the limit that multiple scattering in the tracker planes is small, the transverse momentum resolution is a function only of the space point resolution of the tracker. To a good approximation,

$$\sigma_{p_x} = k_1 \sigma_x \tag{9.1}$$

where k_1 is some constant and σ_x is the space point resolution. In practice, there is a weak linear dependence of transverse momentum resolution on transverse momentum and on longitudinal momentum due to multiple scattering. The sim-

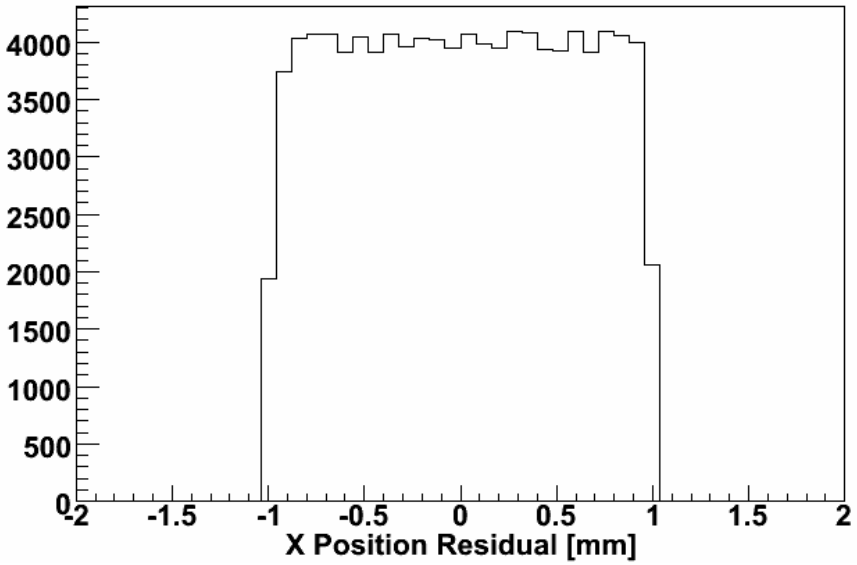

(*a*)

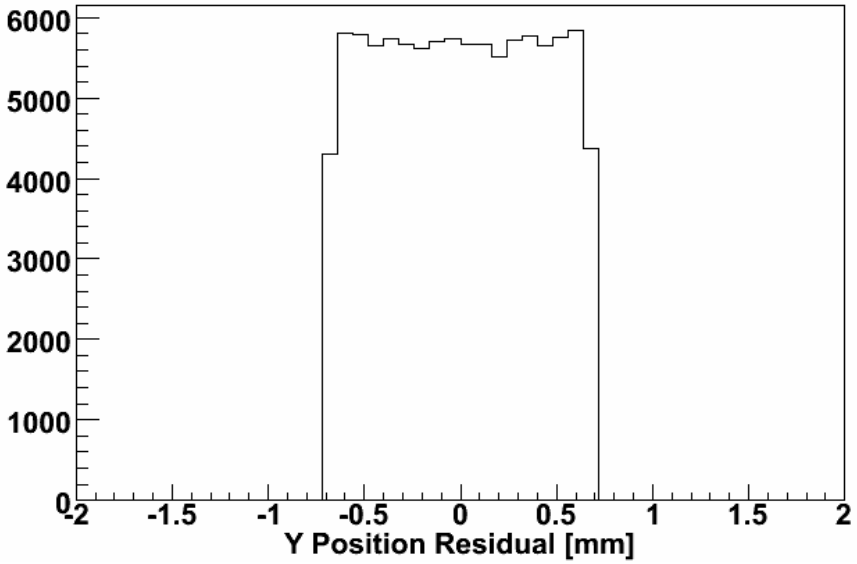

(*b*)

Figure 9.1: Simulated position residuals of the tracker.

ulated distribution of errors in p_x and p_y is shown in Figures 9.2 and 9.3, given by Gaussian distributions with RMS

$$\sigma(p_x) = s_{px} + (p_z - p_0)c_{pz} + p_t c_{pt}, \quad (9.2)$$

$$\sigma(p_y) = s_{py} + (p_z - p_0)c_{pz} + p_t c_{pt} \quad (9.3)$$

and mean 0.

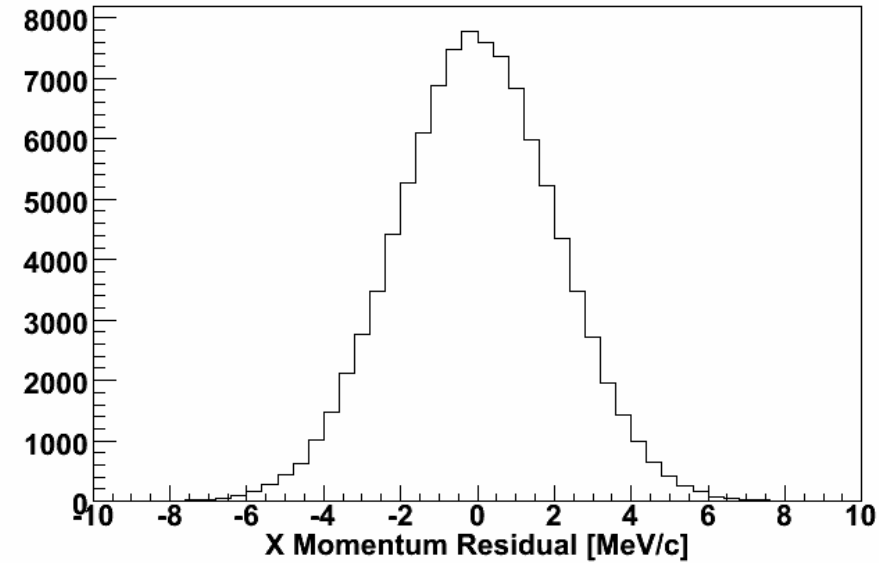

(*a*)

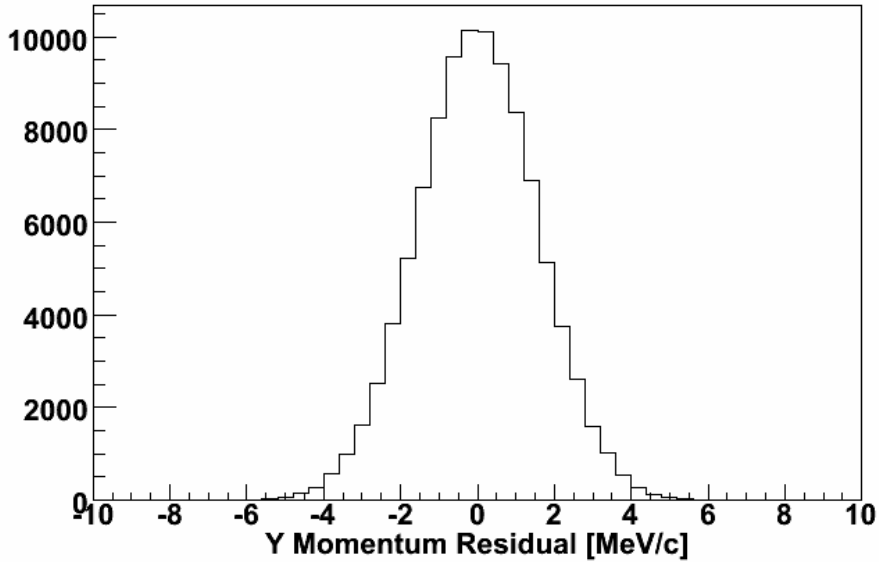

(*b*)

Figure 9.2: Transverse momentum residuals of the tracker.

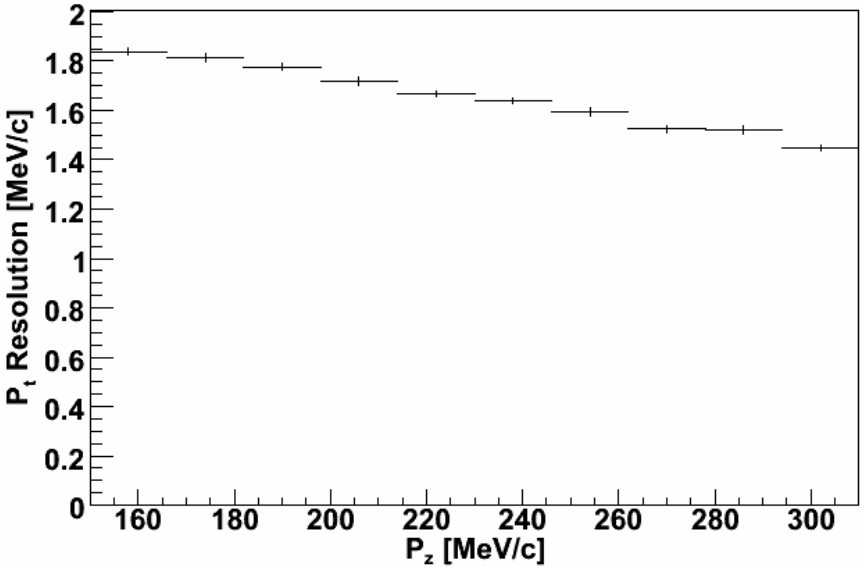

(*a*)

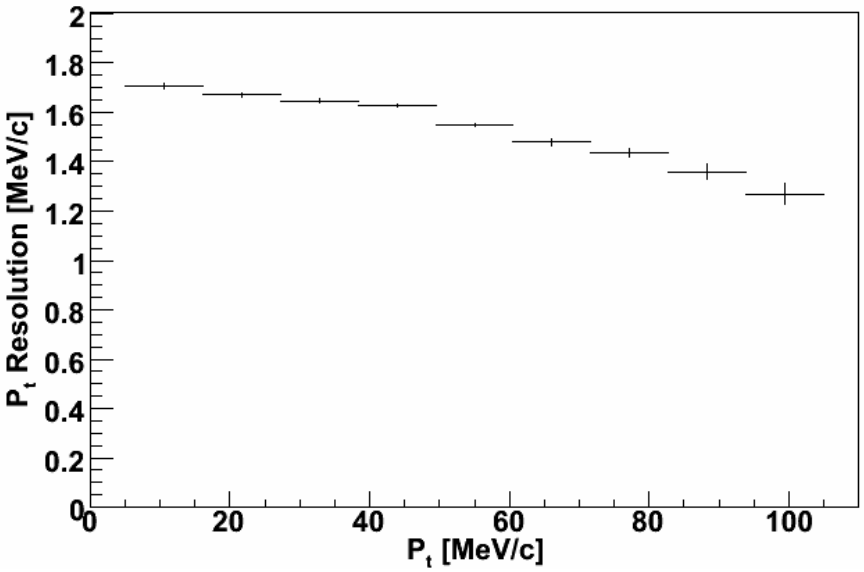

(*b*)

Figure 9.3: Correlation of p_t residuals with true p_z and p_t.

The longitudinal momentum measurement is highly correlated to the longitudinal and transverse momentum of the tracks themselves. The longitudinal momentum is inversely proportional to the angle, in cylindrical polar coordinates, through which a muon travels in the tracker. The error on the measurement of this angle is dependent on its size, so tracks with high momentum, that cut out a smaller angle, will be measured with a worse momentum resolution. The error on the measurement of the angle is also dependent on the radius of the helix. Measurement of the track momentum will be worse for muons that follow a tighter helix. The momentum resolution is approximately given by [75]

$$\sigma_{p_z} = \frac{k_2}{\Delta z}\frac{p_z^2}{p_t}\sqrt{2}\sigma_x \tag{9.4}$$

where Δz is the distance between tracker stations. In practice, several tracker station spacings are used and the orientation of the tracker results in an asymmetry between x and y. In Figure 9.4 the simulated tracker resolution is shown for p_z as a function of true p_t and true p_z, with errors given by a Gaussian smearing with mean 0 and RMS

$$\sigma_{p_z} = s_{pz} + k_{pz}\frac{p_z^2}{p_t}. \tag{9.5}$$

In Figure 9.5 E residuals are displayed for a beam with 6 mm emittance matched to the upstream tracker and a 10 MeV energy spread.

As will be shown, the correlations between the phase space variables and the measurement error have an important, detrimental, effect on the emittance measurement.

9.1.2 TOF Model

The Time of Flight (TOF) counters are expected to have an absolute resolution of 70 picoseconds and a possibly worse resolution relative to the RF phase. A reconstruction of the TOF counters is in preparation, but no such reconstruction has been performed to date. In this thesis the TOF residuals are assumed to have a gaussian distribution with a RMS width of 70 ps and mean 0, as shown in Figure 9.6.

9.1.3 Combined TOF and Tracker Simulation

In MICE, the momentum and position of particles is measured at the MICE TRPs, while the time of flight of muons is measured at the MICE TOFs. To calculate six dimensional emittance, the timing measurement must be extrapolated to the reference plane of the tracker.

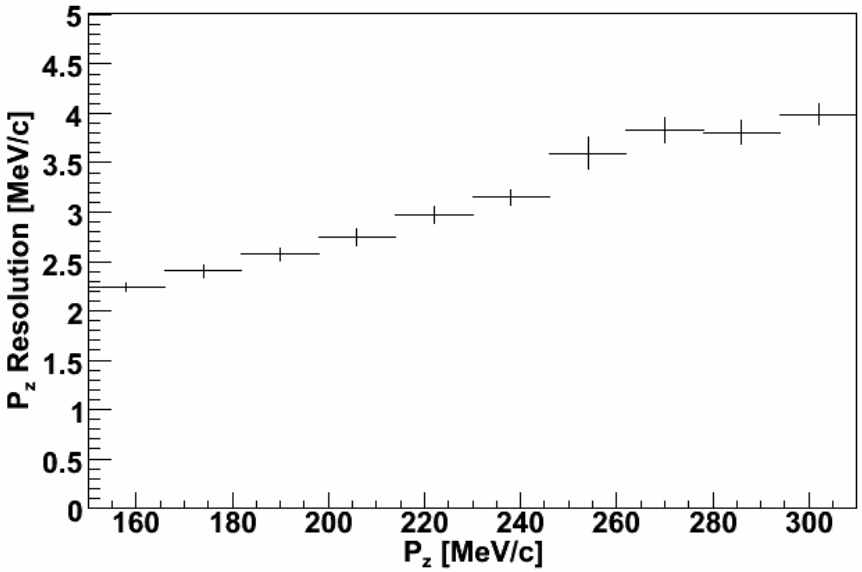

(*a*)

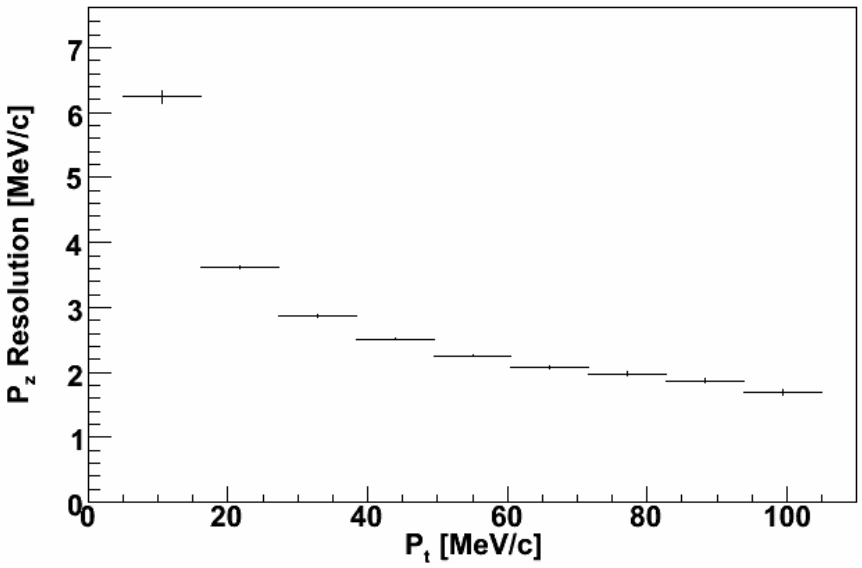

(*b*)

Figure 9.4: Correlations of p_z residuals with true p_z and true p_t.

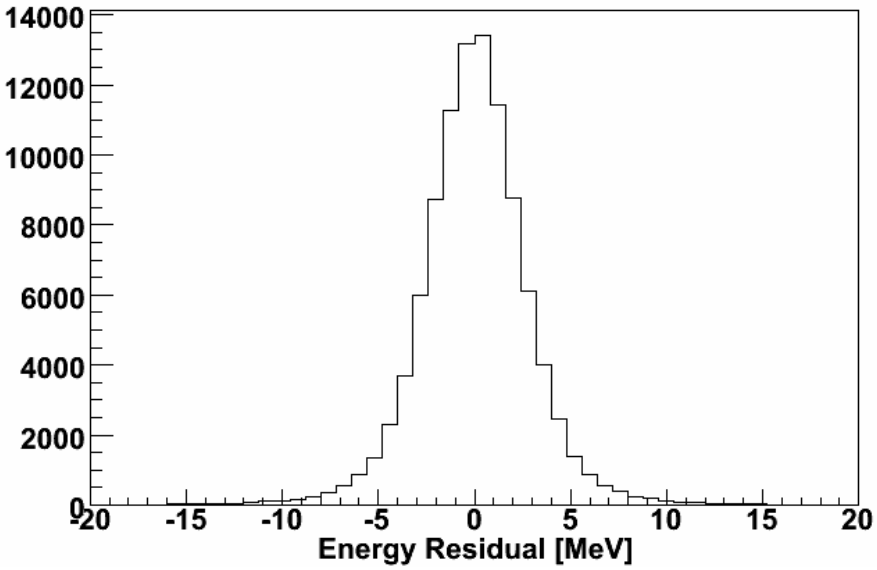

Figure 9.5: Simulated energy residuals of a MICE tracker.

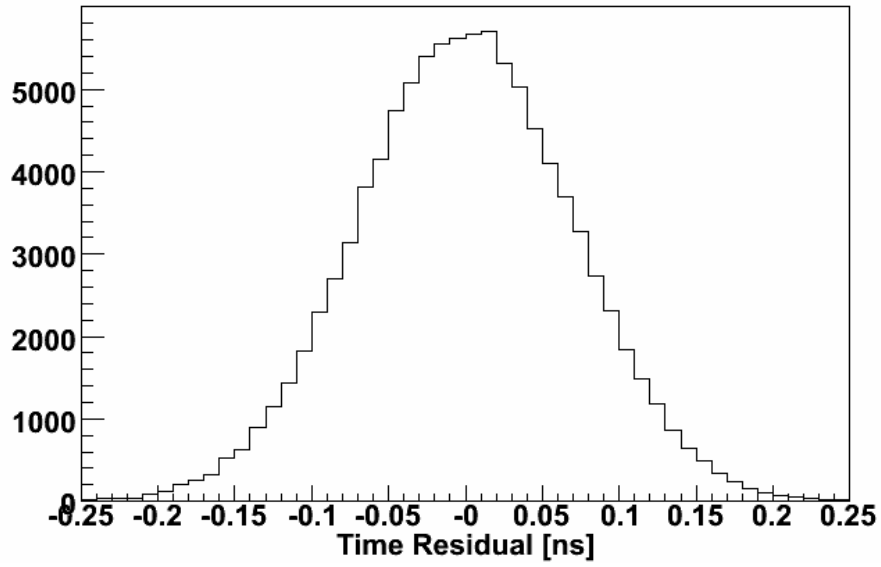

Figure 9.6: Simulated time measurement resolution of a MICE Time of Flight counter.

The accuracy of this extrapolation is effected by two factors. The finite resolution of the trackers introduces a correlation between the error on the energy and momenta and the error on the time measurement. Also, the energy straggling that occurs in material that lies between the tracker and the Time of Flight counter introduces an additional error.

In Figure 9.7 the two sources of error are compared and the combined effect is examined together with the intrinsic TOF resolution. Matched 6 mm Gaussian beams were transported from TOF 1 to the upstream TRP and from the downstream TRP to TOF 2. For this study, the TOFs and TRPs were taken to sit 6.611 m and 4.69 m respectively from the centre of MICE. A 7.6 mm lead diffuser was also present in the upstream tracker and the beam used was matched to account for its presence. The study was performed with material present and ideal detectors, without material and with both material and realistic detectors, in order to examine separately the effect of material, of tracker resolution and the two effects combined.

The time measurement was extrapolated from the TOF by tracking from the TRP to the TOF using the reconstructed momentum and position, in order to estimate the momentum and position of the track at the TOF; assigning the tracks with the measured value of the time; and then tracking back to the TRP. Tracking was achieved using numerical integration of the equations of motion with fourth order Runge Kutta and assuming no multiple scattering and mean energy loss in materials calculated according to the Bethe Bloch formula.

The dominant error arises due to the finite resolution of the TOF; at the upstream tracker the presence of material also makes a significant contribution to the time resolution. The tracker resolution has a negligible effect. The RMS errors introduced are listed in Table 9.2.

Note that there is also a mean error on the time measurement at the upstream TRP of 24 ps. Muons that suffer a large multiple scatter in the lead are measured with a reduced p_z in the tracker due to the higher transverse momentum. Hence, the beam has a higher mean p_z before the diffuser than after, which is not simulated.

9.2 Emittance Measurement Shift

In a particle by particle experiment such as MICE, a systematic shift in the measured emittance arises from statistical errors in the phase space variables. This is because MICE is measuring the width of a distribution so that statistical errors tend to add a systematic offset to the beam width measurement analogous

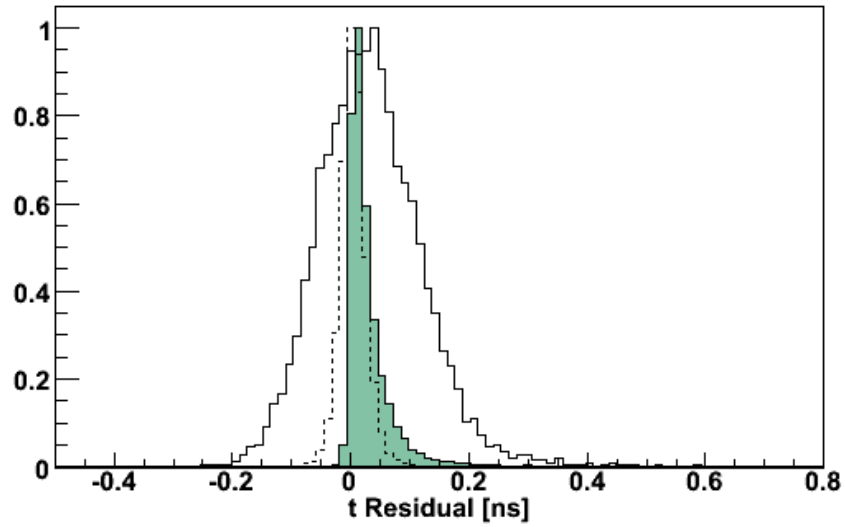

(*a*)

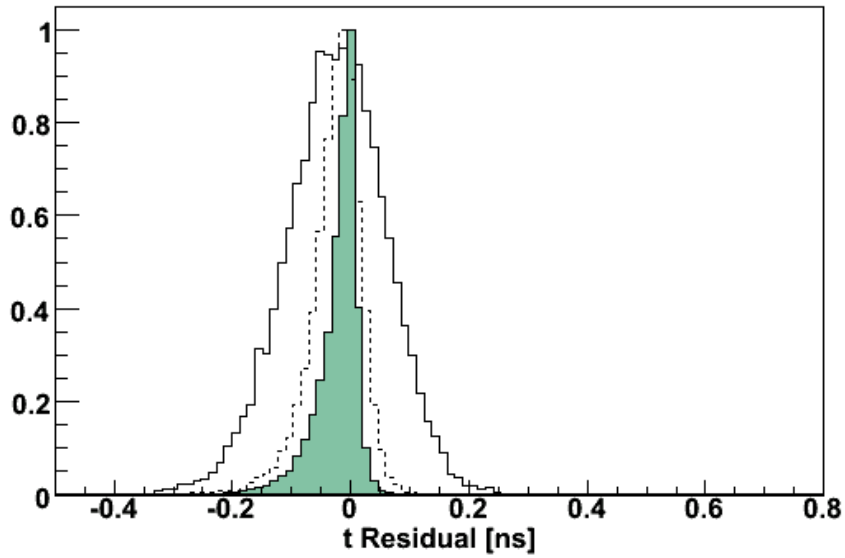

(*b*)

Figure 9.7: Simulated time resolution at the upstream (a) and downstream (b) MICE TRPs. The filled histogram shows the contribution due to the resolution of the trackers; the dashed histogram shows the contribution due to stochastic processes in materials; and the full, unfilled histogram shows the two effects combined with a TOF measurement smeared by a Gaussian distribution.

	$\sigma(t)$ (Downstream) [ps]	$\sigma(t)$ (Upstream) [ps]
Tracker	24	26
Material	12	50
TOF	70	70
Tracker, material and TOF	77	90

Table 9.2: Simulated RMS errors on time resolution at the TRP due to various effects.

to addition of errors in quadrature. Statistical errors arise due to fluctuations in the actual stochastic processes in the beam, the choice of input beam itself and fluctuations in the systematic shift in the measurement. Fluctuations in the cooling performance of the channel were discussed in the previous chapter.

9.2.1 Systematic Error

It is possible to estimate the size of this systematic shift and remove it, assuming a well calibrated detector.

In the case of a one dimensional distribution where errors δu_i on the measurement of a variable, u_i^m, are uncorrelated to the true value of the variable, u_i^t, the error on the variance is given by the familiar rule of addition in quadrature, [76]

$$< (u_i^m)^2 >=< (u_i^t)^2 > + < \delta u_i^2 > . \tag{9.6}$$

In the multidimensional case, the error on the covariance of phase space variables u_i and u_j becomes [76]

$$< u_i^m u_j^m >=< u_i^t u_j^t > + < \delta u_i \delta u_j > . \tag{9.7}$$

Note that here the measurement error is again assumed to be independent of the position of the muon in phase space. It is observed that the measured covariance is always larger than the actual covariance. The measured probability distribution of muons is the convolution of the true distribution and the distribution of the errors and such a smearing will always increase the measured emittance over the true emittance.

However, as demonstrated above, the phase space measurement errors are not, in general, independent of the position in phase space of the tracks. In this case the simple relationship above is not true. For example, in the one dimensional case, if there were a negative correlation between the true muon position and the measurement error, muons with negative positions would tend to have a positive

error; and muons with positive positions would tend to have a negative error. Then, overall, the width of the measured distribution would be smaller than the width of the actual distribution.

In general, the measured phase space variable of a particle is related to the true phase space variable by $u_i^m = u_i^t + \delta u_i$. Then the covariance of phase space variables for some ensemble of n_μ particles can be written as

$$\begin{aligned} < u_i^m u_j^m > - < u_i^m >< u_j^m > &= < (u_i^t + \delta u_i)(u_j^t + \delta u_j) > - \\ & \quad < u_i^t + \delta u_i >< u_j^t + \delta u_j > \end{aligned} \tag{9.8}$$

where the standard notation for raw moments in a discrete probability distribution,

$$< x >= \frac{1}{n_\mu} \sum_{all\mu}(x), \tag{9.9}$$

is used. Expanding the brackets,

$$\begin{aligned} < u_i^m u_j^m > - < u_i^m >< u_j^m > &= < u_i^t u_j^t > - < u_i^t >< u_j^t > \\ & \quad + < u_i^t \delta u_j > - < u_i^t >< \delta u_j > \\ & \quad + < u_j^t \delta u_i > - < u_j^t >< \delta u_i > \\ & \quad + < \delta u_i \delta u_j > - < \delta u_i >< \delta u_j > \end{aligned} \tag{9.10}$$

This can be written in matrix form,

$$\mathbf{V}^m = \mathbf{V}^t + \mathbf{R} + \mathbf{R}^\mathbf{T} + \mathbf{C}, \tag{9.11}$$

where $\mathbf{V}^t$ is the true covariance, $\mathbf{V}^m$ is the measured covariance and $\mathbf{R}$ and $\mathbf{C}$ are correction terms.

It is assumed subsequently that $\mathbf{R}$ and $\mathbf{C}$ are small. Then the true distribution will be close to the measured distribution, such that values of $\mathbf{R}$ and $\mathbf{C}$ can be determined from a Monte Carlo study using the covariance matrix of the measured distribution. This is an appropriate approximation; the R_{ij}/V_{ij} and C_{ij}/V_{ij} are at most of order 1 % for the detector model outlined above.

9.2.2 Transverse Emittance Shift

Transverse emittance shifts calculated using a full detector simulation are shown in Figure 9.8; in this study, detector reconstruction code was written by Malcolm Ellis, while analysis code and the systematic error removal technique was written by the current author. Over a range of emittances, 20 samples of 50000 events were simulated through MICE with p_z 200 MeV/c. The difference between true and measured emittance is shown for the upstream and downstream trackers.

At transverse emittances of 6 mm, a systematic offset of +0.3 % is observed. At lower emittances, correlations between phase space coordinates and detector errors are dominant leading to a negative systematic offset. Larger transverse momentum errors are observed for muons at higher radii and those with larger transverse momentum in the upstream tracker. It is thought that this may be due to multiple scattering, which may have had a stronger effect on muons that have large radii and large transverse momenta at the TRP. Further optimisation of reconstruction software may reduce these correlations. However, the correction technique outlined above is seen to remove the systematic offset.

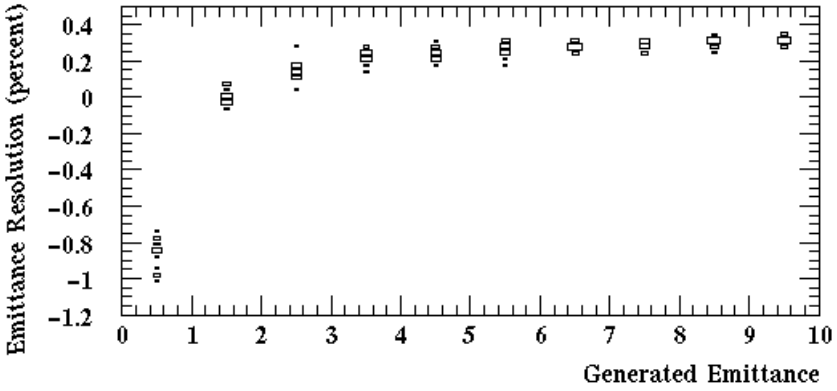

(*a*)

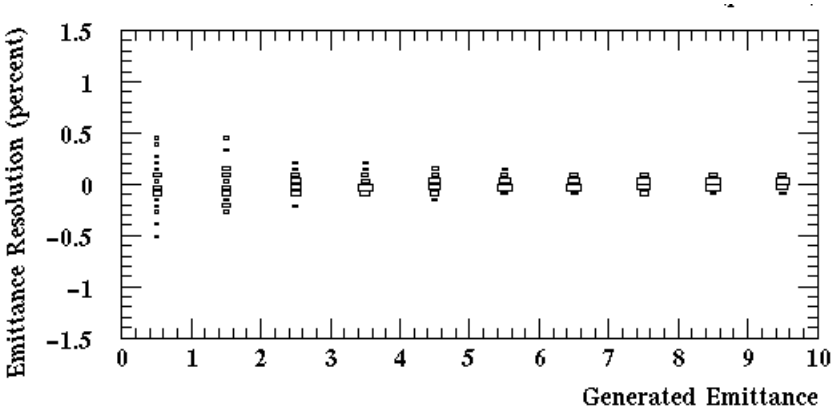

(*b*)

Figure 9.8: Emittance residuals for the upstream tracker (a) before correction (b) after correction for a monochromatic beam with various transverse emittances.

9.2.3 Longitudinal Emittance Shift

The emittance residuals for an ensemble of 100 samples, each containing 1000 events, is shown in Figure 9.9 calculated for longitudinal and 6D phase space,

before and after applying the correction procedure outlined above. The residuals are calculated for the worst case of the upstream tracker, where the lead diffuser makes the longitudinal emittance resolution worse as outlined above. In this case, a matched bunch with transverse emittance of 6 mm and longitudinal emittance of 0.25 ns/c was tracked from the TOF to the TRP and the phase space measurement was simulated as outlined above. The measured bunch emittance was calculated for each sample from the bunch and compared with the true emittance.

The shift was corrected by calculating the matrices **R** and **C**using the full sample of 100000 muons and an approximation to the true covariance matrix was calculated for each sample using the formula derived above. The corrected residuals are shown in 9.9 and the systematic shift on the emittance measurement is seen to be removed for longitudinal and six dimensional phase spaces. In Table 9.3 the standard deviation and mean of the emittance residual is given for both the uncorrected and corrected cases for transverse, longitudinal and six dimensional emittance at the upstream tracker.

	$\sigma(\delta\epsilon_\perp)$	$<\delta\epsilon_\perp>$	$\sigma(\delta\epsilon_{//})$	$<\delta\epsilon_{//}>$	$\sigma(\delta\epsilon_{6D})$	$<\delta\epsilon_{6D}>$
Uncorrected	0.0147	0.0261	0.00205	0.00528	0.0101	0.0222
Corrected	0.0148	-0.0008	0.00213	-0.00029	0.0106	-0.00043

Table 9.3: Mean and RMS of the emittance residual after simulated measurement of the emittance of 100 bunches of 1000 muons in the upstream tracker for a transverse emittance of 6 mm and longitudinal emittance of 0.1 ns before and after the correction.

9.2.4 Statistical Error Magnitude

The statistical errors on emittance measurement, listed in Table 9.3, can be compared with the intrinsic statistical variation in the emittance change, listed in Table 8.1. The two tables list errors on two different quantities. In the former case the statistical and mean measurement errors on measurement of emittance in the upstream tracker is listed. In the latter case the statistical and mean measurement errors on emittance change in the two trackers is listed. The statistical error on emittance change is dominated in transverse phase space by the inherent statistics of the cooling channel, even when studying the same sample of muons. In longitudinal phase space, the statistical error on emittance change is

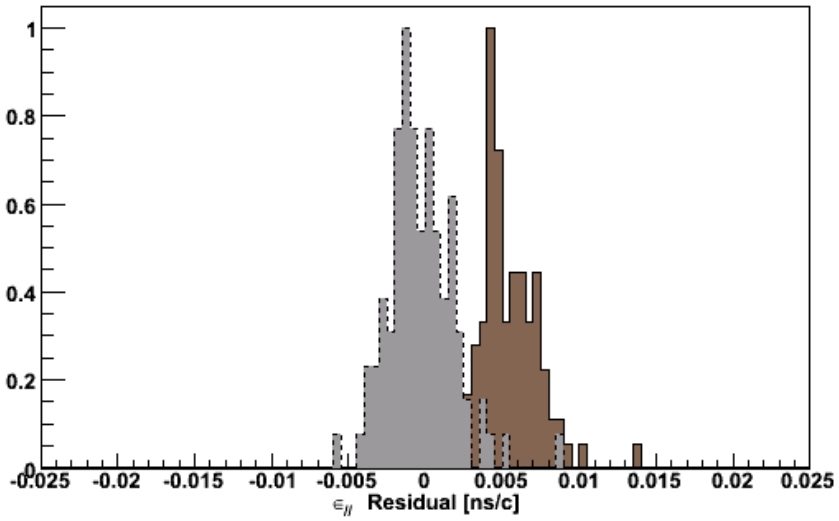

(*a*)

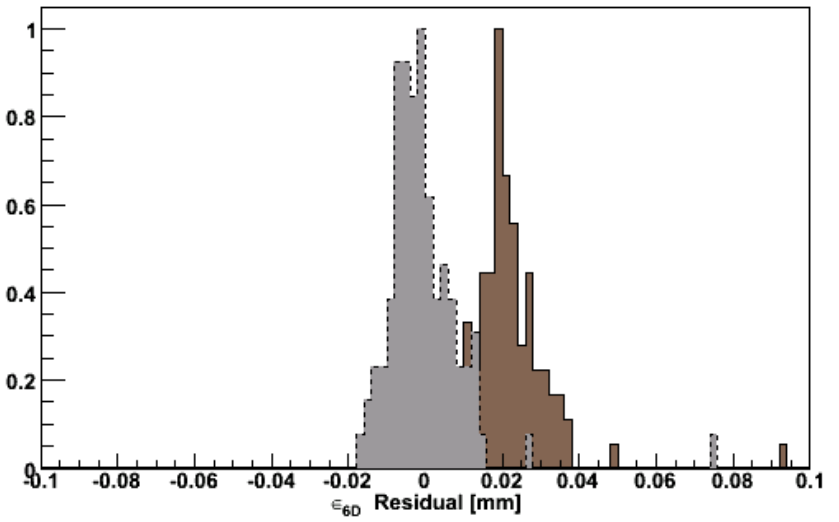

(*b*)

Figure 9.9: Emittance residuals for (a) longitudinal phase space (b) six dimensional phase space for a transverse emittance of 6 mm and longitudinal emittance of 0.1 ns. In each case the full line with brown fill is the uncorrected emittance residual while the dashed line with grey fill is the corrected emittance residual.

roughly the same for the two cases, and both the statistical error on the emittance measurement and processes in the cooling channel are expected to contribute.

9.2.5 Measured Cooling Performance

The aim of MICE is to measure the change in transverse emittance with a resolution of 1 %. Two requirements must be fulfilled to achieve this: firstly, sufficient events must be collected so that the statistical errors are smaller than 0.1 %; and secondly the systematic shift must be removed. A technique has been outlined for removing the systematic offset, depending on a well calibrated detector system. MICE intends to sample of order 10^6 muons per run in order to reduce the statistical errors below the required level.

Chapter 10

Conclusions

The Neutrino Factory facility is designed to enable the measurement of neutrino oscillations to an unprecedented precision and cooling is an important part. Muons are produced as tertiary particles and as such have a large beam emittance which can only be reduced by ionisation cooling. This enables a higher flux of muons to be accelerated, resulting in a larger neutrino flux which leads to better precision in measurement of the oscillation parameters.

The baseline design for muon cooling in a Neutrino Factory calls for ionisation cooling in a solenoidal focussing channel. In order to reduce the number of solenoids required and reduce decay losses, the solenoids, RF cavities and absorbers must be placed in a very short length, resulting in a very challenging accelerating structure to construct. Ionisation cooling has never been observed before and so MICE has been proposed as a Proof-of-Principle experiment to demonstrate that such a cooling channel is possible. In MICE, the cooling channel will be constructed, the position of muons in phase space will be measured before and after the channel and the beam emittance will be measured and compared with simulation.

In this thesis the simulation of the MICE cooling channel has been detailed. The accuracy of the solenoidal and RF model has been studied and the parameters necessary for simulation with an accuracy much greater than the resolution of the detectors have been found. In addition, the tracking parameters used in the GEANT4 library have been described and the accuracy of the tracking has been examined. It may be necessary to extend this work to include three dimensional field maps due to the possible presence of iron in the floor of the MICE hall. This will be determined by measurement of the fields when MICE is constructed. In addition, further work is required to understand the tracking accuracy in free space, or if necessary patch the GEANT4 simulation to allow tracking to arbitrary accuracy.

The tracking of a matched bunch through the MICE fields has been studied, including a comparison of G4MICE with linear beam optics under various different regimes. The first study of longitudinal beam dynamics in MICE has been performed and the resonance structure of the MICE lattice has been examined. A further study of the emittance change in MICE due to second and higher order terms would be interesting; in addition, the effect of resonances on the emittance change would be of interest.

The multiple scattering and energy loss model of GEANT4 has also been examined, and the accuracy of the representation of the absorbers has been studied. In particular, the effect on cooling of including spherical windows with varying thickness as compared to cylindrical windows has been examined and found to be small. The cooling performance of the full MICE channel has been studied including a brief investigation of the statistical errors. Further studies of statistical errors would be productive; in particular a Monte Carlo study to understand the statistics required for the measurement of cooling better than 1% would be of great interest.

The measurement of emittance in the MICE channel has also been studied. In particular, the error introduced by extrapolating the timing measurement from the TOF to the Tracker Reference Plane has been examined and the effect of detector resolution on emittance has been studied. It has been noted that statistical errors in the measurement of phase space variables introduces a systematic shift in the emittance measurement and a correction algorithm has been suggested assuming good calibration of the detectors. The analysis presented here assumed that the beam had been sampled from the available events correctly, but such a sampling procedure may in itself introduce errors that would be of interest to understand. In addition, it would be desirable to repeat this procedure with a more realistic detector reconstruction. Finally, the inclusion of errors from misidentification of impurities from beamline pions upstream and decay electrons downstream would also be of interest.

Here the simulation of the MICE experiment has been examined and various aspects studied in detail. The real test will come from the construction of a real cooling channel. The demonstration of ionisation cooling will be not only important to the Neutrino Factory collaboration, but also vital to the muon collider community and of concern to the muon physics community. At its most basic level, the MICE collaboration seeks to understand how to produce muons, a form of matter known for only 80 years, in unprecedented quantities; it is hoped that this will have benefits beyond the physics community and in the wider world.

Bibliography

[1] Ed. S. Geer and M. Zisman, Neutrino Factory and Beta Beam Experiments and Development (aka Feasibility Study IIa), Technical Report BNL-72369-2004, 2004.

[2] Y. Ashie et al., Evidence for Oscillatory Signature in Atmospheric Neutrino Oscillation, Phys. Rev. Lett. 93, 101801, 2004.

[3] S. Ahmed et al., Measurement of the Total Active ^{8}B Solar Neutrino Flux at the Sudbury Neutrino Observatory with Enhanced Neutral Current Sensitivity, Phys. Rev. Lett. 92, 181301, 2004.

[4] Measurement of Neutrino Oscillation with KAMland: Evidence for Spectral Distortion, The KamLAND Collaboration, Phys. Rev. Lett. 94, 081801 (2005).

[5] K. Elsener, CNGS - CERN Neutrinos to Gran Sasso, Proc. Les Recontres de Physique de la Vallee d'Aoste, 2000.

[6] MINOS Collaboration, Neutrino Oscillation Physics at Fermilab: the NuMi-MINOS Project, Fermilab report NuMI-L-375, 1998.

[7] MINERvA Collaboration, The Physics Case and Technology of the MINERvA Experiment, MINERvA Document 218-v4, available from http://minerva.fnal.gov, 2004

[8] K2K collaboration, Measurement of Neutrino Oscillation by the K2K Experiment, Phys. Rev. D 74, 072003, 2006.

[9] T2K Collaboration, Neutrino Oscillation Experiment at JHF, Letter of Intent, 2003.

[10] Particle Data Group, Review of Particle Physics, J. Phys. G 33, 1, 2006.

[11] Ed. A. Baldini et al., Beams for European Neutrino Experiments (BENE): Midterm Scientific Report, CERN 2006-005, ISBN 92-9083-273-8, 2006.

[12] Ed. S. Ozaki, R. Palmer, M. Zisman, and J. Gallardo, Feasibility Study-II of a Muon-Based Neutrino Source, BNL-52623, 2001.

[13] J.S. Berg et al., A Cost-Effective Design for a Neutrino Factory (aka Feasibility Study IIb), MUCOOL Note 327, 2005.

[14] NuFactJ Group, A Feasibility Study of a Neutrino Factory in Japan, available from http://www-prism.kek.jp/nufactj/, 2001.

[15] ISS Collaboration, International Scoping Study of a Neutrino Factory and Superbeam Facility, (to be published), Proc. 4th APAC, 2007.

[16] S. Machida, Non-Scaling FFAG Issues, Talk at KEK ISS Meeting, 2006.

[17] K. Long, RAL Summer School for Young High Energy Physicists, 2005.

[18] C.R. Prior: Proton Drivers for a Neutrino Factory, Proc. 1st Neutrino Factory School, 2002.

[19] Computer Pion Yields from a Tantalum Rod Target: Comparing MARS15 and GEANT4 across Proton Energies, Nuclear Physics B Proceedings Supplements, Volume 155, 295-296, 2005.

[20] R. Bennett et al., Studies of a Target System for a 4-MW, 24-GeV Proton Beam, 20th INTC Meeting, 2004.

[21] D. Neuffer, High Frequency Buncher and Phase Rotation for the Muon Source, MUCOOL Note 181, 2000.

[22] A. Moretti et al., Effects of high solenoidal magnetic fields on rf accelerating cavities, PRSTAB 8, 072001, 2005.

[23] MICE Collaboration, MICE Technical Reference Document, available from http://www.mice.iit.edu/, 2006.

[24] S. Berg, A. Garren, R. Palmer, G. Rees, D. Summers and Y. Zhao, A Pulsed Muon Synchrotron for a Neutrino Factory, MUCOOL Note 259, October, 2002.

[25] A. Bogacz and V.A. Lebedev, Recirculating Linac Muon Acceleration for Neutrino Factory, NIM A 503, 306-308, 2003

[26] E. Keil and A.M. Sessler, Muon Acceleration in FFAG Rings, NIM A 538, 159-177, 2005.

[27] R. Palmer, Optimization of Cooling vs Accelerator Acceptance, Talk at KEK ISS Meeting, 2006.

[28] G. Penn and J.S. Wurtele, Beam Envelop Equations for Cooling of Muons in Solenoid Fields, PRL 85, 764-767, 2000.

[29] C. Wang and K. Kim, Recursive Solution for Beam Dynamics of Periodic Focussing Channels, PRE 63, 056502. 2001.

[30] M. Berz and K. Makino, COSY Infinity User's Guide and Reference Manual, 2001.

[31] C.J. Johnstone, M. Berz, D. Errede, K. Makino, High-order map Treatment of Superimposed Cavities, Absorbers, and Magnetic Multipole and Solenoid Fields, NIM A 519, 162-174, 2004.

[32] A.J. Dragt, Numerical Third-Order Transfer Map for Solenoid, NIM A 298, 441-459, 1990.

[33] V.I. Arnold, Mathematical Methods of Classical Mechanics, Springer, ISBN: 0387968903, 1997.

[34] S. Y. Lee, Accelerator Physics, World Scientific, ISBN 981-02-3710-3, 1999.

[35] A.B. El-Kareh and J.C.El-Kareh, Electron Beams, Lenses and Optics 2, 20, 1970.

[36] H. Goldstein, Classical Mechanics, Addison Wesley, ISBN: 0321188977, 2001.

[37] K. Floetmann, Some Basic Features of Beam Emittance, PRSTAB 6, 034202, 2003.

[38] E.B. Holzer, Figure of Merit for Muon Cooling - An Algorithm for Particle Counting in Coupled Phase Space Planes, CERN NuFact Note 111, 2003.

[39] J. van Zeijts, Arbitrary Order Transfer Maps for RF Cavities, Proc. PAC 1995, 2356-2358, 1996

[40] D. Abell, Numerical Computation of High-Order Transfer Maps for RF Cavities, PRSTAB 9, 052001, 2006

[41] G. Arfken, Mathematical Methods for Physicists, 3rd ed. Orlando, Florida: Academic Press, 1985.

[42] R.C. Fernow, Classification of Periodic Alternating-Solenoid Lattices, MUCOOL Note 244, 2002.

[43] D. Neuffer, Principles and Application of Muon Cooling, Fermilab Note FN-378, 1983.

[44] B. Palmer et al., The RFoFo Ionization Cooling Ring for Muons, MUCOOL Note 314, 2005.

[45] V. Balbekov, Solenoid-Based Ring Coolers, MUCOOL Note 232, 2002.

[46] D. Summers,Muon Ionization Cooling for a 8 TeV Lepton Collider, 73rd Annual Meeting of the Southeastern Section of the APS, 2006.

[47] V. Elvira, P. Lebrun and P. Spentzouris, MUCOOL Note 193, Simulation of a Helical Cooling Channel Using GEANT4, 2001.

[48] P. Hanlet et al., High Pressure RF Cavities in Magnetic Fields, Proc. EPAC 2006, 2006.

[49] P. Drumm, Progress and Status of the MICE Project, Proc. EPAC 2006.

[50] J. Cobb, Private correspondence, 2006.

[51] G. Gregoire, π/μ Separation, Talk at the MICE Upstream PID Review, 2006.

[52] T. Roberts, More on Upstream Simulation, MICE Collaboration Meeting 9, 2004.

[53] R. Sandstrom, EMCal Design, MICE Collaboration Meeting 15, 2006.

[54] A. Khan et al., MICE Scintillating Fibre Tracker Prototype, MICE NOTE 90, 2004.

[55] S. Agostinelli et al., GEANT4 - A Simulation Toolkit, NIM A 506, 250-303, 2003.

[56] Unified Modelling Language Version 1.4.2, International Standards Organisation, ISO/IEC 19501:2005, 2005.

[57] R.E. Kalman, A New Approach to Linear Filtering and Prediction Problems, Transactions of the ASME - Journal of Basic Engineering Vol. 82, 35-45, 1960.

[58] T. Roberts, G4Beamline User's Guide, available from http://g4beamline.muonsinc.com, 2006 (development ongoing).

[59] K.L. Brown, Ch. Iselin, D.C. Carey, Decay Turtle, CERN 74-2, 1974.

[60] R.C. Fernow, ICOOL: A simulation code for ionization cooling of muon beams, Proc. 1999 Particle Accelerator Conference, 3020, 1999

[61] R.C. Fernow, Recent developments on the muon-facility design-code ICOOL, Proc. 2005 Particle Accelerator Conference, 2651, 2005.

[62] Raytrace users manual, S. Kowalski & H. Enge, 1986.

[63] R.C. Fernow, Modeling Solenoids Using Coil, Sheet and Block Conductors, MUCOOL Note 281, 2003.

[64] K. Halbach and R. F. Holsinger, SuperFish - A Computer Program for Evaluation of RF Cavities with Cylindrical Symmetry, Particle Accelerators 7, 213-222, 1976.

[65] GEANT4 Collaboration, GEANT4 User's Guide for Application Developers (v4.8.1), available from http://geant4.web.cern, 2006 (development ongoing).

[66] W. Press et al, Numerical Recipes in C, Cambridge University Press, 1987.

[67] L. Richardson, The deferred approach to the limit, Philosophical Transactions of the Royal Society of London, Series A 226: 299349, 1927.

[68] M. Galassi et al, GNU Scientific Library Reference Manual (2nd Ed.), ISBN 0954161734.

[69] J.H. Cobb, A Compendium of MICE Beta Functions, MICE Note 83, 2004.

[70] J. S. Berg, Amplitude Dependence of Time of Flight and its Connection to Chromaticity, arXiv:physics/0607275, submitted to NIM A, 2006.

[71] U. Fano, Inelastic collisions and the Moliere theory of multiple scattering, Phys. Rev. 93:117-120, 1954.

[72] R.C. Fernow, Scattering in ICOOL, MUCOOL Note 336, 2006.

[73] J. Cobb, unpublished note, 2005.

[74] A. Blondel and R. Edgecock, Time structure and useful muon rates for MICE at RAL, MICE Collaboration Meeting, 2002.

[75] T. Matsuhita, Tracker Workshop, 2006

[76] R.J. Barlow, Statistics: A Guide to the Use of Statistical Methods in the Physical Sciences, J. Wiley, ISBN: 0471922951, 1993.